

Analysis of Large-Scale Systems: Coupled Multispecies Gradient Flows and Distributed Control

Thesis by
Lauren Elaine Conger

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended May 18, 2026

© 2026

Lauren Elaine Conger
ORCID: 0000-0002-2151-6044

All rights reserved

In honor of Madison Pedrotty.

ACKNOWLEDGEMENTS

My time at Caltech has been wonderful because of all the people with whom I have had the privilege to work, study, and grow. From the first year, my cohort fostered friendships and mathematical foundations. I am especially grateful to Eitan and Matthieu for their guidance in analysis, and to my department colleagues and friends Laure, Chris, Nico, Ricardo and Tinashe. Many thanks to my officemate, teammate, and friend Margaret for being a constant source of inspiration and fortitude. John's group members and alumni — especially Shih-Hao, Jing, Lisa, James, and Carmen — provided support during the early days of my PhD and welcomed me into the world of research. Thanks to Minh, Kathrin, Sangmin, Dohyeon, Ritvik, Jinghao, Minakshi, Elizabeth, Jeremy and Eyyup for fun mathematical discussions, hikes, and dinners. Thank you to all the wonderful administrators who have supported me through the program, especially Monica, Sumaia, Nuvia, Minzhi, Syndey, Carmen, Maria, and Diana.

Outside of research, the formation and friendship of the Catholic community in Pasadena has been amazing. I am grateful to Alessandra, Emily, Brendan, Abby, Hayley, Jorge, Katt, Stephanie, Christopher, Anakarina, Jerome, Joe, Maria, Maria, Maria, Fr. Tony, Sr. Mariana, Sr. Juliana, Nick, Mairead, Jesús, Emely, and Ana for your friendship — and to Daniel and Krista especially for sharing with me the joy of liturgical music. Fr. Nicolas, Fr. Pablo, Fr. John: thank you for your mentorship and spiritual guidance. I am also grateful to the St. Philip young adult choir and Octech, with whom I have had the pleasure of singing. To my triathlon adventure crew — especially Ryan, Marianne, Annette, and Carina — for many fun runs, training sessions, and races. To the Caltech cross country and track and field teams and Kelli, for a new perspective on training.

Thanks to my gradient flow graduate student and postdoc colleagues I have met at various conferences and events - Anna, Alessandro, Charles, Djordje, Claire, Jethro, Fabius, Giacomo, Nico, Holger, Vas, Sixou, Guillaume, and Max, and many others. I have enjoyed learning from you all at our online reading groups and in person.

Thanks are owed to the texts and results that shaped my mathematical development: A Hitch-hiker's Guide to Infinite Dimensional Analysis, Gradient Flows in Metric Spaces (AGS), Kinetic Equilibration Rates, and Fatou's Lemma for Weakly Converging Measures under the Uniform Integrability Conditions. From these I learned the basics of gradient flow calculations and analysis, and mathematical storytelling.

Thank you to my thesis committee for your feedback and insights: Franca Hoffmann, Eric Mazumdar, Katy Craig, and Adam Wierman. Thanks to André Schlichting for his discrete gradient flow insights and mathematical mentorship alongside José Carrillo, Oliver Tse, and Giulia Cavagnari; to my

collaborators Antoine Leeman and Lily Ratliff; to Giuseppe Savaré for generously sharing his expertise and mathematical guidance, to Steven Brunton for encouraging me to pursue academia in my early PhD years and for teaching me how to review papers; to Joel Tropp for encouraging me to pursue mathematics; and to Richard Murray for convincing me to come to Caltech and supporting me through the qualifying exam process. To Adam Wierman: thank you for years of thoughtful mentorship on academic careers, for your generous investment in students, and for teaching me to motivate my work and communicate it well.

To my advisors: John, thank you for teaching me how to pose the simplest problem that illuminates an important phenomenon. Thank you for endlessly generating ideas, for many entertaining discussions and debates, and for recognizing that I was capable of thriving in graduate school. Eric, whose remarkable intuition and instinct for interesting problems has taught me how to think at a high level and build mathematical results from the top down: thank you for introducing me to game theory and encouraging me to pursue my interests. Thank you for your excellent example of how to give engaging talks, for helping me find my strengths in research, and for teaching me how to write and pursue well-motivated work. Franca, you enabled my dream of becoming a mathematician: thank you for patiently introducing me to gradient flow analysis, for your example of bridging mathematical tools across fields, for your excitement about my research, for seeing my strengths and reminding me of them during times of doubt. Thank you for many wonderful adventures together, mathematical and otherwise — hiking Mt. Baldy, ice skating, performing the Geographical Fugue, among others. Thank you for teaching me to be rigorous, for our shared appreciation of the beauty of mathematics, for your global vision of the world, and for fostering a welcoming group with many shared meals. It is an honor to be Franca and Eric's first PhD student to graduate; I look forward to following their academic journeys for the entirety of their faculty careers.

In gratitude, finally, to Alessandra, for her support and being a wonderful friend and roommate, to Peter, my siblings Hailey, David, Sarah & Cole, Claire, and my parents Jody and Diane — for their unconditional love and support.

*Grant me, O Lord my God, a mind to know you, a heart to seek you,
wisdom to find you, conduct pleasing to you, faithful perseverance in
waiting for you, and a hope of finally embracing you. Amen.*

- St. Thomas Aquinas

ABSTRACT

Large-scale multi-agent systems — from autonomous vehicle fleets and power grids to biological populations and machine learning algorithms — have grown too vast and heterogeneous for individual-level modeling or centralized control. This thesis develops new mathematical tools from two complementary perspectives: one that aggregates agents into evolving probability distributions and analyzes their collective behavior, and one that designs distributed controllers acting directly on individual components of large dynamical systems.

In the first approach, we treat the joint distribution of n agent populations as evolving to minimize a collection of coupled cost functionals, yielding a coupled system of partial differential equations. The natural analytical framework is that of gradient flows in metric spaces — the infinite-dimensional analogue of gradient descent. While gradient flow theory is well-established for a single species, the multispecies setting presents new difficulties: the joint dynamics do not necessarily have a gradient flow structure, the natural candidate for a steady state is a Nash equilibrium rather than a minimizer, and strong displacement convexity of each individual energy is generally insufficient for convergence. We introduce λ -monotonicity for n -species systems in the Wasserstein-2 metric space, extending the game-theoretic notion of monotonicity from Euclidean space to the measure-valued setting. Under this condition, we prove exponential convergence to a unique steady state — which, when the dynamics arise from coupled gradient flows, is the unique Nash equilibrium of the associated n -player game. For existence of solutions in general metric spaces, we introduce the Variational Movement Scheme (VMS), a fully implicit discrete-time update whose solution is a Nash equilibrium at each step, and prove existence via a zeroth-order monotonicity condition combining barycentric κ -convexity and η -interaction dissipativity. Taking the discrete time step to zero yields a continuous evolution variational inequality and a contraction estimate at rate $\lambda = \kappa - \eta$. These results establish foundational multispecies gradient flow theory: existence of solutions and long-time behavior.

In the second approach, we shift from a global to local perspective to design controllers for individual components of large-scale dynamical systems. Building on the System Level Synthesis (SLS) framework — which reparameterizes the controller via closed-loop maps (CLMs) from disturbances to states and inputs — we develop tools for distributed, globally optimal constrained control. First, we characterize controllability and observability under structural constraints (spatial locality, communication delays, actuation limits) through a constrained Gramian and controllability volume. We establish explicit rank conditions under which structural constraints incur no performance loss: locality and communication constraints can satisfy these conditions, while actuation delays monotonically reduce the reachable volume. Second, when system dynamics are unknown, we establish the first non-asymptotic convergence rate for set membership estimation (SME) under

general convex disturbance supports, proving diameter convergence with rate $\tilde{O}(1/T)$, which improves on the $\tilde{O}(1/\sqrt{T})$ rate of least squares estimation in sample complexity. Third, we extend SLS to infinite-dimensional evolution equations over Hilbert spaces, establishing weak-form CLMs, an operator-theoretic system level parameterization, and an equivalence between CLMs and linear feedback controllers — enabling an optimize-then-discretize paradigm for PDE control that avoids the state-space explosion of the standard discretize-then-optimize approach. We include a biology-inspired example that shows how our method can provide better control than spatial discretization. These results build a toolset for implementing distributed, local controllers with rigorous stability guarantees.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] L. Conger, F. Hoffmann, and G. Savaré. Dissipative variational evolutions in metric spaces, 2026.
- [2] L. Conger, F. Hoffmann, R. Baptista, and E. Mazumdar. Computing Optimal Transport Plans via Min-Max Gradient Flows. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 2210–2217, Dec. 2025. doi: 10.1109/CDC57313.2025.11312926. ISSN: 2576-2370.
- [3] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Monotone Multispecies Flows, June 2025. doi: 10.48550/arXiv.2506.22947. arXiv:2506.22947.
- [4] L. Conger, A. P. Leeman, and F. Hoffmann. Convex Constrained Controller Synthesis for Evolution Equations. In *2025 American Control Conference (ACC)*, pages 1169–1176, July 2025. doi: 10.23919/ACC63710.2025.11108072. ISSN: 2378-5861.
- [5] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games, Nov. 2024. doi: 10.48550/arXiv.2411.07403. arXiv:2411.07403 [math].
- [6] L. Conger, Y. Lin, A. Wierman, and E. Mazumdar. Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pages 8491–8497, Dec. 2024. doi: 10.1109/CDC56724.2024.10886635. ISSN: 2576-2370.
- [7] Y. Li, J. Yu, L. Conger, T. Kargin, and A. Wierman. Learning the Uncertainty Sets of Linear Control Systems via Set Membership: A Non-asymptotic Analysis. en. In *Proceedings of the 41st International Conference on Machine Learning*, pages 29234–29265. PMLR, July 2024.
- [8] L. Conger, F. Hoffmann, E. Mazumdar, and L. Ratliff. Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows. en. *Advances in Neural Information Processing Systems*, 36:45971–46006, Dec. 2023.

In [1], Giuseppe Savaré shared notes outlining the zero-sum setting for proving existence of solutions. I extended the argument to the general multispecies setting, leaving open existence of the Nash equilibrium update. Giuseppe proposed a solution and generalization to dissipative operators in metric space. I fleshed out the proofs and results, added examples, and wrote the rest of the paper. Franca and Giuseppe provided comments, ideas, and writing feedback.

In [2], after Eric Mazumdar proposed formulating the Kantorovich optimal transport problem as a min-max gradient flow, I determined which results to prove and completed approximately 75% of the lemmas and theorems, including the main result. I developed and implemented the numerical

method proposed in the paper, and did the majority of the writing, with collaborative input from Franca Hoffmann and Ricardo Baptista on the introduction and the remaining 25% of the proofs.

In [3], I developed the notion of monotonicity used in this paper, completed all proofs and code, and wrote the initial draft. Franca and Eric provided suggestions for revision and editing.

In [4], I proposed the idea of formulating system level synthesis for PDEs to Franca, identified what needed to be proved, and developed the initial proofs. Franca assisted in determining the domains of all operators and refining the results. I wrote the initial draft of the paper and completed all coding for the new method, with writing contributions from my coauthors thereafter.

In [5], I determined which results to prove and completed the proofs in all sections except Section 7 of the appendix, which I completed jointly with Franca. I wrote the paper and developed all code, with revisions based on feedback from the other authors.

In [6], I drove the high-level direction of the paper and jointly proved the results with Yiheng Lin. I wrote the majority of the paper and conducted all numerical simulations.

In [7], Yingying Li and Jing Yu derived the main results. I contributed to the numerical simulations and to verifying the proofs.

In [8], I assisted Franca in completing the proofs, wrote the majority of the paper, and conducted all numerical simulations. Lily Ratliff and Eric provided modeling insights and problem formulation ideas.

TABLE OF CONTENTS

Acknowledgements	iv
Abstract	vi
Published Content and Contributions	viii
Table of Contents	ix
List of Illustrations	xv
List of Tables	xx
Chapter I: Introduction	1
1.1 Part I: Multispecies Gradient Flow Analysis	3
1.1.1 Gradient Flow Introduction	3
1.1.2 Well-Posedness and Asymptotic Behavior of a Single Species Gradient Flow	7
1.1.3 Wasserstein-2 Metric	12
1.1.4 Results: Asymptotics of Coupled Wasserstein-2 Multispecies Evolutions	17
1.1.5 Results: Well-Posedness of Coupled Multispecies Evolutions	22
1.1.6 Results: Gradient Flow Applications	28
1.2 Part II: Distributed Control	39
1.2.1 System Level Synthesis (SLS) Parameterization	41
1.2.2 Results: Controllability and Observability	49
1.2.3 Results: Learning Dynamics from Data via Set Membership	52
1.2.4 Results: Control of Linear Evolution Equations	54
INTRODUCTION BIBLIOGRAPHY	59

I Multispecies Gradient Flow Analysis 65

Chapter II: Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games	66
Abstract	67
2.1 Introduction	68
2.1.1 Related Literature	70
2.1.2 Contributions	72
2.1.3 Paper Structure	73
2.2 Preliminaries	73
2.3 Main Results	77
2.3.1 Assumptions	77
2.3.2 The Cooperative Setting	78
2.3.3 The Competitive Setting	79
2.4 Application: Distribution Shift in Machine Learning	80
2.4.1 Assumptions under Timescale Separation	82
2.4.2 Analysis of Competitive Objectives with Timescale Separation	83
2.5 Insights from Numerical Experiments	84
2.5.1 Census Data in Colombia	84

2.5.2	Loan Applications: Feature Modification	86
2.5.3	Performative Prediction	87
2.6	Cooperative Setting (Proof of Theorem 2.13)	89
2.6.1	Ground States and Steady States	89
2.6.2	Functional Inequalities	94
2.7	Competitive Setting (Proof of Theorem 2.14)	97
2.A	Competitive Objective, Fast Algorithm (Proof of Theorem 2.19)	106
2.B	Competitive Objective, Fast Population (Proof of Theorem 2.20)	116
2.C	Auxiliary Lemmas	127
Chapter III: Monotone Multispecies Flows		134
Abstract		135
3.1	Introduction	136
3.1.1	Contributions	140
3.1.2	Related Literature	141
3.1.3	Structure	143
3.2	Preliminaries	143
3.3	Main Results	146
3.3.1	Second Moments	149
3.3.2	Proof of main theorem	153
3.4	Coupled Gradient Flow Setting	156
3.4.1	Example: Two Species Zero-Sum Game with Nonlinear Diffusion	158
3.4.2	Monotonicity Implies Convexity	160
3.4.3	Additivity of Monotonicity	161
3.4.4	Monotonicity of Cross-Interaction Kernels	163
3.5	Applications	166
3.5.1	Multi-Learner System	166
3.5.2	Team Mean Field Games	169
3.5.3	Degenerately Monotone Game with Diffusion	170
3.A	Auxiliary Definitions and Results	172
3.A.1	Second order monotonicity condition	173
3.A.2	Monotonicity vs Convexity and Measures vs Vectors	175
3.B	Interaction Kernel Computations	177
3.C	Optimal and Admissible Plans in the Monotonicity Inequality	182
Chapter IV: Dissipative Variational Evolutions in Metric Spaces		186
Abstract		187
4.1	Introduction	188
4.1.1	Contributions	196
4.1.2	Related Literature	197
4.1.3	Plan of the paper	200
4.2	Preliminaries: metric spaces, slopes, and convexlike functions	202
4.2.1	Barycentric Convexity and Related Notions	202
4.2.2	Metric Derivatives	205
4.2.3	Inequalities with upper and lower Dini derivatives	206
4.2.4	Examples	207
4.3	Equilibrium problems	210

4.3.1	Saddle points of antysymmetric bifunctions	215
4.3.2	Examples	217
4.4	Metric λ -dissipative evolutions	222
4.4.1	Driving bifunctions and EVI flows.	222
4.4.2	Equivalent formulations	225
4.4.3	Discrete b-EVI	228
4.5	Solution to the Discrete b-EVI	231
4.5.1	Examples	235
4.6	Multispecies Systems	239
4.6.1	Euclidean Multispecies Gradient Flows	239
4.6.2	Multispecies Coupled Wasserstein-2 Gradient Flows	240
4.A	Proof of Cauchy Estimate	244
Chapter V: Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows . .		249
Abstract		250
5.1	Introduction	251
5.2	Problem Formulation	253
5.2.1	Case (i): Aligned Objectives	255
5.2.2	Case (ii): Competing Objectives	255
5.3	Results	255
5.3.1	Gradient Flow Structure	256
5.3.2	Analysis of Case (i): Aligned Objectives	259
5.3.3	Analysis of Case (ii): Competing Objectives	260
5.4	Numerical Examples	262
5.4.1	Competitive Objectives	262
5.4.2	Naïve Behavior	263
5.5	Future Directions, Limitations, and Broader Impact	263
5.A	General structure and preliminaries	265
5.A.1	Structure of the dynamics	265
5.A.2	Definitions and notation	268
5.A.3	Steady states	270
5.B	Proof of Theorem 5.9	272
5.C	Proof of Theorem 5.11	279
5.D	Proof of Theorem 5.12	289
5.E	Additional Simulation Results	294
5.E.1	Aligned Objectives	294
5.E.2	Multiple Dynamical Populations	295
5.E.3	Sampled Gradients	296
5.E.4	Two-dimensional Distributions	297
Chapter VI: Computing Optimal Transport Plans via Min-Max Gradient Flows		299
Abstract		300
6.1	Introduction	301
6.2	Problem Setup	303
6.2.1	Notation	303
6.2.2	Displacement Convexity	303
6.2.3	Optimal Transport Problem	304

6.3 Results	305
6.3.1 Displacement Convexity of Energy	306
6.3.2 Danskin's Theorem	307
6.3.3 Long-Time Behavior	311
6.4 Numeric Implementation	313
PART I BIBLIOGRAPHY	318

II Distributed Control 334

Chapter VII: Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints	335
Abstract	336
7.1 Introduction	337
7.2 Motivation & Background	338
7.2.1 Classical Controllability and Observability	339
7.2.2 System Level Synthesis (SLS)	340
7.3 Constrained Systems	340
7.4 Results	342
7.4.1 Connection to classical methods	343
7.4.2 Computing Constrained Controllability	343
7.4.3 Maintaining Performance Under Constraints	345
7.5 Numerical Results	348
7.5.1 Distributed Control of Linearized Swing Equations	348
7.5.2 Computation of Rank Condition Versus Volume	349
7.6 Concluding Remarks	350
7.6.1 Proof of Lemma 7.9	350
7.6.2 Proof of Lemma 7.10	350
Chapter VIII: Learning the Uncertainty Sets for Control Dynamics via Set Membership: A Non-Asymptotic Analysis	351
Abstract	352
8.1 Introduction	353
8.2 Problem Formulation and Preliminaries	355
8.2.1 Problem Formulation	355
8.2.2 Set Membership Estimation (SME)	356
8.3 Set Membership Convergence Analysis	359
8.3.1 Convergence Rate of SME with Known $w_{\max}$	359
8.3.2 SME with Unknown $w_{\max}$	361
8.4 Proof Sketch of Theorem 8.8	363
8.5 Applications to Robust Adaptive Control	366
8.6 Numerical Experiments	367
8.7 Concluding Remarks	369
8.A Additional notations	370
8.B More discussions on least square and set membership	371
8.B.1 Least square estimation	371
8.B.2 Set membership	372

8.C	More discussions on Assumptions 8.5 and 8.7	373
8.C.1	More discussions on Assumption 8.5	373
8.C.2	More discussions on Assumption 8.7	374
8.D	Proof of Theorem 8.8	375
8.D.1	Technical lemmas: set discretization	375
8.D.2	Proof of Lemma 8.16	376
8.D.3	Proof of Lemma 8.17	378
8.D.4	A more precise upper bound for Theorem 8.8	385
8.E	Proof of Corollary 8.9	385
8.F	Proof of Corollary 8.10	386
8.G	Proof of Theorem 8.11	386
8.H	Proofs of Theorem 8.12, Corollary 8.13, and Theorem 8.14	388
8.H.1	Proof of Theorem 8.12	388
8.H.2	Proof of Corollary 8.13	391
8.H.3	Proof of Theorem 8.14	392
8.I	Simulation details and additional experiments	393
8.I.1	Baseline: LSE's confidence regions	393
8.I.2	Figure 8.1: SME and LSE uncertainty set visualization	394
8.I.3	Figure 8.2	394
8.I.4	Figure 8.3	394
8.I.5	Figure 8.4	395
Chapter IX:	Convex Constrained Controller Synthesis for Evolution Equations	396
Abstract		397
9.1	Introduction	398
9.2	Problem Setup	399
9.2.1	Notation	399
9.2.2	Strong and Weak Forms of the Dynamics	400
9.2.3	System Level Synthesis	401
9.3	Results	402
9.3.1	State Feedback	403
9.3.2	Output Feedback	404
9.3.3	Optimization and Structural Constraints	405
9.3.4	Parallel Solving for Integral Transforms	406
9.4	Control Implementation	408
9.4.1	Simulation Results	408
9.5	Conclusion	410
9.5.1	Acknowledgements	411
9.A	Theorem Proofs	411
9.B	Implementation Details	414
PART II	BIBLIOGRAPHY	419
Index		428

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 A barycentric map $\mathbf{x}$ maps from the simplex $\triangle_J = \triangle_3$ to the space $X = \mathbb{S}^2$. The simplex vertices $(e_J^j)_j$ map to points in X , and interpolations in $\triangle_J$ map to interpolations in X	25
1.2 Poverty index scores in 1995 and 1998 are plotted, with the red line the threshold below which households qualify for government aid programs. In summer 1997, information about how the poverty index score is computed from census data, along with the threshold, was made publicly available, and a distribution shift appeared in 1998.	29
2.1 After the criteria for government aid was released in 1997, local officials misreported income data to increase the number of constituents qualifying for aid. The PDE (2.3) is able to capture the sharp drop at the classifier threshold. The convergence rate for the loss of the population and algorithm are 0.00995 and 0.0102; the convergence rate for $\mathcal{W}_2(\rho_t, \rho^{(98)})$, where $\rho^{(98)}$ is the steady state distribution, is also 0.0114 which is similar to the expected rate of 0.01. The expected rate is computed using convexity properties of the KL term.	85
2.2 While the accuracy of the classifier is similar under both interaction models, the precision differs; this indicates that understanding the intra-agent interactions is important for understanding how errors impact subpopulations, in this case, those agents with algorithm label “qualified.”	86
2.3 Agent densities are split based on their true label (eligible vs ineligible). The color denotes the algorithm output regarding loan qualification (qualified vs unqualified). The repulsive kernel causes the agents to spread apart along attributes 1 (normalized age) and 2 (normalized past due), while the attractive-repulsive kernel causes more swarm-like behavior. This is not evident from classifier performance only, indicating the importance of understanding the population dynamics explicitly.	87
2.4 A gradient descent approach outperforms a state-of-the-art technique of learning and using a linear mapping between the classifier parameters and the mean of the strategic distribution, illustrating the importance of having more detailed population models. .	88

3.1	The relationships between main results (in blue), assumptions (in violet), and the monotonicity condition (in orange) are illustrated. Assuming that solutions to (3.1) are absolutely continuous curves (Assumption 3.11) allows us to compute the time derivative of the Wasserstein-2 metric via [56, Thm 3.14] to show contraction. Then via [52, Lemma 7.3], a unique steady state $\rho^\infty \in \overline{\mathcal{P}_2}$ exists (Theorem 3.17a). Convergence of $\rho(t)$ to ρ^∞ follows by similar arguments (Theorem 3.17b). If the dynamics are a coupled gradient flow, that is, take the form (3.2), then using Assumption 3.11, Assumption 3.12 and displacement convexity of each F_i (Lemma 3.28a), we show that ρ^∞ is critical point and the unique Nash equilibrium of F (Theorem 3.17c). The convergence of the velocity fields (Proposition 3.18) and second moment bounds (Lemma 3.21) for solutions to (3.1) follow from monotonicity.	153
3.2	(Multi-Learner System) The fraudulent population (blue) and valid population (orange) both aim for high probability of approval from the classifiers. The black lines labeled x_1 and x_2 mark the 50% classification probability threshold for h_1 and h_2 respectively, with the arrows pointing in the direction of increasing probability of loan approval. In this setting, the fraudulent population has an interaction kernel, and the fraudulent population splits into two modes, strongly preferring the blue classifier.	168
3.3	(Multi-Learner System) In the setting without any interaction kernels, the fraudulent population does not split. Additionally, there are more fraudulent agents obtaining services from bank 2 than in the setting with the interaction kernel.	168
3.4	(Degenerately Monotone Game with Diffusion) The upper-bound for the convergence rate across the tested values of b is $\lambda = 0.31$. As the coupling between the two players strengthens, the oscillations of F increase in frequency, but still decay at roughly the same rate.	171
3.5	(Degenerately Monotone Game with Diffusion) In the setting where $b = 150$, the density of each species is plotted from $t = 0$ to $t = 4$, during which time the energy has oscillated a few times. The densities for ρ_1 (blue) and ρ_2 (orange) spread out due to diffusion, eventually minimizing a quartic-shaped basin. The species ρ_2 aims to overlap with ρ_1 while ρ_1 aims to be in the opposite quadrant as ρ_2 , due to the coupling term $\iint x_1^\top x_2 d\rho_1(x_1)d\rho_2(x_2)$, which ρ_1 aims to minimize and ρ_2 aims to maximize.	171
4.1	The nomenclature and coefficients for convexity, monotonicity, accretivity, and dissipativity as used in this work are illustrated in particular settings.	194
4.2	A multispecies coupled gradient flow system in the Wasserstein-2 metric is not necessarily a joint gradient flow system. Particular choices of energies $F^{(i)}$ result in min-max flows or a single-species gradient flow.	196

4.3	The strip $Q_{0,T}^\varepsilon$	245
5.1	When x versus ρ best-responds, we observe the same final state but different intermediate states. Modes appear in the strategic population which simpler models cannot capture.	262
5.2	In this experiment the population and classifier have similar rates of change, and the distribution change for ρ exhibits behaviors from both the fast ρ and fast x simulations; the right-hand mode does not peak as high as the fast ρ case but does exceed its final height and return to the equilibrium.	263
5.3	Although the classifier starts with a larger cost by taking the naive strategy, the final loss is better. This illustrates how our model can be used to compare robustness of different strategies against a strategic population.	264
5.4	The dynamics include a consensus kernel, which draws neighbors in z -space closer together. We see that the population moves to make the classifier performer better, as the two distributions become more easily separable by the linear classifier.	295
5.5	The population ρ aims to be classified with the τ population, while the classifier moves to delineate between the two. We observe that τ adjusts to improve the performance of the classifier while ρ competes against it. The distributions are plotted at time $t = 0$, corresponding to $\tilde{\rho}$ and $\tilde{\tau}$, and time $t = 20$, corresponding to ρ and τ	296
5.6	When the classifier is updating at the same rate as the population, we do not see a significant change in the evolution of both species, suggesting that as long as the gradient estimate for the classifier is correct on average, the estimate itself does not need to be particularly accurate.	297
5.7	When the classifier is best-responding to the population, we observe that using $n = 4$ samples leads to different behavior for both the classifier and the population, compared with a more accurate estimate using $n = 40$ samples.	297
5.8	We use (5.41) for the classifier functions, using a Gaussian initial condition and regularizer for ρ . We see the distribution moving toward the region with higher probability of misclassification.	298
6.1	The transport map generated with a kernel density estimator method from [153] (top) is limited by the number of points ($N = 1,000$), whereas estimating the density by counting voxels (bottom, our method) allows for a greater number of points (20,000) and faster computation time.	315

6.2	In the setting of two log-concave input measures, the coupling cost starts small due to the initialization of $X_i \approx Y_i$, and increases to approach the optimal coupling cost. The regularization penalty weight Λ increases more slowly over time, as expected due to the $O(\sqrt{t})$ estimate, and the KL divergence between the marginals and the inputs decreases as the particles drift to minimize the entropy regularization terms.	315
6.3	The source and target distributions for the experiment (b) are shown. Note that the input marginals are not log-concave, going beyond Assumption 6.3.	316
6.4	We test our method on two input measures which are non-log-concave (Left side is μ and right side is ν). When running our method with two different entropy regularizers, using $\text{KL}(\rho \mu, \nu)$ for the $Y^{(2)}$ particles and $\text{KL}(\mu, \nu \rho)$ for the $X^{(2)}$ particles (labeled "both KL") is compared with using the same for all particles. While the displacement convexity of $\text{KL}(\cdot \mu)$ is well known, the displacement convexity of $\text{KL}(\mu \cdot)$ is not.	317
7.1	When actuation delay is zero, the addition of locality constraints on the system does not change the controllability volume for locality greater than five. The controllability decreases with increasing actuation delay.	348
7.2	As the state dimension grows, the amount by which the rank condition computation time is faster than the controllability volume time increases.	349
8.1	A visualized toy example of uncertainty set comparison between SME in (8.2) and LSE confidence regions in [110, 2] for a one-dimensional system $x_{t+1} = A^*x_t + B^*u_t + w_t$, with $w_t, u_t \in [-1, 1]$ generated i.i.d. from a truncated Gaussian distribution. Detailed experiment settings are in Section 8.I. Figure (a) compares the diameters of the uncertainty sets from SME and LSE 90% confidence bounds. Figure (b) and (c) visualize the the uncertainty sets after $T = 5$ and $T = 250$ data points.	354
8.2	Figures (a)-(b) compares the diameters of SME, UCB-SME, and SME with loose disturbance upper bounds that are 2, 3, 5, and 10 times larger than the true disturbance bound $w_{\max}$, as well as the baseline uncertainty set from the 90% confidence region of LSE. Figure (c) shows the convergence to the true bound $w_{\max}$ of the lower estimation $\bar{w}_{\max}$ in (8.4) and the UCB $\hat{w}_{\max}$ generated by the UCB-SME algorithm in Figures (a)-(b)	364
8.3	Diameters of the uncertainty sets constructed by SME, UCB-SME, and LSE for systems with different dimensions.	368
8.4	Linear quadratic tracking of robust adaptive MPC based on SME, LSE's confidence regions, and the accurate model (OPT).	369

9.1	Traditional control approaches discretize PDE dynamics and then solve a finite-dimensional optimal control problem (discretize-then-optimize). We propose optimizing in the infinite-dimensional space, and then discretizing the solution (optimize-then-discretize). This requires an infinite-dimensional SLS framework. . .	398
9.2	Optimal state response (continuous green shade) and control magnitude (red dot) after $t = 1$, and $t = 2$	409
9.3	The state and control responses to each of the 15 disturbances are computed in parallel, and then summed to obtain the final control values. The closed-loop maps are constrained to have support locally around each disturbance.	410

LIST OF TABLES

<i>Number</i>	<i>Page</i>
3.1 (Multi-Learner System) At times $T = 6, 12$, bank 1 has a larger share of the fraudulent population (ρ_1) than bank 2, with a larger margin in the kernel setting. At time $T = 20$, the valid population (ρ_2) prefers bank 2 with or without the interaction kernel, but the interaction kernel changes the bank preference for the fraudulent population.	168
3.2 (Coupled Gradient Flow Setting) The key difference between monotonicity and convexity is the operator $\bar{\nabla}$ which is <i>not</i> a gradient operator; it is an entry-wise operator which takes a gradient with respect to different states.	177
3.3 (Coupled Gradient Flow Setting) Similar to the $\mathbb{R}^d$ setting, the difference between monotonicity and convexity is the entry-wise operators. The second order inequalities have two Hessian-like terms in the measure setting, unlike the $\mathbb{R}^d$ setting.	178
6.1 Using method (III) results in the smallest L^2 error among the three methods because it takes advantage of $\tilde{E}$ being a better flow in the direction from ν ($s = 1$) to μ ($s = 0$) and E being a better flow in the direction from μ ($s = 0$) to ν ($s = 1$). The large error for the $\tilde{E}$ flow is due to the poor performance for $\tilde{E}$ in the μ to ν direction, and this is avoided in method (III).	317
9.1 Relative error and performance gain at each time step.	409
9.2 Performance comparison with finite-dimensional SLS [7].	410

Chapter 1

INTRODUCTION

Large scale heterogeneous multi agent systems have permeated society in nearly every aspect of human life; human-algorithm interactions occur constantly: in traffic with autonomous vehicles, online social networks, shopping and pricing algorithms, and incessant queries of large language models. In these high-dimensional complex systems, strategic behavior and nonlocal interactions play a crucial role in characterization and control. An *agent* is any strategically evolving entity, such as a person, algorithm, AI bot, or even a biological cell. These systems of agents appear across many applications, such as tumors with diverse cell species [25], adversarial agents exerting significant control on machine learning algorithms [34], physical phenomena such as fluid flow, biological systems [56, Ch 8], traffic [48], and in emerging areas like generative modeling [42] and opinion dynamics [28].

These interactions occur at massive scale, with millions of people commuting, posting on social media, purchasing staples, and asking agents for information every day. To accurately model and influence strategic behavior at unprecedented scales, traditional techniques that model individual people and agents are no longer feasible. The scale and complexity of these systems demands a new modeling paradigm — one that remains tractable without sacrificing the behavioral richness necessary for meaningful analysis and control.

This thesis pursues two complementary approaches: one, aggregate modeling and analysis, and two, granular control.

The first approach replaces individual agent modeling with a population-level description, representing the collective evolution of agents as the evolution of a probability density governed by a partial differential equation (PDE). This approach is well-established in fluids and in biology, where modeling at the atomic or cell level is infeasible; in contrast modeling at the distributional level results not only in analytical feasibility, but also yields insights into aggregate phenomena that are obscured at the microscopic scale. However, in many of the applications considered here, agents exhibit significant heterogeneity, and a single-population description fails to capture the diversity of behaviors or the interactions between distinct agent types. A multispecies framework is needed, where each species is a characterization of a behaviorally distinct subpopulation and the coupling encodes their interactions.

A natural analytical tool for studying both optimization over the space of distributions and the evolution of distributions is a gradient flow, the infinite-dimensional analogue of gradient descent.

Many results surrounding gradient descent and its variants, such as regularized gradient descent, stochastic gradient descent, and accelerated gradient descent, form the backbone of optimization in machine learning algorithms. While its infinite-dimensional counterpart is well-understood for a single species, the multispecies setting has many fundamental open problems. Even basic questions such as well-posedness of solutions are unresolved in general settings. Developing the analytical foundations necessary for modeling and control in this framework requires substantial new theory. This thesis addresses these open problems by establishing new frameworks and the first results for existence of solutions and asymptotic behavior for more general classes of multispecies gradient flow systems.

The second approach reverses the perspective: rather than aggregating agents into a population, we zoom in to design controllers that act directly on individual components of a large-scale dynamical system. For systems evolving in a massive state space — such as power grids, communication networks, or coupled physical systems — computing a globally optimal controller is computationally intractable. Even when such a controller exists in principle, its implementation requires real-time knowledge of the entire system state, which is physically impossible when information propagates on the same timescale as the dynamics themselves.

A recent parameterization known as System Level Synthesis (SLS) addresses this tension by enabling distributed control: controllers that can be computed with global optimality guarantees while being constrained to act on only locally available information. This framework decouples the computational and implementation challenges inherent to large-scale control, making it a powerful foundation for tractable, deployable controllers.

Nevertheless, fundamental questions remain. When locality constraints are imposed, it is unclear how much optimality is sacrificed — quantifying this performance gap is essential for calibrating the tradeoff between feasibility and quality. Furthermore, as large-scale systems are increasingly represented by continuous or distributional models, it is natural to ask how the resolution at which one models the system — aggregate or granular — interacts with the design of optimal controllers. Finally, when the system dynamics are unknown or subject to adversarial disturbances, classical controller synthesis may fail entirely, raising the question of whether dynamics can be reliably identified and controlled even in hostile environments. This thesis addresses these challenges by developing computational methods for controller synthesis within the SLS framework, extending the framework to nonlinear dynamics, and providing rigorous quantification of how structural constraints such as locality affect achievable performance.

The introduction section contains background material and overviews of the key results in this thesis. Part I includes the novel framework and results at the aggregate-level using gradient flow approaches, and Part II discusses the granular-level distributed control theory background and results.

1.1 Part I: Multispecies Gradient Flow Analysis

Complex systems composed of heterogeneous, strategic agents interacting with each other and algorithms impact our daily lives. Predicting and influencing their behavior requires rigorous theoretical analysis of these systems, but existing methods are largely insufficient for capturing inter-species interactions between distinct agent populations.

To account for these interactions and building upon existing literature, we propose expressing these evolutions as coupled *gradient flows*, infinite dimensional descent schemes which optimally evolve distributions to minimize functionals, analogous to gradient descent in Euclidean space. The gradient flow framework has advanced modeling and analysis for a number of applications, such as bacterial chemotaxis [43, 41, 30], diffusion phenomena, stochastic agent systems driven by potential functions and nonlocal interaction potentials, and molecular particle transport [35, 40, 14]. Recent years have seen a surge in the application of gradient flow theory including work in Monte Carlo sampling [29, 46], image processing (image registration, warping, shape classification, image segmentation and image restoration) [7, 33, 47, 59], modeling the behavior of plastic materials in material science [50], among many others. This is now a well-understood framework for single species dynamics [3], but largely unexplored for multispecies systems. Most commonly, asymptotic analysis for coupled PDE systems is done case-by-case [9, 26, 32], leveraging special structure and addressing only a relatively small family of examples.

Part I develops a general framework for coupled multispecies systems by drawing from game theory and analysis of PDEs, leveraging both the expressivity of infinite-dimensional objects and the rich complexity of interactions between distinct groups. The results pave the way for predicting and guaranteeing system behaviors, like stability and pattern formation, in biological, machine learning, and control systems.

1.1.1 Gradient Flow Introduction

This section outlines the fundamental mathematical concepts underlying gradient flows. Let (X, d) be a finite- or infinite-dimensional metric space, and let $F : X \rightarrow \mathbb{R}$ be an energy or cost functional. In many applications, the goal is to find and characterize the minimizers of F . For $X = \mathbb{R}^d$, one computational approach to find a minimizer is gradient descent; in general metric spaces, a gradient flow describes the evolution of a point in a metric space in the direction of steepest descent of F . The central questions are:

- Does a minimizer of F exist in X ? Is it unique?
- Does the solution to a gradient flow on F converge to a minimizer, and under what conditions?

To build intuition for the general setting, we first answer these questions in the Euclidean setting. For a sufficiently smooth function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, a gradient flow evolution is equivalent to gradient descent:

$$\dot{x} = -\nabla f(x), \quad x \in \mathbb{R}^d. \quad (1.1)$$

First, we ask under what conditions f admits a minimizer. The extreme value theorem guarantees that on any proper, closed, and bounded set $D \subset \mathbb{R}^d$, a continuous function on $\mathbb{R}^d$ has a minimizer $x^* = \arg \min_{x \in D} f(x)$. From this theorem, we can determine conditions under which the function has a minimizer on unbounded sets. Let the sublevel sets of f be defined as

$$\sigma_-(c) = \{x \in \mathbb{R}^d : f(x) \leq c\}.$$

We say f is *coercive* if the nonempty sublevel sets of f are bounded for all $c \in \mathbb{R}$. Coercivity ensures that f tends to infinity as $\|x\| \rightarrow \infty$. Given any nonempty sublevel set $\sigma_-(c_0)$, coercivity guarantees $\sigma_-(c_0)$ is contained in a compact set, and the extreme value theorem yields a minimizer $x^* \in \sigma_-(c_0)$. Since all x outside of $\sigma_-(c_0)$ satisfy $f(x) > c_0$, one can conclude that x^* is the global minimizer of f .

One sufficient condition for coercivity, resulting in both uniqueness of the minimizer and convergence of the gradient descent scheme (1.1) to x^* , is strong convexity of f ; we say f is λ -convex if

$$\begin{aligned} f(\alpha x + (1 - \alpha)y) &\leq \alpha f(x) + (1 - \alpha)f(y) - \frac{\lambda}{2}\alpha(1 - \alpha) \|x - y\|^2 \\ \forall \alpha \in [0, 1], \forall x, y \in \mathbb{R}^d, \quad &\text{(zeroth order)} \end{aligned} \quad (1.2)$$

and we say f is *strongly convex* if $\lambda > 0$. If f is once- or twice- differentiable, we can derive the following first and second order conditions for λ -convexity, where the "order" corresponds to the number of gradients in the condition:

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \lambda \|x - y\|^2, \quad \forall x, y \in \mathbb{R}^d, \quad \text{(first order)} \quad (1.3)$$

$$\nabla^2 f(x) \succeq \lambda I, \quad \forall x \in \mathbb{R}^d. \quad \text{(second order)} \quad (1.4)$$

Notice that the second order condition relies only on one test point $x \in \mathbb{R}^d$, whereas the zeroth and first order conditions rely on two points $x, y \in \mathbb{R}^d$. In this sense, the second order condition is a *local* convexity condition (f is convex at x), and the zeroth and first order conditions are *nonlocal* conditions. We will utilize analogs of all three notions later in more general metric spaces.

Strong ($\lambda > 0$) convexity implies uniqueness of the minimizer x^* . Suppose another minimizer $x^{**} \neq x^*$ exists; then the first order convexity condition implies

$$\langle \nabla f(x^*) - \nabla f(x^{**}), x^* - x^{**} \rangle \geq \lambda \|x^* - x^{**}\|^2,$$

which is a contradiction because $\nabla f(x^*) = \nabla f(x^{**}) = 0$. To show that strong convexity results in convergence of (1.1) to the unique minimizer of f , consider the following Lyapunov functions:

$$V_1(x) = \frac{1}{2} \|x - x^*\|^2, \quad V_2(x) = \frac{1}{2} \|\nabla f(x)\|^2.$$

The first function V_1 requires knowledge of the minimizer, while the second requires differentiability of f . Taking time derivatives of each along solutions to (1.1),

$$\begin{aligned} \frac{d}{dt} V_1(x(t)) &= -\langle x(t) - x^*, \nabla f(x(t)) - \underbrace{\nabla f(x^*)}_{=0} \rangle \\ &\stackrel{(1.3)}{\leq} -\lambda \|x(t) - x^*\|^2 = -2\lambda V_1(x(t)), \\ \frac{d}{dt} V_2(x(t)) &= -\langle \nabla f(x(t)), \nabla^2 f(x(t)) \cdot \nabla f(x(t)) \rangle \\ &\stackrel{(1.4)}{\leq} -\lambda \|\nabla f(x(t))\|^2 = -2\lambda V_2(x(t)). \end{aligned}$$

Applying Grönwall's inequality and writing the definitions of V_1 and V_2 results in

$$\|x(t) - x^*\| \leq e^{-\lambda t} \|x(0) - x^*\|, \quad \|\nabla f(x(t))\| \leq e^{-\lambda t} \|\nabla f(x(0))\|.$$

When $\lambda > 0$, the above estimates show that $x(t)$ converges exponentially in norm to the minimizer x^* , and that the norm of the gradient also converges exponentially to zero, both with rate λ . For strongly convex functions, gradient descent therefore forms the basis of an algorithm for finding the minimizer of a function.

In the more general setting, we would like to identify minimizers of functionals, define strong convexity, and define a notion of gradient descent for functionals on metric spaces, for example, probability measures or points on the simplex. A number of challenges arise in more general settings, particularly when the metric space is infinite-dimensional.

Existence of minimizers. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, conditions such as coercivity and continuity or strong convexity (which is stronger than coercivity) and lower-semicontinuity can be used to show existence of a minimizer [57, Cor 2.20]. In general metric spaces, because closed and bounded sets are not necessarily compact in infinite dimensions, compactness is required in addition to coercivity. Moreover, coercivity and lower-semicontinuity can fail for infinite dimensional settings while holding in analogous finite dimensional settings. For example, consider the entropy function on finite-dimensional measures in comparison to the entropy function on infinite-dimensional measures. The functional on finite-dimensional measures is $E : \Delta_d \rightarrow \mathbb{R}$, where $\Delta_d := \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i = 1, x_i \geq 0 \forall i\}$ is the d -dimensional simplex, and the infinite-dimensional functional is $\tilde{E} : \mathcal{P}_1(\mathbb{R}^d) \rightarrow \mathbb{R}$, where $\mathcal{P}_1(\mathbb{R}^d)$ is the set of probability measures on $\mathbb{R}^d$ with a finite

first moment. Let $\mathcal{L}$ be the Lebesgue measure in $\mathbb{R}^d$, and let $\rho \ll \mathcal{L}$ denote that ρ is absolutely continuous with respect to $\mathcal{L}$. The functionals are defined as

$$E(x) = \sum_{i=1}^d x_i \log x_i, \quad \forall x \in \Delta_d,$$

$$\tilde{E}(\rho) = \begin{cases} \int \rho(x) \log \rho(x) dx & \rho \ll \mathcal{L} \\ +\infty & \text{otherwise} \end{cases}, \quad \forall \rho \in \mathcal{P}_1(\mathbb{R}^d).$$

In the finite-dimensional setting, E is continuous in x with respect to any Euclidean norm, and the domain $\Delta_d \subset [0, 1]^d$ is closed and bounded, and therefore compact due to the finite dimensionality. It has a minimizer at $x_i = [1/d, \dots, 1/d]$.

In contrast, the functional $\tilde{E}$ does not have a minimizer over its domain. This functional is 0-convex in the sense of displacement convexity (see (1.12)), but lower-semicontinuity only holds [60] with respect to sequences (ρ_n) that satisfy

$$\int f d\rho_n \rightarrow \int f d\rho \quad \text{for all } f \in C^\infty, \quad \int \|x\| \rho_n(x) \leq C, \quad \tilde{E}(\rho_n) \leq C.$$

Because of the requirement that $\int \|x\| \rho_n(x) \leq C$, E is *not* lower-semicontinuous on all of $\mathcal{P}_1(\mathbb{R}^d)$. Similarly, existence of a lower bound holds only on subsets of $\mathcal{P}_1(\mathbb{R}^d)$ with uniformly bounded first moments. To observe this, let $\rho_\varepsilon(x) = \varepsilon$ for $x \in [0, 1/\varepsilon]$ and $\rho_\varepsilon(x) = 0$ otherwise, so that $\int \|x\| d\rho_\varepsilon(x) = (2\varepsilon)^{-1}$ and $\tilde{E}(\rho_\varepsilon) = \log(\varepsilon)$. As $\varepsilon \rightarrow 0$, the first moment $(2\varepsilon)^{-1}$ goes to $+\infty$ and $\tilde{E}(\rho_\varepsilon)$ goes to $-\infty$. The field of calculus of variations includes many tools for proving existence of minimizers for various metric spaces, and we will utilize results from this area in constructing results for gradient flows in infinite dimensions.

Notion of convexity and interpolations. The notion of convexity (1.2) uses the interpolation $x_\alpha := \alpha x + (1 - \alpha)y$ between $x, y \in \mathbb{R}^d$. One can show that this interpolation is the *shortest path* between x and y in any norm, that is, of all possible continuous paths from x to y , the curve x_α has the shortest length. When considering an appropriate notion of convexity for general metric spaces, the interpolation choice $\alpha x + (1 - \alpha)y$ can pose issues. First, metric spaces are not required to have notions of addition, and even with notions of addition, one must ensure that $\alpha x + (1 - \alpha)y$ is in the metric space for all values of $\alpha \in [0, 1]$. More importantly, this interpolation is not necessarily the shortest path in a metric space. Later we will define *geodesics*, which are shortest paths in metric spaces, and use these paths to define a notion of convexity called *geodesic convexity*, also known as *displacement convexity*.

Notion of gradient. For a differentiable function $f \in C^1(\mathbb{R}^d, \mathbb{R})$, the definition of a gradient of f at

a point x is the slope of f at x in each coordinate of x :

$$\nabla f(x) := \lim_{h \rightarrow 0} \frac{1}{h} \begin{bmatrix} f(x + he_1) - f(x) \\ \vdots \\ f(x + he_d) - f(x) \end{bmatrix},$$

where $e_i \in \mathbb{R}^d$ is the zero vector with the i^{th} element set to one. In the above definition, the point x is perturbed in all possible directions in $\mathbb{R}^d$. To formulate a gradient for a functional on an infinite-dimensional space, two issues arise. One, there are an infinite number of directions in which one can perturb the function input in order to compute the gradient. Secondly, and perhaps more importantly, the directions in which the function input can be perturbed depend on the choice of metric.

For example, consider a measure $\rho \in \mathcal{P}_2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$. If we utilize the L^1 metric, the perturbation to ρ cannot include any direction with a Dirac, since the resulting measure does not remain in L^1 . In contrast, if we chose the $\mathcal{W}_2$ metric (See Definition (1.10)), we *can* use perturbations with Diracs. We cannot perturb the measure to having an infinite second moment, whereas that perturbation is feasible in the L^1 metric. We will later see that the choice of metric determines the direction of steepest descent in gradient flows.

While many new mathematical tools are required to move from Euclidean space to general metric spaces for gradient descent evolutions, the intuition from Euclidean space will guide definitions and proof techniques, particularly for multispecies systems. The next section outlines the key structures and techniques for gradient flow analysis in metric spaces.

1.1.2 Well-Posedness and Asymptotic Behavior of a Single Species Gradient Flow

In this section, we provide an overview of techniques for showing well-posedness of a gradient flow evolution and analysis of its long-time behavior. The method we use for well-posedness relies heavily on the existence of minimizers, and additionally the natural candidates for asymptotic profiles of a gradient flow are the minimizers of the corresponding energy functional. Therefore, we begin by investigating existence of minimizers. Then we can determine if minimizers are also steady states of the evolution equation; this often requires that the minimizers be regular enough. Then we check under what conditions solutions to the evolution equation exist, using a time-discretized approximation of the gradient flow. Lastly, we then show the solutions converge to a steady state, possibly given by a minimizer of the energy functional, as $t \rightarrow \infty$.

Existence of Minimizers. To show the existence of minimizers of a given functional $G : X \rightarrow \mathbb{R}$ on the metric space (X, d) , we illustrate the *direct method in the calculus of variations* using the assumptions given in [3, Section 2.1]. We first show the existence of a minimizing sequence

$(x_n) \in X$, second, show that the limit of the sequence exists, and third, that it is the minimizer of G . We select a topology σ which is potentially weaker (in the sense that sequential compactness of sublevel sets is easier to show, while lower semicontinuity of d is harder to guarantee) than the topology induced by d , such that d is sequentially σ -lower semicontinuous. This means

$$(u_n, v_n) \xrightarrow{\sigma} (u, v) \quad \Rightarrow \quad \liminf_{n \rightarrow \infty} d(u_n, v_n) \geq d(u, v).$$

In the above, the lower semicontinuity assumption on d limits how weak σ can be. The following three key assumptions adjusted from [3, Section 2.1] are sufficient for the argument:

1. Bounded and proper: $\inf_{x \in X} G(x) > -\infty$ and G is not uniformly equal to $+\infty$. This ensures that there exists some (x_n) such that $G(x_n) \rightarrow \inf_{x \in X} G(x)$, and $\sup_n G(x_n) < +\infty$.
2. Compactness: Sublevel sets $(S_c)_{c \in \mathbb{R}}$ of G are relatively σ -sequentially compact, where S_c is defined as

$$S_c = \{u \in X \mid G(u) \leq c\}.$$

Then the minimizing sequence (x_n) admits a σ -convergent subsequence because it is, by definition, contained in a sublevel set of G which is σ -compact.

3. Lower semicontinuity: for any sequence $(x_n) \in (X, d)$ such that $x_n \xrightarrow{\sigma} \bar{x}$,

$$\liminf_{n \rightarrow \infty} G(x_n) \geq G(\bar{x}).$$

From the compactness assumption, we have the existence of a convergent subsequence $(x_{n_k}) \xrightarrow{\sigma} \bar{x}$, which allows us to apply lower-semicontinuity to show that the limit $\bar{x}$ achieves the minimum value of G .

In [3, Section 2.1], Assumptions 2 and 3 are weakened so that compactness and lower semicontinuity only need to hold for d -bounded sequences, that is, for (x_n) satisfying $\sup_{n,m} d(x_n, x_m) < \infty$; the above assumptions are selected to fit easily with the direct methods argument.

In Euclidean space, strong convexity implies coercivity. In the metric setting, F is λ -convex along a curve $u : [0, 1] \rightarrow X$ if

$$F(u_s) \leq (1-s)F(u_0) + sF(u_1) - \frac{\lambda}{2}s(1-s)d^2(u_0, u_1), \quad \forall s \in [0, 1]. \quad (1.5)$$

If for any two points $x, y \in X$, there exists some u satisfying $u_0 = x$ and $u_1 = y$ such that F is strongly convex ($\lambda > 0$) along u and lower-semicontinuity holds, then a unique minimizer exists [3, Lemma 2.4.8]. While this is a sufficient condition, existence of a minimizer can still occur in other settings. Once we have shown that F has a minimizer x^* , we can ask whether the solution to a gradient flow converges to x^* . Before understanding the long-time asymptotics, we first need to formulate the gradient flow dynamics and show that solutions exist.

EVI Gradient Flow Formulation. To gain intuition for the gradient flow construction, we first consider a Hilbert space setting. For finite-dimensional or Hilbert-space energy functionals, we use lower-case f and for general metric spaces, we use F . In a Hilbert space $X = H$ with metric $d(x, y) = \langle x - y, x - y \rangle^{1/2} = \|x - y\|$, the weak form of a gradient flow for a λ -convex functional f can be defined by an *evolution variational inequality (EVI)*:

$$\left\langle \frac{d}{dt}x(t), x(t) - y \right\rangle + \frac{\lambda}{2}d^2(x(t), y) + f(x(t)) \leq f(y), \quad \forall y \in H.$$

The above formulation allows for f that are not differentiable; weakening the notion of solution, compared to a pointwise notion of solution involving the gradient of f , enables a more general class of functions f for which the gradient flow can be defined. In the setting where f is differentiable, the ODE $\frac{d}{dt}x(t) = -\nabla f(x(t))$ is a solution to the EVI, as seen by substituting $\frac{d}{dt}x(t)$, which results in the first-order λ -convexity inequality for f :

$$\langle \nabla f(x(t)), x(t) - y \rangle \geq f(x(t)) - f(y) + \frac{\lambda}{2}d^2(x(t), y), \quad \forall y \in X. \quad (1.6)$$

In spaces without an inner product structure, a different notion of $\langle \frac{d}{dt}x(t), x(t) - y \rangle$ is needed. Using the fact that

$$\left\langle \frac{d}{dt}x(t), x(t) - y \right\rangle = \frac{1}{2} \frac{d}{dt} \|x(t) - y\|^2 = \frac{1}{2} \frac{d}{dt} d^2(x(t), y),$$

the EVI can be written in a form without an inner product:

$$\frac{1}{2} \frac{d}{dt} d^2(x(t), y) + \frac{\lambda}{2} d^2(x(t), y) + F(x(t)) \leq F(y), \quad \forall y \in X, \quad t \geq 0. \quad (\text{EVI})$$

This EVI now requires only a notion of distance d and λ -convexity, and does not require a Hilbert space structure or notion of gradient. We will now show how to prove existence of a solution to this EVI, starting from a discrete-time minimization perspective.

Minimizing Movement Scheme. Consider the strongly convex and differentiable $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with unique minimizer x^* . When numerically implementing the gradient descent evolution $\frac{d}{dt}x = -\nabla f(x)$ we can select an implicit or explicit scheme for discretization in time, with time step size τ :

$$x^{k+1} = x^k - \tau \nabla f(x^k), \quad (\text{explicit})$$

$$x^{k+1} = x^k - \tau \nabla f(x^{k+1}) = \arg \min_{x \in \mathbb{R}^d} \left\{ f(x) + \frac{1}{2\tau} \|x - x^k\|^2 \right\}. \quad (\text{implicit})$$

The explicit formulation requires an upper Lipschitz bound $\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\|$ such that $\tau L < 2$ in order to ensure convergence of the scheme to the global minimizer [57, Prop 6.21].

However, the implicit formulation is stable without an upper Lipschitz bound, and additionally has an argmin formulation that does not depend on the gradient of f . This is particularly useful for the general metric setting, where the form of ∇F is often difficult to determine explicitly. We can generalize this formulation to the metric setting, and approximate the gradient step via

$$X_\tau^{k+1} = \arg \min_{x \in X} \left\{ F(x) + \frac{1}{2\tau} d^2(x, X_\tau^k) \right\}. \quad (1.7)$$

This is called the *minimizing movement scheme* (or JKO scheme¹ in the Wasserstein-2 setting), because each step minimizes F with a penalty for the distance from the previous point. We would like to take the limit as $\tau \rightarrow 0$ of the sequence (X_τ^k) generated by (1.7) to obtain a solution to (EVI). Any sequence (X_τ^k) generated by (1.7) with a λ -convex F also satisfies the discrete EVI,

$$\begin{aligned} \frac{1}{2\tau} \left(d^2(X_\tau^k, Y) - d^2(X_\tau^{k-1}, Y) \right) + \frac{1}{2\tau} d^2(X_\tau^k, X_\tau^{k-1}) + \frac{\lambda}{2} d^2(X_\tau^k, Y) \\ \leq F(Y) - F(X_\tau^k), \quad \forall Y \in X. \end{aligned} \quad (1.8)$$

Next, we show that the sequence (X_τ^k) generated by the minimizing movement scheme (1.7) satisfies the discrete EVI. Since $F(\cdot) + \frac{1}{2\tau} d^2(\cdot, X_\tau^{k-1})$ is $\left(\lambda + \frac{1}{\tau}\right)$ -convex along generalized geodesics² and X_τ^k is the unique minimizer, for all $\alpha \in [0, 1]$,

$$\begin{aligned} F(X_\tau^k) + \frac{1}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) &\leq \alpha \left(F(Y) + \frac{1}{2\tau} d_X^2(Y, X_\tau^{k-1}) \right) \\ &\quad + (1 - \alpha) \left(F(X_\tau^k) + \frac{1}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) \right) - \frac{\lambda + \tau^{-1}}{2} \alpha(1 - \alpha) d_X^2(X_\tau^k, Y), \end{aligned}$$

which, rearranging and dividing by α , results in

$$F(X_\tau^k) + \frac{1}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) \leq F(Y) + \frac{1}{2\tau} d_X^2(X_\tau^{k-1}, Y) - \frac{\lambda + \tau^{-1}}{2} (1 - \alpha) d_X^2(X_\tau^k, Y).$$

Taking $\alpha \rightarrow 0$ gives the discrete EVI

$$\begin{aligned} \frac{1}{2\tau} \left(d_X^2(X_\tau^k, Y) - d_X^2(X_\tau^{k-1}, Y) \right) + \frac{1}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) + \frac{\lambda}{2} d_X^2(X_\tau^k, Y) \\ \leq F(Y) - F(X_\tau^k). \end{aligned}$$

This weak formulation of (1.7) will allow us to take the limit $\tau \rightarrow 0$. Importantly, we need the following convexity condition. For any $\hat{x}, \hat{y}, X_\tau^k \in X$, there must exist a curve u with $u_0 = \hat{x}$, $u_1 = \hat{y}$ such that

$$x \mapsto F(x) + \frac{1}{2\tau} d^2(x, X_\tau^k) \text{ is } \left(\frac{1}{\tau} + \lambda \right) \text{-convex along } u \text{ for all } 0 < \tau < \frac{1}{\lambda^-},$$

¹Named after Jordan, Kinderlehrer, and Otto, who first derived the scheme for the Wasserstein-2 metric [37].

²The squared Wasserstein-2 metric is not 1-convex along geodesics, but it is 1-convex along *generalized geodesics*. Therefore convexity is considered along this class of curves. For details, see [3, Chapter 9].

where $\lambda^- := \max\{-\lambda, 0\}$ and $\frac{1}{\lambda^-} := +\infty$ if $\lambda > 0$. In order for this to hold, F must be λ -convex along some set of curves. Secondly, $d^2(\cdot, X_\tau^k)$ must be 1-convex along that same set of curves³. Equipped with this convexity condition, we now show that there exists some limiting curve $x(t)$ as $\tau \rightarrow 0$. We cannot directly compare $X_{\tau_1}^k$ with $X_{\tau_2}^k$ due to the different time discretizations; to overcome this, we construct piecewise constant interpolations with continuous-time solutions that are comparable for τ_1 and τ_2 . The piecewise interpolant, for all $t \in [0, T]$, is

$$X_\tau(t) = X_\tau^k, \quad \text{if } t \in [\tau k, \tau(k+1)).$$

In order to guarantee the existence of some limiting curve x as $\tau \rightarrow 0$, we will utilize a Cauchy estimate of the form

$$d(X_{\tau_1}(t), X_{\tau_2}(t)) \leq C(t, \tau_1, \tau_2) \tag{1.9}$$

where, for any fixed $t \geq 0$, $C(t, \tau_1, \tau_2) \rightarrow 0$ as $|\tau_1 - \tau_2| \rightarrow 0$. The estimate (1.9) relies on the $(\frac{1}{\tau} + \lambda)$ convexity of $F(\cdot) + \frac{1}{2\tau}d^2(\cdot, X_\tau^k)$; while F need not be strongly convex with $\lambda > 0$, convexity for some $\lambda \in \mathbb{R}$ is needed to construct C ; see [53] for details on how to construct such an estimate. The function $C(\cdot, \tau_1, \tau_2)$ is increasing in time if $\lambda < 0$ and decreasing if $\lambda > 0$. Since $d(X_{\tau_1}(t), X_{\tau_2}(t)) \rightarrow 0$ as $|\tau_1 - \tau_2| \rightarrow 0$, there exists a unique $x(t)$ which is the limit of $X_\tau(t)$ as $\tau \rightarrow 0$, pointwise in t . Taking $\tau \rightarrow 0$ in the discrete EVI (1.8), one can show that the limit $x(t)$ satisfies (EVI); more detailed computations are available in a more general setting in the proof of Theorem 4.42. This shows existence of a solution $x(t)$ to the weak gradient flow formulation (EVI).

Long-time Behavior. Finally, we want to show convergence to the minimizer. Recall our lower semicontinuity, coercivity, and compactness assumptions for F , which guarantee the existence of a minimizer. Consider (EVI) with test point $y = x^*$, the minimizer of F , which we know exists,

$$\frac{1}{2} \frac{d}{dt} d(x(t), x^*)^2 + \lambda d^2(x(t), x^*) + F(x(t)) \leq F(x^*).$$

Since $F(x^*) \leq F(x(t))$ by definition of x^* being a minimizer, we apply Grönwall's theorem to

$$\frac{1}{2} \frac{d}{dt} d^2(x(t), x^*) \leq -\lambda d^2(x(t), x^*),$$

to obtain

$$d(x(t), x^*) \leq e^{-\lambda t} d(x(0), x^*), \quad \forall t \geq 0.$$

When F is strongly convex ($\lambda > 0$), we see that $x(t)$ converges to x^* exponentially fast with rate λ .

³This poses an issue for our favorite metric space, Wasserstein-2 (1.10), since $\mathcal{W}_2^2(\cdot, X_\tau^k)$ is not 1-convex along geodesics, that is, curves of shortest distance between any two points as measured with the Wasserstein-2 distance. It is, however, convex along generalized geodesics with base point X_τ^k . For more details, see [3, Sections 9.1-9.2] and Remark 4.44 in Chapter 4.

1.1.3 Wasserstein-2 Metric

In this section, we introduce the Wasserstein-2 gradient flow setting. Let $\mathcal{P}_2(\mathbb{R}^d)$ be the space of probability measures with bounded second moments on $\mathbb{R}^d$. For any $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, the Kantorovich formulation of the Wasserstein-2 metric is given by

$$\mathcal{W}_2^2(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|z - z'\|^2 d\gamma(z, z') \quad (1.10)$$

where $\Gamma(\mu, \nu) \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$ is the set of all joint probability distributions with bounded second moments and marginals μ, ν , i.e. $\mu(dz) = \int \gamma(dz, z') dz'$ and $\nu(dz') = \int \gamma(z, dz') dz$.

Another way of writing the Wasserstein-2 Metric is via the *Benamou-Brenier formulation* (also known as the *dynamic formulation* of $\mathcal{W}_2$), expressing the distance between two probability measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ as the path solving a continuity equation with minimal kinetic energy:

$$\mathcal{W}_2^2(\mu, \nu) = \inf_{(\rho_t, w_t)} \left\{ \int_0^1 \int \|w(t, x)\|^2 d\rho(t, x) dt ; \right. \\ \left. \partial_t \rho_t + \nabla \cdot (\rho_t w_t) = 0 \text{ s.t. } \rho_0 = \mu, \rho_1 = \nu \right\}.$$

The Benamou-Brenier formulation is useful for geodesic calculations, as the optimal (μ, w) are also solutions to a geodesic equation. Indeed, any optimal curves (ρ_s, w_s) for the minimization problem above satisfy the following geodesic conservation of mass and momentum equations:

$$\partial_s \rho_s + \operatorname{div}(\rho_s w_s) = 0, \quad \partial_s(\rho_s w_s) + \operatorname{div}(\rho_s w_s \otimes w_s) = 0. \quad (1.11)$$

The solution ρ_s is often called the displacement interpolant. One key tool we will use is *geodesic convexity*, which is convexity via (1.5) where the curve is chosen to be a geodesic (shortest path) between ρ_0 and ρ_1 . Geodesic convexity in the Wasserstein-2 space $(\mathcal{P}_2(\mathbb{R}^d), \mathcal{W}_2)$ is known as *displacement convexity*. We denote by $\Gamma^*(\mu, \nu)$ the set of all $\mathcal{W}_2$ optimal plans between μ and ν :

$$\Gamma^*(\mu, \nu) = \left\{ \gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d) : \mathcal{W}_2^2(\mu, \nu) = \int \|z - z'\|^2 d\gamma(z, z') \right\}.$$

A functional $G : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ is *displacement convex* with constant λ if for all $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$ we have

$$G(\rho_s) \leq (1-s)G(\rho_0) + sG(\rho_1) - s(1-s)\frac{\lambda}{2}\mathcal{W}_2^2(\rho_0, \rho_1)^2, \quad (1.12)$$

where $\rho_s = [(1-s)x + sy]_{\#} \gamma$ is the displacement interpolant between ρ_0 and ρ_1 , with $\gamma \in \Gamma^*(\rho_0, \rho_1)$. If the energy F is twice differentiable along geodesics, then the condition $\frac{d^2}{ds^2} F(\rho_s) \geq \lambda \mathcal{W}_2^2(\rho_0, \rho_1)^2$ implies F is strongly displacement convex with constant $\lambda > 0$. For more details, see [49] and [68, Chapter 5.2]. Lastly, the chain rule for functionals with inputs in the Wasserstein-2 space is given by

$$\frac{d}{ds} F(\rho_s) = \int \delta_\rho F[\rho_s](z) \partial_s \rho_s(z) dz.$$

The notation $\delta_\rho F[\rho](z)$ denotes the first variation of F with respect to ρ , defined as

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (F(\rho + \varepsilon\psi) - F(\rho)) = \int \delta_\rho F[\rho](z) \psi(z) dz, \quad \forall \psi : \mathbb{R}^d \rightarrow \mathbb{R}, \quad \int \psi = 0.$$

Wasserstein-2 Gradient Flow Satisfies EVI. For some lower-semicontinuous, λ -displacement convex energy F , suppose the measure $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ evolves according to the Wasserstein-2 gradient flow

$$\partial_t \rho = -\nabla_{\mathcal{W}_2} F(\rho) = \operatorname{div}(\rho \nabla \delta_\rho F(\rho)). \quad (1.13)$$

The identification of $-\nabla_{\mathcal{W}_2} F(\rho)$ with $\operatorname{div}(\rho \nabla \delta_\rho F(\rho))$ was shown in [37, 55]. These important works introduced the minimizing movement scheme as a method for showing existence of solutions to a gradient flow in the Wasserstein-2 space, and the resulting framework was used to identify the explicit form of $\nabla_{\mathcal{W}_2}$. We will use s to index ρ along geodesics and t to index ρ along solutions to (1.13). We formally show that (1.13) satisfies (EVI) in the general metric setting. The time derivative of $\mathcal{W}_2^2(\rho_t, \mu_t)$, for two curves $\mu, \rho : [0, 1] \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ satisfying

$$\partial_t \rho_t + \operatorname{div}(\rho_t v_t) = 0, \quad \partial_t \mu_t + \operatorname{div}(\mu_t w_t) = 0,$$

is, for $\gamma(t) \in \Gamma^*(\rho_t, \mu_t)$,

$$\frac{d}{dt} \mathcal{W}_2^2(\rho_t, \mu_t) = \int \langle v_t(z) - w_t(z'), z - z' \rangle d\gamma(t, z, z').$$

See [12, Theorem 3.14] and [67, Prop 23.9] for rigorous statements of the derivative of $\mathcal{W}_2^2$ along curves. By selecting $v_t(z) = -\nabla \delta_\rho F[\rho](z)$, ρ_t satisfies the gradient flow dynamics (1.13), and selecting $w_t = 0$ so that μ is constant, the EVI (EVI) becomes

$$\begin{aligned} \int \langle \nabla \delta_\rho F[\rho_t](z), z - z' \rangle d\gamma(t, z, z') + \lambda \mathcal{W}_2^2(\rho_t, \mu) + F(\rho_t) &\leq F(\mu), \\ \forall \mu \in \mathcal{P}_2(\mathbb{R}^d), \end{aligned}$$

which is exactly the first-order Wasserstein-2 displacement convexity condition, analogous to (1.6) in the Hilbert space setting.

Convergence via HWI and Bakry-Émery. The EVI framework discussed in the previous section can be used to provide a contractivity estimate

$$d(x(t), y(t)) \leq e^{-\lambda t} d(x(0), y(0)).$$

In the context of analyzing long-time behavior, the following two methods have been used to show convergence in the Wasserstein-2 setting. Recall our two Lyapunov functions from the

Euclidean setting: $V_1(x) = \frac{1}{2} \|x - x^*\|^2$ and $V_2(x) = \frac{1}{2} \|\nabla f(x)\|^2$. The two analogous quantities in the Wasserstein-2 setting are

$$V_1(x) \longleftrightarrow \frac{1}{2} \mathcal{W}_2^2(\rho, \rho^*), \quad V_2(x) \longleftrightarrow D(\rho) := \frac{1}{2} \int \|\nabla \delta_\rho F[\rho](z)\|^2 d\rho(z),$$

where $\rho^* \in \mathcal{P}_2(\mathbb{R}^d)$ is the minimizer of F . The quantity D is referred to as the energy dissipation functional, as it captures the rate at which energy decreases along a gradient flow: specifically, $\frac{d}{dt} F(\rho_t) = -D(\rho_t)$. The two methods are (1) via the HWI inequality, and (2) via the Bakry-Émery approach. We will introduce the HWI inequality and illustrate how it can be used to show uniqueness of a minimizer, convergence to the minimizer in the Wasserstein metric, and the rate of energy dissipation; the Bakry-Émery method requires an assumption of convergence to the minimizer, but avoids smoothness assumptions on the measures that are required for the HWI method⁴. It can be used to show the exponential convergence of D to zero as well as the rate of energy dissipation. For both approaches, the existence of a minimizer ρ^* must be shown separately.

HWI Inequality. We assume $\frac{d^2}{ds^2} F(\rho_s) \geq \lambda \mathcal{W}_2^2(\rho_0, \rho_1)$ for any admissible $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$, with ρ_s a geodesic between ρ_0 and ρ_1 . We also assume that ρ_0, ρ_1 and ρ_s admit densities so that there exists some convex, differentiable $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\rho_1 = \nabla \varphi_\# \rho_0$. Following similar arguments as in [10] and [54, Section 5] and using the geodesic equations (1.11),

$$\left. \frac{d}{ds} F(\rho_s) \right|_{s=0} = \int \partial_s \rho_s(z) \partial_\rho F[\rho_s](z) dz \Big|_{s=0} = \int \langle w_0, \nabla \delta_\rho F[\rho_0](z) \rangle d\rho_0(z),$$

where the second equality follows by integration by parts. Using the integral Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} \left. \frac{d}{ds} F(\rho_s) \right|_{s=0} &\geq - \left(\sqrt{\int \|w_0(z)\|^2 d\rho_0(z)} \right) \left(\sqrt{\int \|\nabla \delta_\rho F[\rho_0](z)\|^2 d\rho_0(z)} \right) \\ &= -\mathcal{W}_2(\rho_0, \rho_1) \sqrt{D(\rho_0)}. \end{aligned} \quad (1.14)$$

We compute a Taylor expansion of $F(\rho_s)$ when considered as a function in s , using the exact form of the remainder:

$$F(\rho_1) = F(\rho_0) + \left. \frac{d}{ds} F(\rho_s) \right|_{s=0} + \int_0^1 (1-t) \left(\left. \frac{d^2}{ds^2} F(\rho_s) \right|_{s=t} \right) dt.$$

The velocity field at $s = 0$ is given by $w_0(z) = \nabla \varphi(z) - z$ where $\rho_1 = \nabla \varphi_\# \rho_0$, with $\nabla \varphi$ an optimal Brenier map, which exists when at least one of ρ_0, ρ_1 admits a density; see [10, eq 3.1] for details. Therefore

$$\int \|w_0(z)\|^2 d\rho_0(z) = \int \|\nabla \varphi(z) - z\|^2 d\rho_0(z) = \mathcal{W}_2^2(\rho_0, \rho_1).$$

⁴See [11, Section 2] for HWI and log-Sobolev arguments “fraught with perils of nonsmoothness, infinite dimensions, and degenerate convexity.”

Using the convexity bound $\frac{d^2}{ds^2}F(\rho_s) \geq \lambda \mathcal{W}_2^2(\rho_0, \rho_1)$ and the estimate (1.14) for $\frac{d}{ds}F(\rho_s)|_{s=0}$, we have the HWI inequality

$$F(\rho_0) - F(\rho_1) \leq \mathcal{W}_2(\rho_0, \rho_1)\sqrt{D(\rho_0)} - \frac{\lambda}{2} \mathcal{W}_2^2(\rho_0, \rho_1). \quad (1.15)$$

The HWI inequality immediately implies uniqueness of minimizers for F in the set $\{\rho \in \mathcal{P} : D(\rho) < +\infty\}$. If ρ_0 is such that $D(\rho_0) = 0$, then for any minimizer ρ_1 in the above set we have $F(\rho_0) \leq F(\rho_1)$ with equality if and only if $\mathcal{W}_2(\rho_0, \rho_1) = 0$. Hence, ρ_1 cannot be a minimizer unless $\rho_0 = \rho_1$. We conclude that there is at most one minimizer, and it has zero dissipation. Selecting $\rho = \rho_0$ and $\rho_1 = \rho^*$ where ρ^* is the minimizer of F (if it exists), the HWI inequality becomes

$$\begin{aligned} F(\rho) - F(\rho^*) &\leq \mathcal{W}_2(\rho, \rho^*)\sqrt{D(\rho)} - \frac{\lambda}{2} \mathcal{W}_2^2(\rho, \rho^*)^2 \\ &\leq \max_{t \geq 0} \left(\sqrt{D(\rho)}t - \frac{\lambda}{2} t^2 \right) = \frac{D(\rho)}{2\lambda}. \end{aligned}$$

This results in the generalized Log Sobolev inequality, with $F(\rho | \rho^*) := F(\rho) - F(\rho^*)$,

$$D(\rho) \geq 2\lambda F(\rho | \rho^*). \quad (1.16)$$

Alternatively, by setting $\rho_0 = \rho^*$ and $\rho_1 = \rho$, then $F(\rho^*) < \infty$ and $D(\rho^*) = 0$, and the HWI inequality implies the Talagrand inequality:

$$\mathcal{W}_2^2(\rho, \rho^*) \leq \frac{2}{\lambda} F(\rho | \rho^*), \quad (1.17)$$

for any $\rho \in \mathcal{P}_2$ such that $F(\rho) < \infty$. From the Log Sobolev inequality, selecting $\rho = \rho_t$ the solution to (1.13),

$$D(\rho_t) = -\frac{d}{dt}F(\rho_t) \geq 2\lambda F(\rho_t | \rho^*).$$

This implies

$$\frac{d}{dt}F(\rho_t | \rho^*) \leq -2\lambda F(\rho_t | \rho^*),$$

and applying Grönwall's inequality results in

$$F(\rho_t | \rho^*) \leq e^{-2\lambda t} F(\rho_0 | \rho^*),$$

where here ρ_0 is the initial condition of (1.13). We can now use the Talagrand inequality (1.17) to show exponential convergence in the metric

$$\mathcal{W}_2^2(\rho_t, \rho^*) \leq \frac{2}{\lambda} e^{-2\lambda t} F(\rho_0 | \rho^*).$$

The HWI inequality relates the energy (classically denoted by H in the case of the Boltzmann entropy), the metric (classically denoted by W in the case of the Wasserstein-2 metric) and the energy dissipation (classically denoted by I for the Fisher information, if the energy is the entropy) - hence the name HWI.

Bakry-Émery Approach. The Bakry-Émery strategy achieves exponential decay of the energy by showing the exponential convergence of D to zero. It does not, however, result in exponential convergence in $\mathcal{W}_2$, so in this sense, it is a weaker result than the HWI inequality. The computations are similar to those for the HWI inequality, with the main difference being that the derivative of the energy F is taken along solutions to the gradient flow (1.13) instead of geodesics. We will need the second-order convexity condition for the energy along solutions, and sufficient regularity of F in order to carry out the following calculation:

$$\begin{aligned} \frac{d}{dt}F(\rho_t) &= \int \partial_t \rho_t(z) \delta_\rho F[\rho_t](z) dz \\ &= - \int \langle \nabla \delta_\rho F[\rho_t](z), \nabla \delta_\rho F[\rho_t](z) \rangle d\rho_t(z) = -D(\rho_t). \end{aligned}$$

Taking another derivative, we have

$$\begin{aligned} \frac{d^2}{dt^2}F(\rho_t) &= 2 \int \langle \nabla \delta_\rho F[\rho_t](z), \nabla^2 \delta_\rho F[\rho_t](z) \cdot \nabla \delta_\rho F[\rho_t](z) \rangle d\rho_t(z) \\ &+ 2 \iint \langle \nabla \delta_\rho F[\rho_t](z), \nabla_{z,z'}^2 \delta_\rho^2 F[\rho_t](z, z') \cdot \nabla \delta_\rho F[\rho_t](z') \rangle d\rho_t(z) d\rho_t(z') \\ &\geq \lambda \int \|\nabla \delta_\rho F[\rho_t](z)\|^2 d\rho_t(z) = 2\lambda D(\rho_t), \end{aligned}$$

where $\nabla_{z,z'}^2$ denotes $\nabla_z \nabla_{z'}$ and where the inequality follows from the convexity properties of F ; see [18, Appendix A.1] for convexity calculation details in a more general setting or [10] for a particular choice of energy. This results in an inequality showing the decay of D ,

$$\frac{d^2}{dt^2}F(\rho_t) = -\frac{d}{dt}D(\rho_t) \geq 2\lambda D(\rho_t) \quad \Rightarrow \quad \frac{d}{dt}D(\rho_t) \leq -2\lambda D(\rho_t),$$

and applying Grönwall's inequality, we have

$$-\frac{d}{dt}F(\rho_t) = D(\rho_t) \leq e^{-2\lambda t} D(\rho_0),$$

where here ρ_0 is again the initial condition of (1.13). Integrating in time from t to T ,

$$F(\rho_t) - F(\rho_T) \leq \frac{1}{2\lambda} [e^{-2\lambda t} - e^{-2\lambda T}] D(\rho_0).$$

Finally, we take $T \rightarrow \infty$ and crucially assume that $\rho_t \rightarrow \rho^*$ as $t \rightarrow \infty$; this assumption is one of the key differences between the HWI and Bakry-Émery approaches. Then

$$F(\rho_t | \rho^*) \leq \frac{1}{2\lambda} e^{-2\lambda t} D(\rho_0).$$

Thus, under strong displacement convexity ($\lambda > 0$) we have shown convergence in F and D . While energy dissipation is used for analysis of the long-time behavior in the single-species case, in the

multispecies setting it is not necessarily true that the energy dissipates, even when the system contracts in the metric. Instead, our analysis uses direct differentiation of $\mathcal{W}_2^2$ along solutions to the coupled flow system.

In the next three sections, we summarize our results surrounding the well-posedness and long-time behavior of multispecies coupled PDE gradient flows. The full results can be found in Chapters 2-6, with theoretical results featured in Chapters 2-4 and applications in Chapters 5-6. The notation used in the introductory sections matches that of the corresponding papers, and all changes in notation are specified. In the first section, we summarize asymptotic analysis results for multispecies coupled Wasserstein-2 gradient flows.

1.1.4 Results: Asymptotics of Coupled Wasserstein-2 Multispecies Evolutions

In this section, we discuss the key contributions and results from Chapters 2 and 3, which contain the publications [17, 18]. Multispecies gradient flows arise in many important applications, such as fluid flow, biological systems [56, Ch 8], traffic [48], and increasingly appear in emerging areas like generative modeling [42], opinion dynamics [28], and strategic learning environments [16, 17]. These *multispecies* systems have more than one type of agent or density evolving over time, and each species must be modeled separately. They cannot be analyzed with existing results due to the complex coupling among the different species, and understanding the long-time behavior of the coupled dynamics is a key question in many applications. To address this challenge, we provide conditions under which we determine the existence of and convergence to a unique steady state. From now on, the index i denotes quantities for the i^{th} species. The general dynamics we consider are, where $\rho_i \in \mathcal{P}_2(\mathbb{R}^{d_i})$,

$$\partial_t \rho_i = \operatorname{div}_{x_i} (\rho_i v_i[\rho_1, \dots, \rho_n]) \quad \forall i \in [1, \dots, n]. \quad (1.18)$$

Each species ρ_i evolves according to a velocity field v_i , which depends on all other species. The mass of each species is conserved due to the divergence structure of (1.18), and we normalize each ρ_i to have total mass one. Let $\rho = [\rho_1, \dots, \rho_n]$ and $d = \sum_{i=1}^n d_i$. An important case, which we highlight in this section, is the coupled gradient flow case, where $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ and the dynamics are

$$\left. \begin{aligned} \partial_t \rho_1 &= -\nabla_{\mathcal{W}_2, \rho_1} F_1(\rho_1, \dots, \rho_n) \\ &\vdots \\ \partial_t \rho_n &= -\nabla_{\mathcal{W}_2, \rho_n} F_n(\rho_1, \dots, \rho_n) \end{aligned} \right\}, \quad (1.19)$$

where $-\nabla_{\mathcal{W}_2, \rho_i} = \operatorname{div} (\rho_i \nabla_{x_i} \delta_{\rho_i} F_i[\rho_1, \dots, \rho_n])$. Let ρ_{-i} denote the set of all species with the distribution of species i removed: $\rho_{-i} := [\rho_1, \dots, \rho_{i-1}, \rho_{i+1}, \dots, \rho_n]$. We use the notation $F_i(\mu_i, \rho_{-i}) := F_i(\rho_1, \dots, \rho_{i-1}, \mu_i, \rho_{i+1}, \dots, \rho_n)$. Although each species is evolving according to

a gradient flow with energy $F_i(\cdot, \rho_{-i})$, because the non- i^{th} species are also evolving, $F_i(\rho)$ may or may not be decaying in time. Brackets around the inputs of F are used to indicate mapping to $\mathbb{R}^{d_i}$, for example, $\nabla_{x_i} \delta_{\rho_i} F_i[\rho_1, \dots, \rho_n] : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ and parentheses are used otherwise. This system also has a game theoretic interpretation where each player i has a cost functional F_i which it seeks to minimize. We will prove conditions under which the dynamics (1.19) will converge to a unique steady state, which is also the unique Nash equilibrium of the game $F = (F_i)_{i=1}^n$. Denote the product space $\overline{\mathcal{P}}_2 = \mathcal{P}_2(\mathbb{R}^{d_1}) \times \dots \times \mathcal{P}_2(\mathbb{R}^{d_n})$ and $\overline{W}$ the joint metric over all species, given by $\overline{W}(\rho, \mu)^2 = \sum_{i=1}^n \mathcal{W}_2^2(\rho_i, \mu_i)$. Consider energy functionals $F_i : \mathcal{A} \subseteq \overline{\mathcal{P}}_2 \rightarrow \mathbb{R}$, for a suitable choice of domain $\mathcal{A} \subseteq \overline{\mathcal{P}}_2$. We present the standard definition of a Nash equilibrium in terms of the action ρ_i and objective F_i of each player.

Definition 1.1 (Nash Equilibrium). *A set of measures $\rho^* \in \mathcal{A} \subseteq \overline{\mathcal{P}}_2$ is a Nash equilibrium over $\mathcal{A}$ if,*

$$F_i(\rho_i^*, \rho_{-i}^*) \leq F_i(\rho_i, \rho_{-i}^*) \quad \forall \rho_i \in \mathcal{A}_i, \quad \forall i \in [1, \dots, n]. \quad (1.20)$$

When each species ρ_i is evolving according to a gradient flow with respect to $F_i(\cdot, \rho_{-i})$, some critical differences for the multispecies system compared to a single species system are:

1. The joint system is not a gradient flow. That is, there is not necessarily a single energy functional which dissipates in direction of steepest descent.
2. The candidate for the steady state is a Nash equilibrium rather than a minimizer.
3. Strong geodesic convexity of each $F_i(\cdot, \rho_{-i})$ is not, in general, sufficient for an n -species system to converge to a unique Nash equilibrium.

We start with a special case, a two-species setting in which the two species either cooperate or compete. For these cases, we show that strong convexity is indeed sufficient due to the special structure, even without a joint gradient flow structure. We then discuss the general n -species setting, and introduce *monotonicity*, a game-theoretic concept which allows us to show convergence to a unique steady state, both for coupled gradient flows (of the form (1.19)) and coupled flows not necessarily generated by energy functionals (of the form (1.18)).

Two Species Min-Max (Zero Sum). In the *cooperative* setting, both species minimize the same energy functional – therefore one can view the resulting dynamics as a game in which both players aim to achieve the same objective of minimizing the same energy, resulting in a joint gradient flow. In the *competitive* setting, one player aims to maximize a function which the other player aims to minimize, resulting in gradient descent-ascent dynamics due to the zero-sum game structure. This

setting does not have a joint gradient flow structure. Let the energy in the cooperative setting be defined as $F_a : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \rightarrow \mathbb{R}$ and in the competitive setting as $F_c : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \rightarrow \mathbb{R}$,

$$F_a(\rho_1, \rho_2) = \iint f(z, x) d\rho_1(z) d\rho_2(x) + \mathcal{R}(\rho_1) + \mathcal{U}(\rho_2), \quad (1.21a)$$

$$F_c(\rho_1, \rho_2) = \iint f(z, x) d\rho_1(z) d\rho_2(x) - \mathcal{R}(\rho_1) + \mathcal{U}(\rho_2), \quad (1.21b)$$

where

$$\begin{aligned} \mathcal{R}(\rho_1) &= \alpha \mathcal{H}(\rho_1) + \frac{1}{2} \int W_1 * \rho_1(z) d\rho_1(z) + \int V_1(z) d\rho_1(z), \\ \mathcal{U}(\rho_2) &= \beta \mathcal{H}(\rho_2) + \frac{1}{2} \int W_2 * \rho_2(x) d\rho_2(x) + \int V_2(x) d\rho_2(x), \end{aligned}$$

with $\alpha, \beta \geq 0$. Here, we denote by $f : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$ the function governing coupling forces between the species ρ_1 and ρ_2 , by $\mathcal{H} : \mathcal{P}(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ the entropy functional

$$\mathcal{H}(\eta) = \begin{cases} \int \eta \log \eta & \text{if } \eta \ll \mathcal{L}^{d_i} \\ +\infty & \text{otherwise} \end{cases},$$

for $\mathcal{L}^{d_i}$ the d_i -dimensional Lebesgue measure, by $V_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ external potentials, and by $W_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ interaction potentials. The dynamics in the cooperative setting are

$$\partial_t \rho_1 = -\nabla_{\mathcal{W}_2, \rho_1} F_a(\rho_1, \rho_2), \quad \partial_t \rho_2 = -\nabla_{\mathcal{W}_2, \rho_2} F_a(\rho_1, \rho_2). \quad (1.22)$$

In the competitive case, the dynamics are

$$\partial_t \rho_1 = \nabla_{\mathcal{W}_2, \rho_1} F_c(\rho_1, \rho_2), \quad \partial_t \rho_2 = -\nabla_{\mathcal{W}_2, \rho_2} F_c(\rho_1, \rho_2). \quad (1.23)$$

The analysis of the dynamics in the cooperative setting (1.22) proceeds similarly to the approach in [10] for the single species setting, in which an HWI inequality holds under sufficient convexity assumptions and log-Sobolev and Talagrand inequalities follow, as outlined in Section 1.1.3. However, because the dynamics in the competitive setting (1.23) no longer have a gradient flow structure, the classical gradient flow techniques no longer apply. In the competitive setting, gradient descent dynamics for each player result in convergence to the unique Nash equilibrium under sufficient convexity conditions, which we show using tools from analysis, game theory, and calculus of variations.

Theorem 1.2. (Convergence, informal) Suppose F_c is λ_c displacement concave in ρ_1 and λ_c displacement convex in ρ_2 and is sufficiently smooth with $\alpha, \beta > 0$. Then the following hold:

(a) *There exists a unique critical point $(\rho_1^*, \rho_2^*) = \gamma^* \in \mathcal{P}_2 \times \mathcal{P}_2$ for F_c which is also a steady state for equation (1.23) and the unique Nash equilibrium.*

(i) $\rho_1^* \in L_+^1(\mathbb{R}^{d_1}) \cap C^2(\mathbb{R}^{d_1})$ and $\text{supp}(\rho_1^*) = \mathbb{R}^{d_1}$.

(ii) $\rho_2^* \in L_+^1(\mathbb{R}^{d_2}) \cap C^2(\mathbb{R}^{d_2})$ and $\text{supp}(\rho_2^*) = \mathbb{R}^{d_2}$.

(b) *The solution $\gamma_t := (\rho_1(t), \rho_2(t))$ to the dynamics (1.23) is in $\mathcal{P}_2 \times \mathcal{P}_2$, has uniformly bounded second moments,*

$$\exists K > 0 : \quad \int \|z\|^2 d\rho_1(t, z) + \int \|x\|^2 d\rho_2(t, x) \leq K \quad \forall t \geq 0,$$

and converges exponentially fast in $\overline{W}$ with rate λ_c ,

$$\overline{W}(\gamma_t, \gamma_*) \leq e^{-\lambda_c t} \overline{W}(\gamma_0, \gamma_*).$$

The full statement is in Theorem 2.14. This theorem tells us that the distributions $\rho_1(t)$ and $\rho_2(t)$ converge at a rate corresponding to the displacement concavity-convexity of the energy functional; we expect this to be true from similar analysis of finite-dimensional zero-sum games [13]. To prove the existence of a steady state of the dynamics (1.23), we first show that the second moments of $\rho_1(t)$ and $\rho_2(t)$ are uniformly upper-bounded; then we show contraction in the joint metric, which allows us to conclude the existence of a steady state. Convergence follows from differentiating the Wasserstein-2 distance between the current pair of distributions and the steady state. Then we prove that the steady state is also a critical point and the unique Nash equilibrium using properties of the Euler-Lagrange conditions.

We also prove timescale-separate results for the competitive setting, specifically for a strategic distribution shift model. Timescale separation means one species updates more quickly than the other, and this can change the trajectories for $\rho_1(t)$ and $\rho_2(t)$. This setting is discussed in greater detail in Section 1.1.6.

Multispecies Coupled Flows. We now move back to the general n -species setting and summarize our results in Chapter 3. The energy functionals (F_i) do not necessarily need to have a cooperative or competitive structure. One of the key contributions in analyzing general multispecies coupled gradient flows is our novel definition of *monotonicity*. In finite-dimensional game theory, monotonicity is a condition on cost functions $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, or for dynamical systems, a contractivity condition, or for variational inequalities, a condition ensuring existence of solutions. In the Euclidean setting, recall that player i optimizes $f_i(\cdot, x_{-i})$ over all $x_i \in \mathbb{R}^{d_i}$, with $x = [x_1, \dots, x_n] \in \mathbb{R}^d$. Then λ -monotonicity of (f_i) is defined as

$$\sum_{i=1}^n \langle \nabla_{x_i} f_i(x) - \nabla_{x_i} f_i(y), x - y \rangle \geq \lambda \|x - y\|^2 \quad \forall x, y \in \mathbb{R}^d.$$

As before, we will use this finite-dimensional notion of monotonicity as inspiration for the measure-valued formulation. One can view $\nabla_{x_i} f_i$ as a velocity field, which, in the measure-valued setting, will depend on the measure itself.

The new monotonicity definition uses the following notation. Each species optimizes over $\rho_i \in \mathcal{P}(\mathbb{R}^{d_i})$, which evolves according to a velocity field $v_i[\rho] : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^{d_i}$, which can depend on all species, denoted $\rho = [\rho_1, \dots, \rho_n]$. Denote by $\Gamma_i(\rho_i^{(0)}, \rho_i^{(1)})$ the set of couplings whose marginals are $\rho_i^{(0)}, \rho_i^{(1)} \in \mathcal{P}_2(\mathbb{R}^{d_i})$, and let

$$\Gamma(\rho^{(0)}, \rho^{(1)}) := \Gamma_1(\rho_1^{(0)}, \rho_1^{(1)}) \times \dots \times \Gamma_n(\rho_n^{(0)}, \rho_n^{(1)}).$$

Let Γ^* be the set of optimal transport plans, where

$$\Gamma_i^* = \left\{ \gamma_i : \gamma_i \in \arg \min_{\hat{\gamma}_i \in \Gamma_i(\rho_i^{(0)}, \rho_i^{(1)})} \int \|x_i - y_i\|^2 d\hat{\gamma}_i(x_i, y_i) \right\}.$$

Recall that $\overline{W}$ is the joint metric over all species, given by $\overline{W}(\rho, \mu)^2 = \sum_{i=1}^n \mathcal{W}_2(\rho_i, \mu_i)^2$. We say that a set $\mathcal{A} \in \overline{\mathcal{P}}_2$ is *displacement convex* if for any $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$, the displacement interpolants remain in $\mathcal{A}$ for all $s \in [0, 1]$, that is, $[\rho_1^{(s)}, \dots, \rho_n^{(s)}] \in \mathcal{A}$ for all $s \in [0, 1]$.

Definition 1.3 (Monotonicity). *Let $\mathcal{A} \subseteq \overline{\mathcal{P}}_2$ be a displacement convex set. If for any $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$,*

$$\begin{aligned} & - \sum_{i=1}^n \int \langle x_i - y_i, v_i[\rho^{(0)}](x_i) - v_i[\rho^{(1)}](y_i) \rangle d\gamma_i(x_i, y_i) \\ & \geq \lambda \sum_{i=1}^n \int \|x_i - y_i\|^2 d\gamma_i(x_i, y_i) = \lambda \overline{W}_2^2(\rho^{(0)}, \rho^{(1)}), \end{aligned} \tag{1.24}$$

for all $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$, then the velocity fields $\{v_i : \mathcal{A} \rightarrow \mathbb{R}^{d_i}\}_{i=1}^n$ are λ -monotone in $\mathcal{A}$. Similarly, we say the energy functionals $\{F_i : \mathcal{A} \rightarrow \mathbb{R}\}_{i=1}^n$ are λ -monotone in $\mathcal{A}$ if $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ are λ -monotone.

The measure-valued definition of monotonicity has a similar inner product structure as the finite-dimensional version; this can be interpreted by recognizing that, formally, the tangent space of $\mathcal{W}_2$ at ρ_i is a subset of the $L^2(\rho_i)$ Hilbert space, and the inner product is over objects in that space. The inner product is integrated over the optimal plan with marginals given by the two test point $\rho_i^{(0)}$ and $\rho_i^{(1)}$, and the lower bound depends on the $\mathcal{W}_2$ distance between $\rho_i^{(0)}$ and $\rho_i^{(1)}$. The negative sign that appears on the left-hand side of the condition results from the gradient structure of $\mathcal{W}_2$ having a negative divergence: $\nabla_{\mathcal{W}_2, \rho_i} F_i[\rho] = -\text{div}(\rho_i \nabla_{x_i} \delta_{\rho_i} F_i[\rho])$. Lastly, one can show that the two-player competitive and cooperative games discussed previously satisfy the monotonicity condition.

A set of actions ρ^* is a Nash equilibrium when for each player i , all other players' actions are fixed and the i^{th} player has no incentive to deviate from its current strategy (see Definition 1.1). Our notion of monotonicity can be seen as a generalization of monotonicity for *mixed* Nash equilibria, that is, equilibria in the space of measure-valued actions.

Theorem 1.4 (Convergence, informal). *Let v be λ -monotone in $\mathcal{A}$, with $\lambda > 0$, and $\rho(t) \in \mathcal{A}$ evolve according to (1.18) for all $t \geq 0$. Then*

- (a) *there exists a unique steady state ρ^∞ of (1.18), and $\rho(t)$ converges exponentially fast to ρ^∞ in $\overline{W}$ with rate λ , and*
- (b) *if $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ for all $i \in [1, \dots, n]$, and the energies F_i are lower semicontinuous, then ρ^∞ is the unique Nash equilibrium in $\mathcal{A}$ for the game specified by F .*

This theorem tells us that monotonicity with $\lambda > 0$ results in existence of and exponential convergence to the unique steady state of the dynamics (1.18). If the dynamics have the particular form of a coupled gradient flow system, then the steady state is the unique Nash equilibrium of the corresponding game. For the full statement of this result, see Theorem 3.17.

Monotonicity and Convexity. There is a close relationship between monotonicity and displacement convexity. One can view monotonicity as an extension of convexity to a multispecies system. Strong ($\lambda > 0$) monotonicity of (F_i) in $\mathcal{A}$ implies that each $F_i(\cdot, \rho_{-i})$ is λ -displacement convex for all $\rho_{-i} \in \mathcal{A}_{-i}$, but convexity of each F_i with respect to ρ_i does *not* imply λ -monotonicity. If each ρ_i evolves independently of the other species, then monotonicity and convexity are equivalent. The coupling structure between two species may decrease or destroy the contractivity of the entire system without changing the individual displacement convexity. Hence, displacement convexity is in general not enough to ensure convergence to a unique steady state. These properties are proven and illustrated in examples in Chapter 3.

In the next section, we summarize new results proving existence of solutions of a general evolution equation, with coupled gradient flows in general metric spaces as a specific example. Choosing the Wasserstein-2 metric as a dissipation structure results in the coupled Wasserstein-2 gradient flows discussed in this section.

1.1.5 Results: Well-Posedness of Coupled Multispecies Evolutions

In this section, we discuss the results from Chapter 4 surrounding the existence of solutions for coupled gradient flows; the full paper provides results for the well-posedness of dissipative flows

more generally. We would like to prove existence of solutions to systems of the form

$$\frac{d}{dt}x^{(i)} = -\nabla_{d_i, x^{(i)}} F^{(i)}(x^{(1)}, \dots, x^{(n)}), \quad \forall i \in [1, \dots, n], \quad (1.25)$$

where $\nabla_{d_i, x^{(i)}}$ denotes the steepest direction of $F^{(i)} : D \subset X \rightarrow \mathbb{R}$ in the metric d_i with respect to $x^{(i)}$. In what follows, we denote by (X, d_X) the joint metric space over all species, where $d_X = \sqrt{\sum_{i=1}^n d_i^2}$ and $X = \prod_{i=1}^n X^{(i)}$. For the case where $d_i = \mathcal{W}_2$ for all i , we recover the coupled Wasserstein-2 gradient flow system discussed in the previous section. As discussed in Section 1.1.2 for the single-species setting, taking a limit in the minimizing movement scheme (1.7) (also called the JKO scheme) which solves an approximate time discretization is one way to prove existence of solutions. In Chapter 4, we propose a new discrete time scheme for the evolution of (1.25), denoting $X_\tau^{(i),k}$ the discrete-time approximation of (1.25) for the i^{th} species at time k , with time step size τ . From a discrete-time gradient flow perspective, there are two natural options for how each species can update to move in the direction of steepest descent of $F^{(i)}$. One option is minimizing $F^{(i)}(\cdot, X_\tau^{(-i),k-1})$ with a distance penalty: for every $i \in [1, \dots, n]$,

$$X_\tau^{(i),k} = \arg \min_{x^{(i)} \in X^{(i)}} \left\{ F^{(i)}(x^{(i)}, X_\tau^{(-i),k-1}) + \frac{1}{2\tau} d_i^2(x^{(i)}, X_\tau^{(i),k-1}) \right\}.$$

This update is implicit in $x^{(i)}$ but explicit in $x^{(-i)}$ because of the dependence on $X_\tau^{(-i),k-1}$, so we call this scheme "partially implicit." The second option is to solve

$$X_\tau^k \text{ is a Nash equilibrium in } D \text{ of } \left\{ x \mapsto F^{(i)}(x^{(i)}, x^{(-i)}) + \frac{1}{2\tau} d_i^2(x^{(i)}, X_\tau^{(i),k-1}) \right\}_i, \quad (\text{VMS})$$

which is a fully implicit update. Because the Nash equilibrium problem can be equivalently posed as variational inequality solution, we call this update scheme a *variational movement scheme* (VMS). Proving the existence of minimizers for the partially-implicit update is easier than proving existence of a Nash equilibrium for a set of energies. On the other hand, we expect better stability properties from the fully implicit method (VMS). To illustrate the lack of stability of the first method, consider the min-max game where $X^{(1)} = X^{(2)} = \mathbb{R}$ equipped with the Euclidean norm, and $F^{(1)}(x^{(1)}, x^{(2)}) = -F^{(2)}(x^{(2)}, x^{(1)}) = \frac{\beta}{2}x^{(1)2} + \alpha x^{(1)}x^{(2)} - \frac{\beta}{2}x^{(2)2}$ with $\alpha \in \mathbb{R}$ and $\beta > 0$. The corresponding continuous-time gradient flow solution (1.25) converges exponentially to the Nash equilibrium $(0, 0)$ with rate β . The schemes via the partially implicit and the Nash update are

$$\begin{aligned} \text{(partially implicit)} \quad X_\tau^k &= (1 + \tau\beta)^{-1} \begin{bmatrix} 1 & -\tau\alpha \\ \tau\alpha & 1 \end{bmatrix} X_\tau^{k-1}, \\ \text{(fully implicit)} \quad X_\tau^k &= ((1 + \tau\beta)^2 + \tau^2\alpha^2)^{-1} \begin{bmatrix} 1 + \tau\beta & -\tau\alpha \\ \tau\alpha & 1 + \tau\beta \end{bmatrix} X_\tau^{k-1}. \end{aligned}$$

An eigenvalue calculation for each system, $\sqrt{\alpha^2\tau^2 + 1}(\beta\tau + 1)^{-1}$ and $[\tau^2(\alpha^2 + \beta^2) + 2\beta\tau + 1]^{-1/2}$, reveals that if the coupling term α is larger than the convexity parameter β , the first system is exponentially stable only for a small enough τ , dependent on α and β , whereas the second system is exponentially stable for all $\tau > 0$, $\alpha \in \mathbb{R}$, and $\beta > 0$. The explicit Euler scheme $X_\tau^k = X_\tau^{k-1} - \tau \nabla_x f(X_\tau^{k-1})$ for gradient descent in Euclidean space requires an upper-Lipschitz bound L on the gradient of the energy $\nabla_x f(x)$ for the stability of gradient descent, along with a small enough time step τ relative to the Lipschitz constant L . We observe a similar phenomenon with the partially implicit method requiring a small enough τ dependent on α .

To avoid the requirement of a Lipschitz bound on the energy functionals, we will use the fully implicit scheme (VMS). Proving that a solution exists to this scheme is nontrivial, which we discuss next.

Existence of Discrete VMS Solutions. A number of challenges arise in proving existence of solutions to (VMS). Let the bifunctions $b : D \times D \rightarrow \mathbb{R}$ and $B : D \times D \rightarrow \mathbb{R}$ be given by

$$b(y, x) = \sum_{i=1}^n F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}), \quad \text{for all } x, y \in D,$$

$$B(\tau, \bar{x}; y, x) = \frac{1}{2\tau} d_X^2(\bar{x}, y) - \frac{1}{2\tau} d_X^2(\bar{x}, x) + b(y, x),$$

where $D \subset X$ is the domain for the arguments of b . The fully implicit update (VMS) has solution X_τ^k if it satisfies the saddle point condition: for every $(x, y) \in D \times D$,

$$B(\tau, X_\tau^{k-1}; X_\tau^k, x) \leq \underbrace{B(\tau, X_\tau^{k-1}; X_\tau^k, X_\tau^k)}_{=0} \leq B(\tau, X_\tau^{k-1}; y, X_\tau^k).$$

If the domain D is a bounded set and the metric space has some linear structure, one can hope to apply a fixed point theorem to show existence of the saddle point (X_τ^k, X_τ^k) . However, many metric spaces of interests do not have such a linear structure. Alternatively, if there is some contraction for the map $(x, y) \mapsto \arg \max_{x \in D} \arg \min_{y \in D} B(\tau, X_\tau^{k-1}; y, x)$, a unique solution could be shown; in finite dimensions, this requires a Lipschitz bound for $\nabla_{x_j} f^{(i)}$ which is not easily available in many metric spaces. From the finite-dimensional game theory literature we expect such an equilibrium point to exist under *strong monotonicity*, which is a first-order (gradient) condition on the energy functionals that guarantees contraction of coupled gradient descent dynamics. However, in general metric spaces, we prefer a zeroth-order notion of monotonicity, since the appropriate notion of gradient is difficult to determine. To address these challenges, we employ a notion of convexity that provides a linear structure, *barycentric convexity*, and define *interaction dissipativity*, which in combination with barycentric convexity, describes the overall contractivity of the dynamics.

Barycentric convexity is a generalization of convexity from a two-point interpolation to a J point interpolation. Let $\mathbf{x} : \triangle_J \rightarrow D$ be a barycentric map from the $J - 1$ dimensional simplex




Figure 1.1: A barycentric map $\mathbf{x}$ maps from the simplex $\Delta_J = \Delta_3$ to the space $X = \mathbb{S}^2$. The simplex vertices $(e_j^i)_j$ map to points in X , and interpolations in Δ_J map to interpolations in X .

to some subset D of X . This map defines an interpolation among J points (x_j) in D , where $x_j = \mathbf{x}(e_j^j) =: \mathbf{x}_j \in D$ and $(e_j^j)_j \in \mathbb{R}^{d_J}$ is the canonical basis of $\mathbb{R}^{d_J}$. A system $\mathcal{B}$ is a collection of barycentric maps $\mathbf{x}$ for all $J \in \mathbb{N}_{\geq 2}$.

Definition 1.5 (Barycentrically Convexlike Function). *The functional $\varphi : X \rightarrow (-\infty, +\infty]$ is barycentrically κ -convexlike with respect to a system $\mathcal{B}$ of barycentric maps if for every integer $J \in \mathbb{N}_{\geq 2}$, every $\mathbf{x} \in \mathcal{B}$, and every $\alpha \in \Delta_J$,*

$$\varphi(\mathbf{x}(\alpha)) \leq \sum_{j=1}^J \alpha_j \varphi(\mathbf{x}_j) - \frac{\kappa}{4} \sum_{i,j=1}^J \alpha_i \alpha_j d_X^2(\mathbf{x}_i, \mathbf{x}_j). \quad (1.26)$$

For example, in a Banach space $X = B$, the system of barycentric maps with $J = 2$ is denoted $\mathcal{B}_2$, and $\mathbf{x} \in \mathcal{B}_2$ can parameterize linear interpolations via $\mathbf{x}(\alpha) = (1 - \alpha)\mathbf{x}_1 + \alpha\mathbf{x}_2$. Here $\alpha \in \Delta_2$ is described by a single parameter α : $\alpha = [1 - \alpha, \alpha]$. In a geodesic space, we can define $\mathcal{G}_2$ as a particular choice of $\mathcal{B}_2$ such that $\mathbf{x}(\alpha)$ satisfies $d_X(\mathbf{x}_1, \mathbf{x}(\alpha)) = \alpha d_X(\mathbf{x}_1, \mathbf{x}_2)$ for all $x \in \mathcal{G}_2$, which means that $\mathcal{G}_2$ is the set of geodesic interpolants. Then the inequality (1.26) corresponds to geodesic κ -convexity for φ . Next, we introduce interaction dissipativity.

Definition 1.6 (Interaction Dissipativity). *A bifunction $b : D \times D \rightarrow \mathbb{R}$ is η interaction dissipative if*

$$b(x, y) + b(y, x) \leq \eta d_X^2(y, x) \quad \text{for every } x, y \in D.$$

The coefficient η parameterizes the interaction between the first and second entries in b ; in the multispecies coupled gradient flow setting, this corresponds to the interactions among all the different species through their energy functionals. We introduce a new definition for games in general metric spaces, using the interaction dissipativity condition in combination with barycentric κ -convexity of $y \mapsto b(y, x)$.

Definition 1.7 (Monotonicity in a Metric Space). *The energies $(F^{(i)})$ are $-\lambda$ -monotone with respect to a barycentric system $\mathcal{B}$ with $\lambda = \eta - \kappa$ if*

- (κ convexity) *For every $y \in \mathcal{D}$, $y \mapsto b(y, x)$ is barycentrically κ -convexlike with respect to $\mathcal{B}$.*
- (η interaction dissipativity) *For every $x, y \in \mathcal{D}$, $b(y, x) + b(x, y) \leq \eta d_X^2(x, y)$.*

Equivalently, each $F^{(i)}(\cdot, \cdot) - F^{(i)}(x^{(i)}, \cdot)$ is barycentrically κ -convexlike with respect to $\mathcal{B}$ and, for every $x, y \in \mathcal{D}$,

$$\sum_{i=1}^N F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}) + F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) \leq \eta d_X^2(x, y).$$

We show that, for particular choices of $\mathcal{B}$, this notion of (game-theoretic) monotonicity implies the classical definitions in Euclidean space and in the space of measures equipped with the Wasserstein-2 metric, as defined in [18, 69]. In [18] and Chapter 3, monotonicity is defined with respect to geodesic curves, whereas in [69] it is defined along linear interpolations of the measures. Using different choices of $\mathcal{B}$ allows for the zeroth order definition to imply each of these different first-order conditions. However, this zeroth order Definition 1.7 is stronger than classical monotonicity in Euclidean spaces; it is, in some sense, an upper bound on $(F^{(i)})$ whereas the classical definitions are lower bounds on the gradients of $F^{(i)}$. For an example of a set of energies that illustrate this gap, see Remark 3.46.

Equipped with barycentric convexity and interaction dissipativity, we take two approaches for showing existence of the discrete VMS: one via compactness, and one via convexity and interaction dissipativity. The first provides a primal solution to the variational equilibrium problem (VMS), which requires compactness of superlevel sets of $b(\cdot, x)$. The linear structure of the simplex allows us to prove a version of the Knaster Kuratowski Mazurkiewicz (KKM) theorem in the general metric space setting. While this compactness condition is looser than the second approach, it can be difficult to show. The second approach uses the dual formulation of the variational equilibrium problem, and relies on barycentric κ -convexity of $y \mapsto B(\tau, \bar{x}; y, x)$ and η interaction dissipativity, both of which replace the compactness assumption needed in the primal formulation.

Finally, we would like to take $\tau \rightarrow 0$ in (VMS) to obtain a continuous time weak solution to (1.25). To take this limit, we show that the sequence X_τ^k generated by the VMS scheme satisfies a discrete evolution variational inequality (EVI),

$$\frac{1}{2\tau} \left(d_X^2(X_\tau^k, Y) - d_X^2(X_\tau^{k-1}, Y) \right) + \frac{1}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) + \frac{\kappa}{2} d_X^2(X_\tau^k, Y) \leq b(Y, X_\tau^k),$$

for every $Y \in \mathcal{D}$,

which provides the desired weak notion of a discrete-time solution. Taking the discrete EVI and passing to the limit $\tau \rightarrow 0$ will allow us to recover a continuous-time, weak form of (1.25).

Discrete to Continuous in Time Limit. In order to take $\tau \rightarrow 0$, the continuous-time piecewise constant interpolant $\bar{X}_\tau : [0, \infty) \rightarrow X$ is constructed,

$$\bar{X}_\tau(0) := X_\tau^0, \quad \bar{X}_\tau(t) := X_\tau^k \quad \text{if } t \in ((k-1)\tau, k\tau].$$

We show a Cauchy condition for $d_X(\bar{X}_{\tau_1}(t), \bar{X}_{\tau_2}(t))$ in terms of τ_1 and τ_2 , which results in the existence of a limit $x(t)$ for $(\bar{X}_\tau(t))_{\tau \geq 0}$ as $\tau \rightarrow 0$, pointwise in t . This estimate is constructed using the doubling integration method from [52]. Then, taking the limit of the discrete EVI, the continuous EVI is obtained:

$$\frac{d}{dt} d_X^2(y, x(t)) + \frac{\kappa}{2} d_X^2(y, x(t)) \leq b(y, x(t)), \quad \text{for all } t \geq 0, y \in X,$$

The interaction dissipativity inequality in combination with the continuous EVI enables us to show the $\lambda = \eta - \kappa$ dissipativity estimate

$$d_X(y(t), x(t)) \leq e^{\lambda(t-s)} d_X(y(s), x(s)) \quad \text{whenever } 0 \leq s \leq t.$$

When $\lambda < 0$, we have contraction in the metric, and if the metric space is complete, we can show existence of a unique steady state. Note that the sign convention for λ here is the opposite of that in Sections 1.1.3-1.1.4 because dissipativity is the negative of monotonicity. The summary of our results is stated in the following theorem, where in the multispecies gradient flow setting the bifunction b is

$$b(y, x) = \sum_{i=1}^n F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}), \quad \text{for all } x, y \in D.$$

Theorem 1.8 (Existence and Contraction Estimate). *Suppose for some barysystem $\mathcal{B}$ dependent on $\bar{x} \in D$ and τ_o such that $\tau_o^{-1} + \kappa > 2\eta$,*

- *for every $\tau \in (0, \tau_o)$, every $x \in D$, the functions $y \mapsto B(\tau, \bar{x}; y, x)$ given by*

$$B(\tau, \bar{x}; y, x) = \frac{1}{2\tau} d_X^2(\bar{x}, y) - \frac{1}{2\tau} d_X^2(\bar{x}, x) + b(y, x),$$

are barycentrically $(\tau^{-1} + \kappa)$ convexlike with respect to $\mathcal{B}$,

- *$x \mapsto B(\tau, \bar{x}; y, x)$ is upper hemicontinuous with respect to $\mathcal{B}_2$ (denoting $\mathcal{B}$ restricted to $J = 2$) for any $\bar{x}, y \in D$, and*

- b is η interaction dissipative:

$$b(x, y) + b(y, x) \leq \eta d_X^2(y, x) \quad \text{for every } x, y \in D.$$

Then there exists a continuous curve $x : [0, +\infty) \rightarrow D$ that satisfies the continuous EVI

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), y) + \frac{\kappa}{2} d_X^2(x(t), y) \leq b(y, x(t)) \quad t \in (0, +\infty), \quad \text{for every } y \in D.$$

Any curves x, y which satisfy the EVI above also satisfy the contraction estimate

$$d_X(x(t), y(t)) \leq e^{\lambda(t-s)} d_X(x(s), y(s)), \quad \text{for all } t \geq s > 0.$$

If (D, d_X) is complete and $\lambda < 0$, there exists a unique steady state $x_\infty \in D$ and x_∞ is the unique Nash equilibrium of $(F^{(i)})$.

We remark that Theorem 1.8 applies to more general b beyond the coupled gradient flow setting, encompassing a broad class of PDEs; see Chapter 4.

Next, we move to applications of coupled gradient flow systems, which inspire the theoretical results we have shown.

1.1.6 Results: Gradient Flow Applications

In this section, we summarize two important applications of coupled gradient flows that we develop in Chapters 5-6: strategic distribution shift in machine learning and optimal transport. By posing these problems in the coupled gradient flow framework, two benefits arise. First, we gain application insights from analyzing this structure. Second, we obtain new mathematical results about the behavior of these coupled systems. To illustrate the usefulness of gradient flows as a modeling paradigm, we highlight an example from [17].

A study of Colombia census data [8] from 1995 to 2003 investigates how local officials misreported data in order to obtain lower poverty index scores for their constituents. Households with a poverty index score below a given threshold receive government aid, and the algorithm computing the poverty index score was not publicly available. When it was released in 1997, the distribution of scores shifted from a Gaussian-like shape to having a sharp drop-off above the threshold by 1998. In Figure 1.2, a subset⁵ of the poverty index scores [8] are plotted, illustrating a significant distribution shift, incentivized by the threshold for government aid.

In that case, the threshold was fixed and the algorithm was static. Modern machine learning settings are more complex: algorithms and data distributions evolve together, each responding to the other.

⁵The first 20 million census data entries were loaded and filtered to include only those from 1995 and 1998, with 2.4 million and 1.4 million entries respectively.




Figure 1.2: Poverty index scores in 1995 and 1998 are plotted, with the red line the threshold below which households qualify for government aid programs. In summer 1997, information about how the poverty index score is computed from census data, along with the threshold, was made publicly available, and a distribution shift appeared in 1998.

This motivates the study of coupled gradient flow systems, in which one "species" is the data distribution and the other is the algorithm's parameters. The following section introduces the key mathematical tools for analyzing these interacting systems, drawing particularly on concepts from game theory.

Timescale Separation. A key theoretical aspect of each of the applications in this section is the impact of *timescale separation*. Because of the close connection to game theory, we call each species a *player*. Timescale separation is the consideration of the relative update speed of each player, where one player can update much faster or slower than the other. For example, in strategic distribution shift settings, suppose a government agency applies a classifier to a population of people. The strategic population may learn and change their responses much more quickly than the government updates classification policies. The dynamics are given by, for $\beta_1, \beta_2 \in (0, \infty)$,

$$\begin{aligned}\partial_t \rho_1 &= -\beta_1 \nabla_{\mathcal{W}_2, \rho_1} F_1(\rho_1, \rho_2), \\ \partial_t \rho_2 &= -\beta_2 \nabla_{\mathcal{W}_2, \rho_2} F_2(\rho_1, \rho_2).\end{aligned}$$

When $\beta_1 = \beta_2$, the players evolve on a similar timescale. Conversely, consider the case in which ρ_2 updates much faster than ρ_1 . Suppose F_1 is such that when ρ_2 is fixed, the gradient flow with respect to ρ_1 converges to the minimizer of $F_1(\cdot, \rho_2)$, and $F_2(\rho_1, \cdot)$ has a unique minimizer for all $\rho_1 \in \mathcal{P}_2$. As $\beta_2 \rightarrow \infty$, player two updates instantaneously, and in the limit we write the dynamics as

$$\partial_t \rho_1 = -\nabla_{\mathcal{W}_2, \rho_1} F_1(\rho_1, b_2(\rho_1)), \quad b_2(\rho_1) = \arg \min_{\rho_2 \in \mathcal{P}_2(\mathbb{R}^{d_2})} F_2(\rho_1, \rho_2). \quad (1.27)$$

We call $b_2(\rho_1)$ the *best response* of ρ_2 to ρ_1 ; that is, ρ_2 instantly minimizes F_2 given the value of ρ_1 . Player one evolves according to a gradient flow through the best response, $b_2(\rho_1)$. Likewise, in

the limit where $\beta_1 \rightarrow \infty$, the dynamics are

$$b_1(\rho_2) = \arg \min_{\rho_1 \in \mathcal{P}_2(\mathbb{R}^{d_1})} F_1(\rho_1, \rho_2), \quad \partial_t \rho_2 = -\nabla_{\mathcal{W}_2, \rho_2} F_2(b_1(\rho_2), \rho_2). \quad (1.28)$$

Now ρ_2 evolves according to a gradient flow, where ρ_1 is instantly best-responding by minimizing $F_1(\cdot, \rho_2)$. The questions we now ask are

1. Do (1.27)-(1.28) converge to a steady state? At what rate?
2. If so, are the steady states of (1.27) and (1.28) the same, coinciding with the Nash equilibrium of (F_1, F_2) ?

To address the first question, we investigate commonly-used assumptions for proving convergence of a Wasserstein-2 gradient flow: strong displacement convexity and existence of a minimizer. We answer these questions in the particular setting of a zero-sum game, that is, $F_1 = -F_2$. To do so, we crucially make use of the following two results.

Theorem 1.9 (Berge's Maximum Theorem, Thm 16.31 [2]). *Let $\varphi : X \rightarrow Y$ be a continuous set-valued map between topological spaces with nonempty compact values, and suppose $f : \text{Gr } \varphi \rightarrow \mathbb{R}$ is continuous. Define the value function $m : X \rightarrow \mathbb{R}$ by*

$$m(x) = \max_{y \in \varphi(x)} f(x, y),$$

and the set-valued map $\mu : X \rightarrow Y$ of maximizers by

$$\mu(x) = \{y \in \varphi(x) : f(x, y) = m(x)\}.$$

Then

- *the value function m is continuous, and*
- *if Y is Hausdorff, then the "argmax" set-valued map μ is upper hemicontinuous with nonempty compact values.*

Theorem 1.10 (Danskin, Theorem II [22]). *Let $f : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$ be a continuous, differentiable function, and define $\mu(x) = \arg \max_y f(x, y)$. Then*

$$\nabla_x f(x, \mu(x)) = [\nabla_x f(x, y)] \Big|_{y=\mu(x)}.$$

In both applications, we use Berge's maximum theorem to obtain continuity results for b_2 and b_1 that are needed to argue that $F_1(\rho_1, b_2(\rho_1))$ in (1.27) and $F_2(b_1(\rho_2), \rho_2)$ in (1.28) are sufficiently smooth in ρ_1 and ρ_2 respectively. Then we prove infinite-dimensional versions of Danskin's theorem (also known as the envelope theorem) to show convexity properties of $F_1(\rho_1, b_2(\rho_1))$ and $F_2(b_1(\rho_2), \rho_2)$. To do so, we rely on showing Γ -convergence results for b_2 ; see [17] and Chapter 2. Danskin-type results are particularly useful for the sufficient condition for λ -convexity calculations in the Wasserstein-2 metric space, where $\rho_1^{(s)}$ is a geodesic between $\rho_1^{(0)}, \rho_1^{(1)}$,

$$\frac{d^2}{ds^2} F_1(\rho_1^{(s)}, b_2(\rho_1^{(s)})) \geq \lambda \mathcal{W}_2^2(\rho_1^{(0)}, \rho_1^{(1)}), \quad \text{for all } \rho_1^{(0)}, \rho_1^{(1)} \in \mathcal{P}_2(\mathbb{R}^{d_1}).$$

We prove an infinite-dimensional version of Theorem 1.10 to show that

$$\frac{d}{ds} F_1(\rho_1^{(s)}, b_2(\rho_1^{(s)})) = \left(\int \delta_{\rho_1} F_1[\rho_1, \rho_2](x) \partial_s \rho_1^{(s)}(x) dx \right) \Big|_{\rho_2=b_2(\rho_1^{(s)})},$$

avoiding direct computation of $\delta_{\rho_1} b_2(\rho_1)$, which would be difficult since b_2 is defined only implicitly through an argmax function. Equipped with continuity and convexity properties, we obtain convergence results with classical methods. Our results proving Danskin's theorem with a measure-valued functional can be found in Propositions 2.41 and 2.60.

The rates of convergence depend on the slower of the two species; we show that the rate for (1.27) depends on the convexity of F_1 with respect to ρ_1 , and the rate for (1.28) depends on the convexity of F_2 with respect to ρ_2 . Outside of zero-sum games, less is known for the timescale separated case and results may need to be shown on a case-by-case basis.

To answer the second question, a game theory perspective offers a notion of equilibria that is useful for analysis. The limiting game resulting from timescale separation gives an order of play, that is, one player selects a strategy before the other. A Nash equilibrium results from simultaneous selection of actions; when the selection of actions is ordered, the equilibrium is called a *Stackelberg equilibrium*. Consider the case where ρ_2 best-responds to ρ_1 . In this setting, we say player one is the leader (plays first) and player two (plays second) is the follower, since the action that player two selects depends directly on that of player one. The Stackelberg equilibrium of this game is defined as (ρ_1^*, ρ_2^*) , where

$$\rho_1^* = \arg \min_{\rho_1 \in \mathcal{P}_2(\mathbb{R}^{d_1})} F_1(\rho_1, b_2(\rho_1)), \quad \rho_2^* = b_2(\rho_1^*). \quad (1.29)$$

This equilibrium is a steady-state candidate for the dynamics (1.27), where ρ_2 updates infinitely faster than ρ_1 . This may seem counterintuitive since (1.29) implies player one leads whereas in (1.27) ρ_2 updates faster than ρ_1 . However, the slowness of ρ_1 forces the update of ρ_2 before making any updates, in the same way that leading in a one-shot game forces the best response of

the other player. See [73] for a detailed discussion on the relationship between update speed and leader/follower roles.

In a two-player game with unique Nash and Stackelberg equilibria, the Stackelberg equilibrium is never worse than the Nash equilibrium for the leader. To see this, suppose the leader selects $\rho_1^* = \rho_1^{**}$, where $(\rho_1^{**}, \rho_2^{**})$ is the unique Nash equilibrium of (F_i) . Then player two selects ρ_2^{**} , since the Nash equilibrium is unique. By the definition of Stackelberg equilibrium, $F_1(\rho_1^*, b_2(\rho_1^*)) \leq F_1(\rho_1^{**}, b_2(\rho_1^{**})) = F_1(\rho_1^{**}, \rho_2^{**})$, since the leader optimizes over all choices of ρ_1 .

In the zero-sum case where $F_1 = -F_2$, Von Neumann's theorem says that the Stackelberg and Nash equilibrium coincide [63], proven first for finite-dimensional spaces and later extended to infinite-dimensional spaces with linear structure. Although this was not proven in our work, we expect that neither player has an advantage from leading, and the Nash equilibrium is the Stackelberg equilibrium $\rho^{**} = \rho^*$. This outcome occurs in all of our numerical simulations.

A number of interesting open questions remain. Proving convergence of the timescale separated dynamics with finite values of β_1 and β_2 and convergence in settings with non-zero-sum structure have yet to be shown. Additionally, we expect the steady state of non-zero-sum games to depend on the timescale separation; it is not known how to optimally select rates β_1 and β_2 to drive ρ_1 and ρ_2 to a desirable equilibrium.

Next, we describe how strategic distribution shift in response to an algorithm can be modeled in the coupled gradient flow framework, and present timescale separation convergence results.

1.1.6.1 Strategic distribution shift

Machine learning algorithms in real-world settings often update their parameters over time to improve performance; as more data is collected, it is natural for the algorithm to update based on more recent information. However, in common applications the distribution of data on which the algorithm is training may not be stationary over time. This phenomenon is known as *distribution shift*, and is induced from a variety of causes, including mis-aligned incentives, interactions with other agents, and natural causes.

Our model in [16] proposes an energy functional which the algorithm $\mu \in \mathcal{P}_2$ seeks to minimize and which the population $\rho \in \mathcal{P}_2$ aims to maximize; some terms have dependence on both distributions while others model energies specific to the evolution of agents and the algorithm, respectively. Let us denote

$$\begin{aligned} \mathcal{V}(\rho, \mu) &= \iint f_1(z, x) d\rho(z) d\mu(x) + \iint f_2(z, x) d\pi(z) d\mu(x), \\ \mathcal{R}(\rho) &= \frac{1}{2} \int W * \rho(z) \rho(z) + \alpha KL(\rho | \bar{\rho}), \quad \mathcal{Q}(\mu) = \frac{\kappa}{2} \int \|x - x_0\|^2 d\mu(x). \end{aligned}$$

Here, $\tilde{\rho} \in \mathcal{P}(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ and $\pi \in \mathcal{P}(\mathbb{R}^d)$ are fixed reference measures, x_0 is a fixed reference point in the algorithm parameter space, and $W : \mathbb{R}^d \rightarrow \mathbb{R}$ is a given interaction potential. The energy in the strategic distribution shift model is given by

$$F(\rho; \mu) := \mathcal{V}(\rho, \mu) + Q(\mu) - \mathcal{R}(\rho).$$

The functional $\mathcal{V}(\rho, \mu)$ models the cost which the algorithm seeks to minimize, and the population minimizes or maximizes depending on the setting. For example, when μ represents a binary classifier, the distribution π models individuals carrying the true label 1, and the distribution ρ model individuals carrying a true label 0, where 0 and 1 denote the labels of two classes of interest. The term $\int f_1(z, x)\mu(dx)$ represents a penalty for incorrectly classifying an individual with features z with true label 0 when using the classifier $\mu(x)$. Analogously, the term $\int f_2(z, x)\pi(dz)$ is large if x incorrectly classifies the population π that carries the true label 1. The functions f_1 and f_2 can be chosen according to the application; a standard choice for classification problems is the logarithmic loss [51].

The functional $Q(\mu)$ is a regularizer for the algorithm; this penalizes the classifier for selecting extreme learning parameters, and provides convexity for the loss function. The coefficient $\kappa > 0$ parameterizes the strength of the regularizer.

The functional $\mathcal{R}(\rho)$ contains two terms, the Kullback-Leibler divergence (denoted KL), also called the relative entropy, and the interaction term driven by the interaction potential W . The term $\alpha KL(\rho|\tilde{\rho})$ forces the evolution of $\rho(t)$ to approach $\tilde{\rho}$. In other words, it penalizes (in energy) deviations from a given reference measure $\tilde{\rho}$. In many application settings, we take $\tilde{\rho}$ to be the initial distribution $\rho(t = 0)$. The solution $\rho(t)$ then evolves away from $\tilde{\rho}$ over time due to the other forces that are present. Therefore, the term $KL(\rho|\tilde{\rho})$ in the energy both provides smoothing of the flow and a penalization for deviations away from the reference measure $\tilde{\rho}$.

The self-interaction term $W * \rho$ introduces non-locality into the dynamics, as the decision for any given individual to move in a certain direction is influenced by the behavior of all other individuals in the population. The choice of W is application dependent. Very often, the interaction between two individuals only depends on the distance between them. This suggests a choice of W as a radial function, i.e. $W(z) = \omega(|z|)$. A choice of $\omega : [0, \infty) \rightarrow \mathbb{R}$ such that $\omega'(r) > 0$ corresponds to an *attractive* force between individuals, whereas $\omega'(r) < 0$ corresponds to a *repulsive* force.

However, in many real-world applications of this model, either the population or the algorithm may update much faster than the other. For example, many government policies are updated on a much slower timescale compared with how quickly individuals can adjust their response. In settings like advertising algorithms, companies can update their algorithms faster than users will

adjust search strategies or viewing patterns. We provide convergence results for these two settings for competitive objectives, taking the timescale separations to be large enough that one entity instantly minimizes or maximizes its objective while the other entity evolves over time. Here we consider the case $\mu(t) = \delta_{x(t)}$ where $x(t)$ solves $\frac{d}{dt}x(t) = -\nabla_x L(\rho(t), x(t))$ with loss function $L(\rho, x) = \int f_1(z, x) d\rho(z) + \frac{\kappa}{2} \|x - x_0\|^2$. The energy in these two setting is denoted $G(\rho, x)$. The energy functional is given by

$$G(\rho, x) := \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2 - \mathcal{R}(\rho).$$

When the algorithm updates quickly relative to the population, we consider dynamics given by

$$\partial_t \rho = \nabla_{\mathcal{W}_2, \rho} G(\rho, x)|_{x=b(\rho)}, \quad b(\rho) := \underset{\bar{x} \in \mathbb{R}^d}{\operatorname{argmin}} G(\rho, \bar{x}). \quad (1.31)$$

In the opposite setting where the population is fast relative the algorithm, we can consider the population immediately responding to the algorithm, which results in the dynamics

$$\frac{d}{dt}x = -\nabla_x G(\rho, x)|_{\rho=r(x)}, \quad r(x) := \underset{\hat{\rho} \in \mathcal{P}}{\operatorname{argmax}} G(\hat{\rho}, x). \quad (1.32)$$

In this time-scale separated setting, model (1.31) is a dynamic maximization of G with respect to ρ in Wasserstein-2 space, and an instantaneous minimization of G with respect to the algorithm parameters x . Model (1.32) is an instantaneous maximization of G with respect to ρ and a dynamic minimization of G with respect to the algorithm parameters x .

In our results, we assume that G is λ_b displacement convex in ρ and λ_d convex in x . In support of the main results, we show that

$$\begin{aligned} G_b(\rho) &:= G(\rho, b(\rho)) \quad \text{is } \lambda_b \text{ displacement convex,} \\ G_d(x) &:= G(r(x), x) \quad \text{is } \lambda_d \text{ convex.} \end{aligned}$$

The following theorems present long-time behavior results.

Theorem 1.11 (Fast Algorithm). *Suppose $\lambda_b, \lambda_d > 0$ and G is sufficiently smooth. Consider a solution ρ_t to the dynamics (1.31) with initial condition $\rho_0 \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho_0) < \infty$. Then the following hold:*

- (a) *There exists a unique maximizer ρ_∞ of $G_b(\rho)$, which is also a steady state for equation (1.31). Moreover, $\rho_\infty \in L_+^1(\mathbb{R}^d) \cap C(\mathbb{R}^d)$ with the same support as $\tilde{\rho}$.*
- (b) *The solution ρ_t converges exponentially fast to ρ_∞ in $G_b(\cdot | \rho_\infty)$ and $\mathcal{W}_2$,*

$$G_b(\rho_t | \rho_\infty) \leq e^{-2\lambda_b t} G_b(\rho_0 | \rho_\infty) \quad \text{and} \quad \mathcal{W}_2(\rho_t, \rho_\infty) \leq c e^{-\lambda_b t} \quad \text{for all } t \geq 0,$$

where $c > 0$ is a constant only depending on ρ_0, ρ_∞ and the parameter λ_b .

Theorem 1.12 (Fast Population). *Suppose $\lambda_b, \lambda_d > 0$ and G is sufficiently smooth. Then it holds:*

- (a) *There exists a unique minimizer x_∞ of $G_d(x)$ which is also a steady state for (1.32).*
- (b) *The vector $x(t)$ solving the dynamics (1.32) with initial condition $x(0) \in \mathbb{R}^d$ converges exponentially fast to x_∞ in G_d and in the Euclidean norm:*

$$\begin{aligned} \|x(t) - x_\infty\| &\leq e^{-\lambda_d t} \|x(0) - x_\infty\|, \\ G_d(x(t)) - G_d(x_\infty) &\leq e^{-2\lambda_d t} (G_d(x(0)) - G_d(x_\infty)), \end{aligned}$$

for all $t \geq 0$. Moreover, $G_d \in C^1(\mathbb{R}^d)$.

Theorems 1.11 and 1.12 state that we will observe exponential convergence in the timescale-separated cases. The proof of Theorem 1.11 hinges on proving a generalized version of Danskin's Theorem, that is, showing that $\delta_\rho G_b(\rho) = (\delta_\rho G(\rho, x))|_{x=b(\rho)}$. Displacement convexity of G_b follows from this, and then standard HWI techniques apply. In order to prove a similar Danskin's result as an ingredient in the proof of Theorem 1.12, we first show that $G_d(x)$ is differentiable. The best response function $r(x)$ is defined implicitly as the result of a minimization problem, and we employ Γ -convergence tools to obtain the regularity result. As previously discussed, the convergence rate is dependent on the slower of the two species in each case. For details, see Chapter 2.

Numerical Simulations. We apply the above framework to several settings surrounding strategic classification. In comparing our distribution shift model with moment models [51], we find that having access to the full distribution results in better classifier performance, illustrating why distributional modeling is important. We also observe that polarization within a population can occur as an outcome of strategic behavior, and that algorithm learning speed impacts the volatility of the population behavior. This motivates the need for future research in understanding the impacts of learning algorithms on strategic populations.

In the next section, we present a second application that uses coupled gradient flows for solving an optimal transport problem.

1.1.6.2 Min-Max Formulation of the optimal transport Problem

Optimally transporting mass from one distribution $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ to another $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ is an important problem in many applications, ranging from sampling to numerically solving partial differential equations. The Kantorovich optimal transport (OT) problem is

$$\inf_{\rho \in \Gamma(\mu, \nu)} \iint c(x, y) d\rho(x, y). \quad (1.33)$$

In the above, $\Gamma(\mu, \nu)$ is the set of measures whose marginals are μ and ν , and the function $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the transport cost of moving mass from x to y . The Kantorovich OT problem is a constrained linear program; solving for the optimal plan ρ exactly is difficult but critical for applications in machine learning, in particular, generative modeling and sampling. We reformulate this as an unconstrained min-max problem by rewriting the marginal constraints with (nonnegative) KL divergences, the weight of which is maximized,

$$\begin{aligned} & \sup_{\Lambda \in \mathbb{R}} \inf_{\rho} E(\rho, \Lambda; \mu, \nu), \\ E(\rho, \Lambda; \mu, \nu) &= \iint c(x, y) d\rho(x, y) + \Lambda (\text{KL}(\rho_1 | \mu) + \text{KL}(\rho_2 | \nu)), \end{aligned} \quad (1.34)$$

where ρ_1 and ρ_2 are the marginals of the joint probability measure ρ . The Nash equilibrium of the min-max problem coincides with the minimizer of the OT problem, and this formulation allows us to develop a gradient flow particle method to dynamically approximate the optimal coupling. In order to solve (1.33) numerically, we will solve (1.34) using a min-max gradient flow with respect to the metric $\mathcal{W}_2$

$$\begin{aligned} \partial_t \rho(t) &= -\frac{1}{\beta} \nabla_{\mathcal{W}_2, \rho} E(\rho(t), \Lambda(t); \mu, \nu), \\ \frac{d}{dt} \Lambda(t) &= \partial_{\Lambda} E(\rho(t), \Lambda(t); \mu, \nu), \end{aligned} \quad (1.35)$$

where $\nabla_{\mathcal{W}_2, \rho}$ denotes the gradient with respect to ρ in the $\mathcal{W}_2$ metric, and $0 < \beta < 1$ is a temperature parameter specifying how fast ρ updates relative to Λ . When β is close to zero, (1.35) is a two-timescale dynamical system. We will prove convergence in the best-response setting, that is, when $\beta \rightarrow 0$. In this setting, ρ is instantly minimizing $E(\cdot, \Lambda; \mu, \nu)$, and the dynamics are given by

$$\begin{aligned} \frac{d}{dt} \Lambda(t) &= \partial_{\Lambda} E(b[\Lambda], \Lambda; \mu, \nu), \\ b[\Lambda] &:= \operatorname{argmin}_{\rho \in \mathcal{P}_2} E(\rho, \Lambda; \mu, \nu). \end{aligned} \quad (1.36)$$

Let $E_d(\Lambda) := E(b[\Lambda], \Lambda)$; then we can write (1.36) as

$$\frac{d}{dt} \Lambda(t) = \partial_{\Lambda} E_d(\Lambda).$$

For convex and smooth c , the minimizer of $E(\cdot, \Lambda; \mu, \nu)$ was proven to exist for any fixed $\Lambda > 0$ in [45, Corollary 1], and therefore $b[\Lambda]$ is well-defined for any $0 < \Lambda < +\infty$. The other timescale-separation regime, when Λ instantly best-responds, is no longer an entropy-regularized problem because as $\Lambda \rightarrow +\infty$, the optimization functional E Γ -converges to the optimal transport problem (1.33) [45, Theorem 7]. We prove that $b[\Lambda]$ converges to the true minimizer of (1.33), ρ^* , under the assumption that $E(\rho, \Lambda; \mu, \nu)$ has a unique minimizer for all admissible μ, ν , and Λ ; this justifies the choice of gradient descent-ascent as a solving method for the min-max optimal transport problem.

Corollary 1.13 (Convergence of OT Cost and Minimizer). *The best response $b[\Lambda(t)] \rightarrow \rho^*$ in $\mathcal{W}_2$, where ρ^* is an optimizer for (1.33) with cost $c^* = \iint c(x, y) d\rho^*(x, y)$, and $E_d(\Lambda(t)) \rightarrow c^*$ as $t \rightarrow \infty$.*

In comparison to other computational OT methods, our method is not expected to compete with the speed of the Sinkhorn algorithm [21], which is arguably the fastest available method for entropy-regularized OT. Entropy regularization refers to the OT problem with the term $\frac{1}{\Lambda} \text{KL}(\rho \mid \mu \otimes \nu)$ added to the cost function. In practice, the entropy regularized OT problem is solved for a fixed Λ , and then Λ is increased and the optimization problem is solved again. This process repeats until the entropy term is sufficiently annealed. A key question was open until around the time of the publication of our work: how fast should $\Lambda \rightarrow \infty$? Our work gives an upper bound for how fast $\Lambda \rightarrow \infty$ for our setting, and the order matches that of [15] where the question was answered for the discrete-space Sinkhorn algorithm.

Theorem 1.14 (Regularization growth condition). *Under the dynamics (1.36), the regularization parameter Λ satisfies*

$$\Lambda(t) \leq \sqrt{2(c^*t + c_0)} \quad \forall t \geq 0,$$

where $c^* < \infty$ is the optimal value of (1.33) and $c_0 > 0$ is a constant dependent on $\Lambda(0) > 0$.

From this result, we see that $\Lambda(t)$ should increase at a rate at most $O(\sqrt{t})$. Next, we propose a numerical implementation for solving the min-max OT problem. Denoting

$$V(\Lambda) := \text{KL}(b_1[\Lambda] \mid \mu) + \text{KL}(b_2[\Lambda] \mid \nu),$$

we can write

$$E_d(\Lambda) = \iint c(x, y) db[\Lambda](x, y) + \Lambda V(\Lambda),$$

Our proof of Danskin's Theorem allows us to use that (1.36) is equivalent to $\dot{\Lambda} = V(\Lambda)$, which circumvents taking the gradient of $b[\Lambda]$ with respect to Λ . In our numerical implementation⁶, we chose very small β , which results in ρ updating much faster than Λ , and we expect this to approximate the dynamics (1.36) due to the timescale separation results in [5]. We adapt a method from [45] by approximating the distribution ρ by particles $(X_i, Y_i) \sim \rho$. We divide the points $(X_j, Y_j)_{j=1}^{2N}$ into two sets: for one set, denoted $(X_i^{(1)}, Y_i^{(1)})_{i=1}^N$, we initialize $X_i^{(1)} \sim \mu$ and $Y_i^{(1)} = X_i^{(1)} + \xi_i$, where $\xi_i \sim \mathcal{N}(0, \eta)$ with small η , and we evolve $Y_i^{(1)}$ while keeping $X_i^{(1)}$ fixed. For the other set, denoted

⁶Code available at <https://github.com/LEConger/OT-MinMax-GradientFlow>.

$(X_i^{(2)}, Y_i^{(2)})_{i=1}^N$ we initialize $Y_i^{(2)} \sim \nu$ and set $X_i^{(2)} = Y_i^{(2)} + \xi_i$, and evolve $X_i^{(2)}$ while keeping $Y_i^{(2)}$ fixed, resulting in

$$\begin{aligned}\frac{d}{dt}X_i^{(2)}(t) &= -\frac{1}{\beta}\nabla_x c(X_i^{(2)}(t), Y_i^{(2)}(t)) + \frac{1}{\beta}\Lambda(t)\nabla_x \log \frac{\rho_1(t, X_i^{(2)}(t))}{\mu(X_i^{(2)}(t))} \\ \frac{d}{dt}Y_i^{(1)}(t) &= -\frac{1}{\beta}\nabla_y c(X_i^{(1)}(t), Y_i^{(1)}(t)) + \frac{1}{\beta}\Lambda(t)\nabla_y \log \frac{\rho_2(t, Y_i^{(1)}(t))}{\nu(Y_i^{(1)}(t))} \\ \frac{d}{dt}\Lambda(t) &= \text{KL}(\rho_1(t) \mid \mu) + \text{KL}(\rho_2(t) \mid \nu) .\end{aligned}$$

The densities ρ_1 and ρ_2 are approximated by binning the points and using finite differences:

$$\begin{aligned}\rho_1(x) &\approx \frac{1}{2N} \sum_{k=1}^{N_b} \mathbf{1}_{\Omega_k}\{x\} \sum_{m=1}^N (\mathbf{1}_{\Omega_k}\{X_m^{(1)}\} + \mathbf{1}_{\Omega_k}\{X_m^{(2)}\}) \\ \rho_2(y) &\approx \frac{1}{2N} \sum_{k=1}^{N_b} \mathbf{1}_{\Omega_k}\{y\} \sum_{m=1}^N (\mathbf{1}_{\Omega_k}\{Y_m^{(1)}\} + \mathbf{1}_{\Omega_k}\{Y_m^{(2)}\}) ,\end{aligned}$$

where the domain is divided into N_b regions, and the fraction of particles in each region Ω_k is counted. Intuitively, this method flows half of the X_i points from ν to μ and half of the Y_i points from μ to ν , while ensuring that the other half of each set is fixed at the desired target distribution. We illustrate the convergence of this algorithm empirically on a two-dimensional example.

Our min-max formulation of the OT problem enables the analysis for the timescale separated dynamics, revealing the insight that Λ should update no faster than $\mathcal{O}(\sqrt{t})$; this was not known until recently for Sinkhorn and other regularized OT algorithms. While we do not expect our numerical method to be as fast as Sinkhorn, our particle method allows for solutions to the OT problem to have different support than μ and ν , which is not possible with Sinkhorn and is relevant for high-dimensional settings where empirical measures have very sparse support.

1.2 Part II: Distributed Control

Instead of viewing a large-scale system from a high-level, distribution perspective, we now shift to a local perspective. For control systems with massive states, we will develop tools for developing local, distributed controllers that bypass the need for global information and communication. First, we introduce fundamentals of control theory, and then the System Level Synthesis parameterization which will enable our methods.

The aim of control theory is to drive a *state*, $x \in \mathbb{R}^{n_x}$, using some input $u \in \mathbb{R}^{n_u}$, in a way that minimizes cost in the presence of disturbances $w \in \mathbb{R}^{n_x}$. In the setting of linear dynamical system, the dynamics evolve according to

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad \forall t \geq 0,$$

for $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$. There can be various objectives for the control; one could aim to drive the state to a particular destination $\bar{x}$ in T time steps, or minimize a cost that is quadratic in the state x_t and input u_t , or drive x_t as close as possible to a reference signal r_t . Control applications appear everywhere: vehicles, heating and cooling, flight, navigation, power systems, robotics, and economics. In all of these settings, there is some input to the system, such as turbine speed, thrust, or motor actuators that produce some effect in the larger system. Control theory formalizes the dynamics of these systems and provides methods for optimizing the inputs, incorporating features like robustness to noise and model inaccuracy.

In the domain of modern cyber-physical systems, such as decentralized power grids and autonomous swarms, traditional control synthesis encounters significant theoretical and computational limitations. Historically, control design has been formulated by optimizing the controller explicitly. However, imposing pragmatic structural constraints—such as communication delays, spatial locality, and actuation limits—onto this standard parameterization generally results in highly non-convex and computationally intractable optimization problems. Consequently, guaranteeing robust performance in complex, distributed networks remains a challenge.

Furthermore, these inherent limitations impede advancements in contemporary data-driven control and complex physical modeling. When employing non-asymptotic machine learning techniques to identify system dynamics, engineers must account for large or complex uncertainty sets. Conventional control parameterizations lack the structural capacity to integrate these learned uncertainty bounds without forfeiting robust stability guarantees. Similarly, extending these non-convex methodologies to infinite-dimensional physical systems, such as those governed by partial differential equations, is frequently computationally prohibitive.

As a concrete example, consider practitioners dealing with information and actuation delays. For safety purposes, one might want to stop disturbances from propagating outside of a particular region.

In very large scale systems, we would like to solve for the controllers using only local information about the dynamics, but with some sort of global optimality. A classical approach for dealing with information delays is to construct augmented states, where the new state consists of the last few time steps:

$$\hat{x}_t := \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \end{bmatrix}, \quad \hat{x}_{t+1} = \begin{bmatrix} A & 0 & 0 \\ I & 0 & 0 \\ 0 & I & 0 \end{bmatrix} \hat{x}_t + \begin{bmatrix} B \\ 0 \\ 0 \end{bmatrix} u_t + w_t.$$

Then we can place support constraints on the controller $u_t = K\hat{x}_t$ to ensure that information about x is used only when the information is actually available. However, this does not solve the issue of using only local information, or the issue of stopping disturbances from propagating too far.

System Level Synthesis (SLS), proposed in [70], provides a robust theoretical framework to overcome these challenges through a reparameterization. Rather than optimizing the controller directly, SLS shifts the optimization space to the closed-loop system responses. This reformulation effectively transforms historically intractable, structurally-constrained control problems into tractable, convex optimization programs. By leveraging this convex architecture, engineers can systematically enforce sparsity, communication delays, and locality constraints directly upon the system's behavior while maintaining stability guarantees.

Beyond resolving topological constraints, SLS provides a bridge between data-driven learning and robust control synthesis. Its convex framework accommodates the complex uncertainty sets generated by non-asymptotic learning methods, such as set-membership identification, thereby ensuring stable and reliable control even under imperfectly-modeled system dynamics. Additionally, our recent theoretical advancements have demonstrated that this exact convex synthesis framework scales effectively to infinite-dimensional evolution equations, providing a unified and scalable solution to the structural complexities that constrain traditional control methodologies.

This introductory section first provides the SLS parameterization and basic results, such as equivalence to linear feedback control and the linear quadratic regulator in the SLS parameterization. Then we discuss how the Alternating Direction Method of Multipliers (ADMM) enables distributed, localized solving of constrained large-scale systems, all while maintaining global optimality guarantees. *Distributed, constrained solving with global optimality is the key feature of System Level Synthesis.* Then the contributions in this thesis are summarized: first, the construction of notions of observability and controllability for constrained systems, then a set-based method for learning unknown dynamics in the presence of possibly-adversarial disturbances, and finally establishing the SLS parameterization for evolution equations, including partial differential equations (PDEs).

1.2.1 System Level Synthesis (SLS) Parameterization

Throughout the introduction, for clarity of presentation we will present ideas in the setting where the state is fully observed; note that the SLS framework is also applicable in the partially observed setting. The key idea of System Level Synthesis (SLS) is to parameterize the controller u in terms of the disturbance signal w , instead of in terms of the state x . Let T be the time horizon over which we will solve for the control input, with dynamics given by

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad t \in [0, \dots, T-1]. \quad (1.37)$$

We will stack the state, input, and disturbances into vectors

$$x := \begin{bmatrix} x_0 \\ \vdots \\ x_T \end{bmatrix} \in \mathbb{R}^{(T+1)n_x}, \quad u := \begin{bmatrix} u_0 \\ \vdots \\ u_T \end{bmatrix} \in \mathbb{R}^{(T+1)n_u}, \quad w = \begin{bmatrix} x_0 \\ w_0 \\ \vdots \\ w_{T-1} \end{bmatrix} \in \mathbb{R}^{(T+1)n_x}.$$

Note that the first entry of w is x_0 ; this carries the interpretation of the initial condition as a disturbance. We also stack the matrices that drive the dynamics into block matrices, resulting in

$$\mathcal{A} = \begin{bmatrix} 0 & & \\ A & \ddots & \\ & A & 0 \end{bmatrix} \in \mathbb{R}^{(T+1)n_x \times (T+1)n_x}, \quad \mathcal{B} = \begin{bmatrix} 0 & & \\ B & \ddots & \\ & B & 0 \end{bmatrix} \in \mathbb{R}^{(T+1)n_x \times (T+1)n_u}.$$

Now the dynamics (1.37) can be equivalently written as

$$x = \mathcal{A}x + \mathcal{B}u + w. \quad (1.38)$$

Before the SLS parameterization, linear control theory primarily surrounded controllers mapping from x to u . For linear dynamics and quadratic cost functions, optimizing over $(K_t)_{t=0}^{T-1} \in \mathbb{R}^{n_x \times n_u}$ is optimal [39], where $u_t = K_t x_t$. Using (1.38), we instead optimize over maps θ_u from w to u and θ_x from w to x . Now we are equipped to define the closed-loop maps (CLM) $\theta_x \in \mathbb{R}^{(T+1)n_x \times (T+1)n_x}$ and $\theta_u \in \mathbb{R}^{(T+1)n_u \times (T+1)n_x}$,

$$\theta_x = \begin{bmatrix} \theta_x(0,0) & & \\ \vdots & \ddots & \\ \theta_x(T,T) & \dots & \theta_x(T,0) \end{bmatrix}, \quad \theta_u = \begin{bmatrix} \theta_u(0,0) & & \\ \vdots & \ddots & \\ \theta_u(T,T) & \dots & \theta_u(T,0) \end{bmatrix}.$$

The matrix $\theta_x(t, \tau)$ is the effect of the disturbance from τ time steps ago on the state at time t , and $\theta_u(t, \tau)$ is the effect of the disturbance from τ time steps ago on the control input at time t . Note that θ_x and θ_u are block-lower triangular so that the system is *causal*, that is, the state at time t should

depend only on the $\tau > 0$ previous time steps. The state and input can be written in terms of the closed-loop maps,

$$\begin{bmatrix} x \\ u \end{bmatrix} = \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} w. \quad (1.39)$$

Now rewriting the dynamics (1.38) in terms of the closed-loop maps, we have

$$\theta_x w = \mathcal{A}\theta_x w + \mathcal{B}\theta_u w + w,$$

and in order for (1.39) to hold for all possible sequences of disturbances and initial conditions w , we obtain the system level parameterization (SLP)

$$\begin{bmatrix} I - \mathcal{A} & -\mathcal{B} \end{bmatrix} \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} = I. \quad (\text{SLP})$$

The SLP (SLP) is the space over which we optimize for controllers. Instead of optimizing over K where $u = Kx$, we now optimize over θ_x, θ_u where $x = \theta_x w$ and $u = \theta_u w$. Next, we show that the two parameterizations are equivalent.

Equivalence with linear feedback controllers. In classical control theory, for linear dynamics and quadratic cost functions, the optimal control input u is parameterized via

$$u = Kx, \quad K \text{ is lower-block triangular.}$$

Note that here since $u \in \mathbb{R}^{(T+1)n_u}$ is the control for the whole time horizon, the feedback controller K is time-dependent. A natural question is whether the space of controllers parameterized by (θ_x, θ_u) is the same as the space of controllers parameterized by K . We will prove that all trajectories (x, u) generated from $u = Kx$ can be achieved by some (θ_x, θ_u) satisfying the system level parameterization, and all (x, u) generated from (θ_x, θ_u) satisfying the system level parameterization can be realized by $K = \theta_u \theta_x^{-1}$. The takeaway is that parameterizing the controller by K or by (θ_x, θ_u) is equivalent. This is established in the following theorem.

Theorem 1.15 (Theorem 2.1 [4]). *All trajectories (x, u) generated from $u = Kx$, with K causal, can be achieved by some (θ_x, θ_u) satisfying (SLP), and all (x, u) generated from (θ_x, θ_u) satisfying (SLP) can be realized by $K = \theta_u \theta_x^{-1}$.*

Infinite Horizon. Theorem 1.15 can also be shown in the infinite horizon setting using z -transforms and signals. In the infinite horizon setting, we need to verify that (SLP) is feasible. Let us assume that A, B are stabilizable⁷; [4, Lemma 4.2] states that the infinite horizon SLP is feasible if and only

⁷Dynamics parameterized by (A, B) are *stabilizable* if all unstable eigenvectors of A are in the controllable subspace of (A, B) . This means the system can be controlled to the state $0 \in \mathbb{R}^{n_x}$ in finite time.

if A, B is stabilizable, giving us the desired property. Let $x : \mathbb{N}_{\geq 0} \rightarrow \mathbb{R}^{n_x}$, $u : \mathbb{N}_{\geq 0} \rightarrow \mathbb{R}^{n_u}$ be signals and let scripts denote transfer functions:

$$\mathcal{G}(z) = \sum_{i=0}^{\infty} \frac{1}{z^i} G_i \quad \forall |z| > 1.$$

The notation $\mathcal{G} \in \frac{1}{z}\mathcal{RH}_{\infty}$ means $\mathcal{G}$ is strictly proper; in particular, $\mathcal{R}$ denotes real, rational transfer functions and $\mathcal{H}_{\infty}$ denotes proper functions (see [4] for more details and definitions). Then $z\mathcal{G} \in \mathcal{RH}_{\infty}$ indicates that a one-step-time-delayed $\mathcal{G}$ is a rational proper transfer function, which means $\mathcal{G}$ is rational and strictly proper. The dynamics in signal space are

$$zx = Ax + Bu + w,$$

and rearranging with $x = \theta_x w$ and $u = \theta_u w$ results in the SLP

$$\begin{bmatrix} zI - A & -B \end{bmatrix} \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} = I, \quad \theta_x, \theta_u \in \frac{1}{z}\mathcal{RH}_{\infty}.$$

Now that we have the SLP, we can show the same steps as in the proof of Theorem 1.15.

Continuous Space and Time. Theorem 1.15 also holds for continuous-time and continuous-space dynamics; the continuous time case is studied in [36, 27] and we prove results for continuous space in [19]. The dynamics have variables in Hilbert spaces, with state variable $x(t)$, control input $u(t)$, and disturbance $w(t)$; taking the Laplace transform results in

$$\dot{x}(t) = Ax(t) + Bu(t) + w(t), \quad \Rightarrow \quad sX(s) = AX(s) + BU(s) + W(s),$$

where $X(s)$ is the state at frequency s . The operators A and B are now linear operators between Hilbert spaces. The SLP is

$$\begin{bmatrix} sI - A & -B \end{bmatrix} \begin{bmatrix} \theta_x(s) \\ \theta_u(s) \end{bmatrix} = I, \quad \text{for all } \theta_x : \mathbb{R} \rightarrow \mathbb{R}^{n_x \times n_x}, \theta_u : \mathbb{R} \rightarrow \mathbb{R}^{n_u \times n_x}.$$

This setting is useful for PDE and integral equation dynamics, where the system is inherently in continuous space or time. We study the discrete time, continuous-space setting in [19], and show that in some cases optimizing in continuous space and then discretizing for implementation can improve controller performance, and we will discuss this in Chapter 9. We now return to discrete time for the linear quadratic regulator.

Linear Quadratic Regulator. Quadratic cost functions are very commonly used in control theory; we illustrate the SLS parameterization on this familiar setting. The linear quadratic regulator

(LQR) problem derives its name from *linear* dynamics and a *quadratic* cost function. The LQR problem is

$$\min_{(u_t)_{t=0}^{T-1}} \sum_{t=0}^{T-1} x_t^\top Q_t x_t + u_t^\top R_t u_t, \quad Q \succeq 0, R \succ 0,$$

$$\text{such that } x_{t+1} = Ax_t + Bu_t, \quad x_0 = \tilde{x}.$$

The choice of Q and R determine the relative cost of control versus the cost of a nonzero state. Writing the problem in terms of the vectorized state and input,

$$\min_u x^\top \hat{Q}x + u^\top \hat{R}u$$

$$\text{such that } x = \mathcal{A}x + \mathcal{B}u + w,$$

where $\mathcal{A}$ and $\mathcal{B}$ are defined earlier and $\hat{Q}$ and $\hat{R}$ are the diagonal block matrices containing the cost matrices

$$w = \begin{bmatrix} x_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \hat{Q} := \begin{bmatrix} Q_0 & & \\ & \ddots & \\ & & Q_{T-1} \end{bmatrix} \in \mathbb{R}^{Td_x \times Td_x}, \quad \hat{R} := \begin{bmatrix} R_0 & & \\ & \ddots & \\ & & R_{T-1} \end{bmatrix} \in \mathbb{R}^{Td_u \times Td_u}.$$

Since there are no disturbances in the dynamics for this problem, we set the entries of w to zero, except for the possibly-nonzero initial condition x_0 . Rewriting the LQR problem in terms of closed loop maps, we have

$$\min_{\theta_x, \theta_u} w^\top \theta_x^\top \hat{Q} \theta_x w + w^\top \theta_u^\top \hat{R} \theta_u w$$

$$\text{such that } \begin{bmatrix} I - \mathcal{A} & -\mathcal{B} \end{bmatrix} \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} w = w.$$

If the initial condition is known or is drawn from a distribution, using the definition of w in this setting, we can reduce the problem to an expectation over the distribution of initial conditions for x_0 . Let $\theta_x(t) := \theta_x(t, t)$ and $\theta_u(t) := \theta_u(t, t)$. The optimization problem results in

$$\min_{\theta_x, \theta_u} \mathbb{E}_{x_0} \left(\left\| \begin{bmatrix} \hat{Q}^{1/2} & \\ & \hat{R}^{1/2} \end{bmatrix} \begin{bmatrix} (\theta_x(t))_t \\ (\theta_u(t))_t \end{bmatrix} x_0 \right\|_2^2 \right),$$

$$\text{such that } \begin{bmatrix} I - \mathcal{A} & -\mathcal{B} \end{bmatrix} \begin{bmatrix} (\theta_x(t))_t \\ (\theta_u(t))_t \end{bmatrix} = I[:, 0],$$

where $I[:, 0]$ denotes the first block column of the identity matrix. The notation $\|\cdot\|_{i \rightarrow j}$ denotes the j norm induced by the i norm, that is $\|A\|_{i \rightarrow j} = \sup_x \frac{\|Ax\|_j}{\|x\|_i}$. In more general settings, one can select a variety of norms, including the Frobenius norm and induced norms, such as

$$\min_{\theta_x, \theta_u} \left\| \begin{bmatrix} \hat{Q}^{1/2} & \\ & \hat{R}^{1/2} \end{bmatrix} \begin{bmatrix} (\theta_x(t))_t \\ (\theta_u(t))_t \end{bmatrix} \right\|_{2 \rightarrow 2}^2$$

such that $\begin{bmatrix} I - \mathcal{A} & -\mathcal{B} \end{bmatrix} \begin{bmatrix} (\theta_x(t))_t \\ (\theta_u(t))_t \end{bmatrix} = I[:, 0]$.

Using the system level parameterization, θ_x can be written as a function of θ_u , because $I - \mathcal{A}$ is always invertible:

$$\theta_x = (I - \mathcal{A})^{-1}(\mathcal{B}\theta_u + I).$$

However, θ_u is not necessarily determined by θ_x because $\mathcal{B}$ is not necessarily invertible. Thus, the optimization problems above can be reduced to unconstrained optimization problems over θ_u .

Incorporating System-Level Constraints. Before discussing distributed optimization techniques, it is crucial to understand how system-level constraints are mathematically formulated and what they mean for the underlying system. Traditional notions of controllability and observability assume unconstrained access to all states and inputs, which is unrealistic for large-scale, distributed networks like power grids or traffic systems. As detailed in [20], the SLS framework allows us to translate spatiotemporal hardware limitations directly into convex structural constraints on the closed-loop system responses, θ_x and θ_u .

In the noise-free SLS parameterization, $\theta_x(t)$ and $\theta_u(t)$ map the initial state x_0 to the state $x(t)$ and control input $u(t)$ at time t , respectively. We can explicitly encode physical network limitations by restricting the support of these block matrices θ_x, θ_u over time. By imposing element-wise zero-constraints, such as $\theta_x(t)[j, i] = 0$ or $\theta_u(t)[j, i] = 0$, we dictate that information from the i -th subsystem's initial state cannot affect the j -th subsystem's state or control action at time t . Collecting all such spatiotemporal requirements defines a structural constraint subspace $\mathcal{S}$. Enforcing the condition that the stacked system responses belong to $\mathcal{S}$ restricts the system to physically realizable controllers while preserving the convexity of the synthesis problem.

We show in Chapter 7 that imposing these restrictions can, depending on the constraint type, fundamentally alter the reachable state space and the system's ability to be controlled or observed, with different constraints yielding distinct physical consequences:

- **Actuation Delays:** These dictate that controllers cannot output control actions until a specified time has elapsed. The inclusion of these delays causes a strict, gradual loss of controllability

and observability (in the case of sensor delay), physically shrinking the volume of the reachable state space.

- **Communication Constraints:** These govern the allowed routing of information between different subsystems over the network. They limit which controllers have access to different state measurements.
- **Locality Constraints:** These restrict the spatial footprint of disturbances, ensuring that a subsystem prevents disturbances from propagating outside of a strictly defined local neighborhood.

A key theoretical result from [20] demonstrates that, unlike delays, the addition of locality and communication constraints does not necessarily decrease the controllability and observability volume. Under certain sufficient conditions, a strictly localized and communication-constrained controller can achieve the same reachable state space as a fully centralized one. Support constraints on θ_x, θ_u provide a tractable computational method for localized, distributed control architectures.

Next, we discuss the optimization method that enables distributed, parallel solving.

Solving System Level Synthesis with ADMM. The Alternating Direction Method of Multipliers (ADMM) is an iterative method for solving constrained optimization problems of the form

$$\begin{aligned} \min_{v \in \mathbb{R}^d, z \in \mathbb{R}^n} \quad & f(v) + g(z) \\ \text{such that} \quad & Fv + Gz = c . \end{aligned}$$

Each iteration k consists of solving an *unconstrained* optimization over v , then an optimization over z , and then a constraint update:

$$\begin{aligned} v^{k+1} &= \arg \min_{v \in \mathbb{R}^d} f(v) + \frac{\rho}{2} \|Fv + Gz^k - c + \mu^k\|_2^2 \\ z^{k+1} &= \arg \min_{z \in \mathbb{R}^n} g(z) + \frac{\rho}{2} \|Fv^{k+1} + Gz - c + \mu^k\|_2^2 \\ \mu^{k+1} &= \mu^k + Fv^{k+1} + Gz^{k+1} - c . \end{aligned} \tag{1.40}$$

The coefficient ρ is an update temperature. The variables v and z are the primal variables while μ is the dual variable, and can be interpreted as the running sum of residual constraint error:

$$\mu^k = \mu^0 + \sum_{j=1}^k r^j, \quad \text{where } r^k := Fv^k + Gz^k - c \text{ is the residual at step } j .$$

In order for the scheme to provably converge to the global optimum, the following assumption is required [6, Assumption 1].

Assumption 1.16. *The functions $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ and $g : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ are closed, convex, and proper.*

An equivalent criterium is that the epigraphs of f and g are closed, nonempty, and convex. The *epigraph* of a function is the set of pairs of points and function values for which they are in the corresponding level set,

$$\text{epi}(f) = \{(v, t) \in \mathbb{R}^d \times \mathbb{R} \mid f(v) \leq t\}.$$

Lower semicontinuity and convexity implies Assumption 1.16. Under Assumption 1.16, there exists some v^* , z^* , and μ^* which is a stationary point of (1.40) [6]. Additionally, the residual, objective, and dual variables satisfy

$$r^k \rightarrow 0, \quad f(v^k) + g(z^k) \rightarrow p^*, \quad \mu^k \rightarrow \mu^*,$$

where p^* is the global optimal value of the optimization problem.

In the context of the SLS framework, the synthesis task shifts from directly designing a controller to optimizing over the closed-loop system responses. The standard constrained SLS problem for a linear discrete-time system takes the form

$$\begin{aligned} \min_{\theta_x, \theta_u} \quad & J(\theta_x, \theta_u) \\ \text{such that} \quad & [zI - \mathcal{A} \quad -\mathcal{B}] \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} = I, \quad \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} \in \mathcal{S}, \end{aligned} \tag{1.41}$$

where θ_x and θ_u are the state and control system responses, and $\mathcal{S}$ represents structural constraints (e.g., sparsity, delay, or spatial locality patterns) necessary for distributed implementation. Note that the affine equality constraint ensures that the system responses are achievable. J is a closed, convex performance objective (such as the $\mathcal{H}_2$ norm), and to apply ADMM we require $J = J^c + J^r$ to be the sum of column- and row- separable cost functions, J^c and J^r respectively. A cost function J^c (J^r) is column (row) separable if the columns (rows) of θ_x and θ_u can be optimized over independently to obtain the optimal value.

To apply ADMM to (1.41), we introduce variable splitting to decouple the global dynamical constraints from the localized structural constraints. We map the SLS problem to the generic ADMM formulation by defining our primal variable v to handle the system dynamics, and z to act as a consensus copy handling the structure. Let $v := (\theta_x, \theta_u)$ and $z := (\Psi_x, \Psi_u)$. We can define our objective components as:

$$\begin{aligned} f(v) &:= J^c(v) + \mathcal{I}_{\text{affine}}(v), \\ g(z) &:= \mathcal{I}_{\mathcal{S}}(z) + J^r(z), \end{aligned}$$

where $\mathcal{I}_{\text{affine}}(v)$ is the indicator function for the affine achievable response space $\{v \mid [zI - \mathcal{A} - \mathcal{B}]v = I\}$, and $\mathcal{I}_{\mathcal{S}}(z)$ is the indicator function for the support constraints $\mathcal{S}$. To enforce consensus, we set $F = I$, $G = -I$, and $c = 0$, so that $v - z = 0$. Because the affine subspace and the standard SLS structural constraints are closed and convex, f and g naturally satisfy Assumption 1.16.

Substituting these functions into the general ADMM updates (1.40) yields a highly scalable, distributed algorithm for optimal control synthesis:

- **The v -update (column-separable cost and dynamics enforcement):** The support constraints are temporarily relaxed. We minimize the performance cost $J^c(v)$ plus a quadratic proximity penalty $\frac{\rho}{2} \|v - z^k + \mu^k\|_2^2$, subject only to the affine system dynamics.
- **The z -update (row separable cost and support constraint enforcement):** The dynamics constraint is not used and the row separable cost function $J^r(z)$ is minimized subject to the support constraints. Since $\mathcal{I}_{\mathcal{S}}$ is both row and column separable, we can incorporate it into either the v or z update.
- **The μ -update (dual update):** We update the dual variable $\mu^{k+1} = \mu^k + v^{k+1} - z^{k+1}$ to penalize the residual constraint error, driving v and z toward consensus.

Because both f and g meet Assumption 1.16, the residual $r^k = v^k - z^k \rightarrow 0$ as $k \rightarrow \infty$ and the system responses iteratively converge to the globally optimal parameterization. In the implementation, the algorithm terminates when sufficient accuracy has been achieved, for example, numerical precision. As emphasized in [4], this optimization drastically simplifies into parallel, subproblems that can be solved completely independently, which we illustrate next.

Column- and Row-wise Separability and Distributed Synthesis. The true scalability of the System Level Synthesis framework, particularly when paired with ADMM, stems from the column- and row-wise separability of the synthesis problem. As established by [70, 4], for a wide class of optimal control problems, the massive centralized optimization decomposes into a set of independent, parallelizable subproblems. This decoupling occurs because the core components of the SLS formulation interact with the columns of the system responses independently. Let $\theta = \begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix}$, let $\theta_{:j}$ denote the j^{th} column of this stacked response matrix, and let $\theta_{i:}$ denote the i^{th} row. We can analyze the separability across three fronts:

1. **Affine Dynamics Constraint:** The global affine constraint $[zI - A - B]\theta = I$ naturally decomposes by column. For the j^{th} column, the constraint is simply $[zI - A - B]\theta_{:j} = e_j$,

where e_j is the j -th standard basis vector. This means the dynamic feasibility of the j -th subsystem's response to a local disturbance depends only on its own column $\theta_{:j}$. In networks where states are influenced directly only by local neighbors, nonzero entries of $\theta_{:j}$ will be sparse and the problem can be solved locally.

2. **Objective Function:** Standard robust and optimal control objectives are row or column-wise separable, such as the $\mathcal{H}_2$ and $\mathcal{L}_1$ norms which are column-separable, and $\theta \mapsto \|\theta C\|_F^2$ (for a matrix C of compatible dimension) which is row-separable. The total cost can be expressed as a sum of separable costs: $J(\theta) = \sum_{j=1}^n J_j^c(\theta_{:j}) + \sum_{i=1}^n J_i^r(\theta_{i:})$.
3. **Structural Constraints:** The locality, communication, and delay constraints defining the set $\mathcal{S}$ are enforced via element-wise sparsity patterns, which is both column and row separable.

Enabling Distributed ADMM. Because $f(v)$ (dynamics and cost) and $g(z)$ (structure) decouple column-by-column, the ADMM updates (1.40) completely decouple as well. By partitioning our ADMM variables into columns— $v = [v_{:1}, \dots, v_{:n}]$, $z = [z_{1:}, \dots, z_{n:}]$, and $\mu = (\mu_{ij})_{ij}$ —the iterative algorithm breaks down into independent optimizations:

$$\begin{aligned} & \text{for all } j \in [1, \dots, 2(T+1)n_x], i \in [1, \dots, (T+1)(n_x + n_u)] : \\ & v_{:j}^{k+1} = \arg \min_{v_{:j}} \left(J_j^c(v_{:j}) + \mathcal{I}_{\text{affine}}(v_{:j}) + \frac{\rho}{2} \|v_{:j} - z_{:j}^k + \mu_{:j}^k\|_2^2 \right), \\ & z_{i:}^{k+1} = \arg \min_{z_{i:}} \left(J_i^r(z_{i:}) + \mathcal{I}_{\mathcal{S}_{i:}}(z_{i:}) + \frac{\rho}{2} \|v_{i:}^{k+1} - z_{i:} + \mu_{i:}^k\|_2^2 \right), \\ & \mu_{ij}^{k+1} = \mu_{ij}^k + v_{ij}^{k+1} - z_{ij}^{k+1}. \end{aligned}$$

For systems where the states are connected to local neighbors, this mathematical structure shifts the computation from global to distributed. Instead of a central entity solving a massive control problem, each subsystem j acts as its own independent computational node. In the v -update, subsystem j computes its optimal dynamically feasible response to a local disturbance. In the z -update, each subsystem i executes a similar computation. The dual variable μ_{ij} manages the consensus between the two.

In this way, ADMM leveraged atop the SLS framework provides a fully distributed synthesis algorithm. The computational complexity scales with the size of the localized neighborhoods rather than the size of the global network, allowing optimal controllers to be synthesized for arbitrarily large-scale systems.

1.2.2 Results: Controllability and Observability

While the System Level Synthesis (SLS) parameterization elegantly handles the inclusion of structural constraints such as delays, communication limits, and spatial locality via the subspace $\mathcal{S}$,

it is unknown how these restrictions impact the fundamental properties of the system. Classical control theory relies heavily on the concepts of controllability and observability—determining whether a system can be driven to an arbitrary state in finite time, and whether the initial state can be perfectly reconstructed from output measurements. However, these classical definitions implicitly assume centralized access to all states and inputs. When real-world distributed structural constraints are enforced, we need to understand whether control and state estimation performance are lost.

In Chapter 7, we establish a novel framework for characterizing the controllability and observability of systems subject to these practical spatiotemporal constraints. In addition to asking a binary question—"is the system controllable?"—we ask a volumetric one: "how much of the state space remains reachable when the controller is structurally restricted?"

Related Literature: Controllability and Observability under Constraints. The formal study of system controllability and observability originated with the seminal work of Kalman [38]. This foundational theory established algebraic rank conditions and Gramian matrices to determine whether a linear system could be driven to an arbitrary state or perfectly reconstructed from its outputs. However, this classical formulation implicitly assumes a centralized architecture with unconstrained, instantaneous access to all states and inputs. As control applications have scaled, researchers have increasingly recognized the inadequacy of these unconstrained assumptions for large-scale, networked systems, prompting a shift toward understanding how localized information structures restrict fundamental system capabilities.

A major trajectory of recent research addresses the challenge of synthesizing controllers under strict spatial and communication limitations. The development of the SLS framework in [70] provided a computationally tractable method for enforcing sparsity and delay constraints directly on closed-loop system responses. Building on these distributed synthesis architectures, recent theoretical advancements [61, 44] have explored the performance trade-offs inherent in localized control. These works demonstrated that localized controllers—such as distributed model predictive control and decentralized linear quadratic regulators—can achieve near-optimal performance relative to their centralized counterparts, provided the underlying network graph meets certain sufficient structural conditions.

Parallel to the study of performance degradation, another critical strand of literature investigates observability from the perspective of security and functional privacy. Recent work [72] explored how structural constraints can be intentionally leveraged to block observability, thereby preserving the privacy of sensitive network states from adversarial inference. This thesis synthesizes these disparate theoretical threads—bridging the gap between the algebraic guarantees of classical controllability, the scalable synthesis of SLS, and the structural metrics of privacy-preserving networks—to provide

a unified, closed-form volumetric characterization of what is fundamentally reachable and observable when a system is subjected to strict locality, communication, and actuation constraints.

To formalize this mathematically, consider the classical unconstrained linear dynamical system $x_{t+1} = Ax_t + Bu_t$. In the remainder of this subsection, we discuss controllability only, noting that observability follows substituting $A \leftrightarrow A^\top$ and $B \leftrightarrow C^\top$, where C specifies the observation y , given by $y = Cx$. A system is said to be *controllable* in time T if the state can be driven from any initial condition x_0 to any terminal state x_T with finite energy. For an unconstrained system, the controllability Gramian is defined as:

$$W_c = \sum_{i=0}^{T-1} A^i B B^\top (A^\top)^i$$

The magnitude of each eigenvalue of W_c represents the distance in the direction of the corresponding eigenvector that one unit of control energy provides. The overall controllability volume, V_c , is given by the determinant, $V_c = \sqrt{\det W_c}$. This represents the state-space volume reachable with one unit of control input. Each eigenvector of W_c can be viewed as the axis of an ellipsoid, and the product of the eigenvalues of W_c is the volume of the corresponding ellipsoid.

However, this classical formulation assumes that any arbitrary sequence of control inputs can be selected. When we enforce structural constraints, we restrict the closed-loop system responses such that $\begin{bmatrix} \theta_x \\ \theta_u \end{bmatrix} \in \mathcal{S}$. This defines the $\mathcal{S}$ -constrained reachable set, $\mathcal{R}_T(\mathcal{S})$, and a corresponding constrained Gramian $W_c(\mathcal{S})$. It naturally follows that the constrained volume is bounded by the unconstrained volume, $V_c(\mathcal{S}) \leq V_c$.

Our analysis characterizes the physical and geometric consequences of different constraints on this volume. We show that incorporating communication delays, actuation delays, and locality constraints still results in an ellipsoidal reachable set, whose volume can be computed in closed form. The computation follows by rewriting the Gramian in terms of the SLS operators θ_x, θ_u , and imposing the constraint set $\mathcal{S}$. Secondly, we establish explicit rank conditions on the system level constraints under which there is no performance loss for any convex cost function compared to the unconstrained system [20, Theorem 9]. When this rank condition is satisfied, a strictly localized and communication-constrained controller achieves $V_c(\mathcal{S}) = V_c$. These results have the following implications for particular constraints:

- **Actuation and Sensing Delays:** We show that incorporating strict delays inevitably causes a monotonic reduction in the volume of $\mathcal{R}_T(\mathcal{S})$. Because the controller must wait to react to disturbances, the system drifts further along its open-loop trajectory, physically shrinking the

space of states that can be actively stabilized or reached within the horizon with one unit of control energy.

- **Locality and Communication Constraints:** We show that, in contrast with the effects of delays, spatial locality constraints do *not* necessarily degrade the reachable volume. We provide an example of linearized swing equations for which the addition of locality constraints, restricting information to the nearest six neighbors, satisfy the rank condition for maintaining performance.

By mapping these structural limitations directly to the rank conditions and singular values of the constrained Gramians, this work provides system designers with a quantitative metric for evaluating network topologies. It allows engineers to mathematically verify whether a proposed decentralized sensor and actuator placement is sufficient to achieve the required control performance, well before synthesizing the controller.

In this work, we assumed knowledge of the state and input matrices A and B . Next, we consider the setting where A and B are unknown and must be learned from data in the presence of possibly adversarial disturbances.

1.2.3 Results: Learning Dynamics from Data via Set Membership

To effectively apply robust control techniques, including the constrained System Level Synthesis methodologies discussed previously, system designers require accurate models of the underlying dynamics. When the true state and input matrices, A^* and B^* , are unknown, they must be learned from observational data. In robust control, it is not sufficient to simply find a "best fit" point estimate for the dynamics; one must instead identify an *uncertainty set* that is guaranteed to contain the true system parameters. The size and shape of this uncertainty set are crucial: if the set is too large, the resulting robust control actions will be overly conservative and degrade system performance; conversely, if the set fails to contain the true system, robust stability guarantees are lost. In Chapter 8, we establish the first non-asymptotic convergence rate bounds for set membership estimation (SME), which are faster than the rate bounds for least squares estimation.

Related Literature: System Identification and Uncertainty Sets. A substantial body of recent literature spanning machine learning and learning-based control relies on statistical confidence bounds, most notably derived from least squares estimation (LSE). Seminal works in this domain [1, 24, 23] have established rigorous non-asymptotic error bounds for identifying linear systems. These methodologies are deeply rooted in the statistical properties of the system disturbances, requiring assumptions such as the noise being independently and identically distributed (i.i.d.) or strictly sub-Gaussian. While theoretically elegant, reliance on statistical confidence intervals often produces

large, conservative uncertainty sets when applied to physical systems where the deterministic bounds on the disturbances are strictly known but the underlying statistical distributions are not.

SME operates under a deterministic framework, assuming only that the unobserved disturbances w_t are bounded and reside within some known support set $\mathbb{W}$. Given a sequence of data $\{x_t, u_t, x_{t+1}\}_{t=0}^{T-1}$, SME constructs the feasible parameter set comprising all matrices $\theta = [A, B]$ that are compatible with the observed data and the assumed disturbance bounds. While SME has demonstrated strong empirical performance and is widely utilized in robust adaptive control, its theoretical guarantees have historically lagged behind those of LSE. Specifically, the *non-asymptotic convergence rate* of SME for linear systems—characterizing precisely how fast the estimated uncertainty set shrinks as the data sequence length T increases—remained an open question for general disturbance sets. Prior theoretical guarantees were largely restricted to simplified settings where the disturbance support was an ℓ_∞ ball.

Non-Asymptotic Convergence of SME. In Chapter 8, we address this theoretical gap by establishing the first non-asymptotic convergence rate bounds for SME under general convex disturbance supports. Consider the linear dynamical system

$$x_{t+1} = \theta^* z_t + w_t,$$

where $z_t = [x_t^\top, u_t^\top]^\top \in \mathbb{R}^{n_x+n_u}$ is the augmented state-input vector and $\theta^* = [A^*, B^*]$ is the true parameter matrix. Assuming the disturbances w_t are bounded within a set $\mathbb{W} := \{w : \|w\|_\infty \leq w_{\max}\}$, the SME uncertainty set Θ_T at time T is the intersection of all parameter matrices consistent with the data:

$$\Theta_T = \left\{ \theta \mid x_{t+1} - \theta z_t \in \mathbb{W}, \quad \forall t \in [0, \dots, T-1] \right\}.$$

To overcome the statistical dependencies inherent in closed-loop dynamical data, our analysis fundamentally relies on recent advancements in non-asymptotic statistics, specifically leveraging the block-martingale small-ball condition established in [62]. This technique allows us to rigorously guarantee persistence of excitation—showing that the empirical covariance matrix of the augmented state-input trajectories grows linearly with time—even in the presence of highly correlated or worst-case data sequences.

Assumption 1.17 (Tight bound on w_t). *For any $\epsilon > 0$, there exists $q_w(\epsilon) > 0$, such that for any $1 \leq j \leq n$, we have*

$$\min\{\mathcal{P}(w_{t,j} \leq \epsilon - w_{\max}), \mathcal{P}(w_{t,j} \geq w_{\max} - \epsilon)\} \geq q_w(\epsilon),$$

where $w_{t,j}$ denotes the j -th entry of vector w_t . Without loss of generality, we can further assume $q_w(\epsilon)$ to be non-decreasing with ϵ and $q_w(2w_{\max}) = 1$.

Assumption 1.17 tells us that there is some positive probability that w_t visits an ϵ -neighborhood of $w_{\max}$ and $-w_{\max}$ for all coordinates. This enables us to state the following main theoretical result regarding the size of the uncertainty set.

Theorem 1.18 (Non-Asymptotic Convergence of SME). *Consider the system dynamics driven by bounded disturbances $w_t \in \mathbb{W}$ which satisfy Assumption 1.17 with $q_w(\epsilon) \geq \Omega(\epsilon)^8$. The diameter of the SME uncertainty set Θ_T in the Frobenius norm satisfies:*

$$\text{diam}(\Theta_T) \leq \tilde{O}\left(\frac{n_x^{1.5}(n_x + n_u)^2}{T}\right).$$

Distributions that satisfy the assumptions include the uniform distribution on $\mathbb{W}$, the zero-mean truncated Gaussian on $\mathbb{W}$, and the uniform distribution on the boundary of $\mathbb{W}$. This result demonstrates that SME enjoys a convergence rate of $\tilde{O}(1/T)$, which is faster than the standard LSE error bound of $\tilde{O}(1/\sqrt{T})$ in terms of the number of samples T , although it presents a worse dependence on the state and control dimensions n_x and n_u .

These non-asymptotic results lay a critical foundation for the theoretical analysis of downstream control designs. By explicitly characterizing how fast the SME uncertainty sets shrink, we enable novel non-asymptotic guarantees for robust frameworks, directly applying these estimation bounds to synthesize safe and optimal controllers via robust System Level Synthesis (SLS) and robust-adaptive model predictive control.

1.2.4 Results: Control of Linear Evolution Equations

Most contemporary distributed control methodologies, including the discrete-time SLS formulations discussed previously, are designed for finite-dimensional state spaces. Yet, many critical cyber-physical systems are governed by partial differential equations (PDEs) that describe continuous spatial dynamics over time—such as thermo-mechanical systems, fluid flows over aerodynamic surfaces, or wave propagation in flexible structures. These systems are mathematically described by evolution equations of the form

$$\begin{aligned} x_{t+1} &= A^* x_t + B^* u_t + w_t^x, \\ y_t &= C^* x_t + w_t^y, \end{aligned} \tag{1.42}$$

where x_t , u_t and y_t belong to (possibly infinite-dimensional) Hilbert spaces X , U , and Y respectively. All A^* , B^* and C^* are linear operators (e.g., a differential operator, an integral operator, or a matrix). The notation A^* means A^* is the adjoint operator of some operator A . The observation $y_t \in Y$ typically is a lower-dimensional representation of $x_t \in X$, for example representing physical sensors.

⁸This means $q_w(\epsilon)$ is lower-bounded by a linear function of ϵ .

The disturbance $w_t^x \in X$ models either actual physical disturbances to the state or model error, and $w_t^y \in Y$ is the sensor noise. The standard engineering practice for controlling PDEs is to perform a spatial discretization (e.g., via finite elements or Galerkin projection) to reduce the system to a massive, finite-dimensional ODE matrix system, and then apply discrete synthesis tools. This "discretize-then-optimize" approach presents two severe drawbacks. First, the resulting system dimension can be massive, rendering even convex control synthesis computationally intractable. Second, discretization can destroy or obfuscate the true continuous spatial locality of the physical system, artificially tightly coupling nodes due to mesh artifacts. In Chapter 9, we address these challenges by elevating the SLS framework to infinite-dimensional function spaces.

Related Literature: Control of Evolution Equations and PDEs. The control of physical processes governed by evolution equations and partial differential equations (PDEs)—such as fluid dynamics, thermo-mechanical heat transfer, and quantum control—poses a theoretical challenge due to the inherently infinite-dimensional nature of their continuous state spaces. The existing literature addressing the control of these infinite-dimensional systems generally falls into three distinct methodological camps.

The most ubiquitous engineering approach to PDE control relies on early spatial discretization. By applying finite element or Galerkin methods, continuous evolution equations are approximated as massive, finite-dimensional systems of ordinary differential equations (ODEs) [58]. While traditional discrete optimal control techniques can then be applied, this "discretize-then-optimize" pipeline frequently generates an unmanageable explosion in the state dimension, making robust constrained synthesis computationally prohibitive.

A second major branch of literature focuses specifically on boundary control techniques, heavily leveraging backstepping methodologies. By employing specialized Volterra integral transformations, researchers have successfully derived closed-form, stabilizing boundary controllers for specific classes of PDE dynamics, such as the reaction-diffusion equation [65]. However, these techniques are often highly tailored to the specific physics of the target system and do not easily generalize to arbitrary structurally constrained or distributed network topologies.

A third, more foundational body of literature surrounds general control theory formulated directly in infinite-dimensional Hilbert or Banach spaces. While classical theoretical works have rigorously established conditions for robust control [71, 66] and stabilizing operators in these spaces [64, 31], there remains a critical need for tractable computational tools that bridge this pure theory with modern, scalable applications. By formulating the SLS problem as a convex optimization over closed-loop operators directly in the continuous Hilbert space, the methods presented in Chapter 9 avoid the state-space explosion of traditional methods while preserving the ability to rigorously

enforce sensor, communication, and locality constraints prior to any spatial discretization.

SLS Formulation. To give a sense to ‘solving (1.42)’, we use a more general notion of solution. An operator A has domain $\mathcal{D}(A)$ and range $\mathcal{R}(A)$. We say $(x_t, u_t)_{t=0}^T$ solves (1.42) *weakly* if for all $f \in \mathcal{D}(A) \cap \mathcal{D}(B) \subseteq X$ and $g \in \mathcal{D}(C) \subseteq Y$,

$$\begin{aligned}\langle x_{t+1}, f \rangle_X &= \langle x_t, Af \rangle_X + \langle u_t, Bf \rangle_U + \langle w_t^x, f \rangle_X, \\ \langle y_t, g \rangle_Y &= \langle x_t, Cg \rangle_X + \langle w_t^y, g \rangle_Y.\end{aligned}$$

Here, f, g are considered test functions, and applying differential operators to test functions instead of the solution itself allows for less regular solutions. As before, we stack the operators A and B into block-off-diagonal operator matrices $\mathcal{A}$ and $\mathcal{B}$, and for a Hilbert space X and discrete times $t \in \{0, \dots, T\}$, we use the script $\mathcal{X} := X^{\otimes(T+1)}$ to denote the $T+1$ dimensional vector of elements $x^0, \dots, x^T \in X$. This is again a Hilbert space when endowed with the inner product $\langle x, \bar{x} \rangle_{\mathcal{X}} := \sum_{t=0}^T \langle x^t, \bar{x}^t \rangle_X$. Throughout, we consider a time horizon $T \in \mathbb{N}$ and denote

$$x := \begin{bmatrix} x^0 \\ \vdots \\ x^T \end{bmatrix} \in \mathcal{X}, \quad w^x := \begin{bmatrix} x^0 \\ w^0 \\ \vdots \\ w^{T-1} \end{bmatrix} \in \mathcal{X}, \quad u = \begin{bmatrix} u^0 \\ \vdots \\ u^T \end{bmatrix} \in \mathcal{U},$$

and operator matrices $\mathcal{A} : \mathcal{D}(\mathcal{A}) \subseteq \mathcal{X} \rightarrow \mathcal{X}$, $K_x : \mathcal{D}(K_x) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ given by

$$\mathcal{A} := \begin{bmatrix} 0 & A & & \\ & \ddots & \ddots & \\ & & 0 & A \\ & & & 0 \end{bmatrix}, \quad K_x := \begin{bmatrix} K_x^{0,0} & \dots & K_x^{T,T} \\ & \ddots & \vdots \\ & & K_x^{T,0} \end{bmatrix}, \quad (1.43)$$

with $\mathcal{B} : \mathcal{D}(\mathcal{B}) \subseteq \mathcal{X} \rightarrow \mathcal{U}$ defined similarly to $\mathcal{A}$, $K_y : \mathcal{D}(K_y) \subseteq \mathcal{U} \rightarrow \mathcal{Y}$ defined similarly to K_x , and $C : \mathcal{D}(C) \subseteq \mathcal{Y} \rightarrow \mathcal{X}$ given by $C := \text{blkdiag}(C)$, that is, the matrix C repeated on the block diagonals; all operator matrices are of size $(T+1) \times (T+1)$. Denote $\overline{\mathcal{D}} := \mathcal{D}(\mathcal{A}) \cap \mathcal{D}(\mathcal{B}) = X \times (\mathcal{D}(A) \cap \mathcal{D}(B))^{\otimes T} \subseteq \mathcal{X}$. Using the full time horizon, the dynamics can be written, for all $f \in \overline{\mathcal{D}}$ and $g \in \mathcal{D}(C)$,

$$\langle x, f \rangle_{\mathcal{X}} = \langle x, \mathcal{A}f \rangle_{\mathcal{X}} + \langle u, \mathcal{B}f \rangle_{\mathcal{U}} + \langle w^x, f \rangle_{\mathcal{X}}, \quad (1.44a)$$

$$\langle y, g \rangle_{\mathcal{Y}} = \langle x, Cg \rangle_{\mathcal{X}} + \langle w^y, g \rangle_{\mathcal{Y}}. \quad (1.44b)$$

In the traditional linear feedback controller setting, the relationship between u and x is given by

$$\langle u, h \rangle_{\mathcal{U}} = \langle x, K_x h \rangle_{\mathcal{X}} \quad \forall h \in \mathcal{D}(K_x) \text{ for state feedback}, \quad (1.45a)$$

$$\langle u, g \rangle_{\mathcal{U}} = \langle y, K_y g \rangle_{\mathcal{Y}} \quad \forall h \in \mathcal{D}(K_y) \text{ for output feedback}. \quad (1.45b)$$

The matrix structure in (1.43) arises by writing the strong form of the dynamics (1.42) in vector notation, and then taking adjoints of the resulting operators with respect to $\langle \cdot, \cdot \rangle_{\mathcal{X}, \mathcal{U}, \mathcal{Y}}$. Note that the controller being causal means K_x^* and K_y^* are lower block-triangular, and hence K_x and K_y are upper block-triangular.

State Feedback. We establish an operator-theoretic SLS parameterization in the state and output feedback settings; in this introduction, we present only the state feedback summary. We define closed-loop operators $\theta_x : \mathcal{D}(\theta_x) \subseteq \mathcal{X} \rightarrow \mathcal{X}$ and $\theta_u : \mathcal{D}(\theta_u) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ to map between the disturbance and the state and input respectively, in block-operator form

$$\theta_x := \begin{bmatrix} \theta_x^{0,0} & \theta_x^{1,1} & \dots & \theta_x^{T,T} \\ & \theta_x^{1,0} & & \theta_x^{T,T-1} \\ & & \ddots & \vdots \\ & & & \theta_x^{T,0} \end{bmatrix}, \quad (1.46)$$

where $\theta_x^{t,\tau} : \mathcal{D}(\theta_x^{t,\tau}) \subseteq X \rightarrow X$ and $\theta_u^{t,\tau} : \mathcal{D}(\theta_u^{t,\tau}) \subseteq U \rightarrow X$ are defined similarly. The adjoints θ_x^* , θ_u^* of the closed-loop maps (CLM) θ_x and θ_u are lower block-triangular to ensure causality, hence θ_x , θ_u are upper block-triangular, and they parameterize the trajectories $x \in \mathcal{X}$ and $u \in \mathcal{U}$ via

$$\langle x, f \rangle_{\mathcal{X}} = \langle w_x, \theta_x f \rangle_{\mathcal{X}} \quad \forall f \in \mathcal{D}(\theta_x) \subseteq \mathcal{X}, \quad (1.47a)$$

$$\langle u, h \rangle_{\mathcal{U}} = \langle w_x, \theta_u h \rangle_{\mathcal{X}} \quad \forall h \in \mathcal{D}(\theta_u) \subseteq \mathcal{U}. \quad (1.47b)$$

In the following theorem, we state one of the main result of this paper: the relationship between the CLM θ_x, θ_u and the state feedback controller K_x is given by

$$\langle f, \theta_x^{-1} \theta_u h \rangle_{\mathcal{X}} = \langle f, K_x h \rangle_{\mathcal{X}}, \quad (1.48)$$

for appropriate test functions f, h . The SLP in Theorem 1.19 parameterizes all possible causal θ_x, θ_u .

Theorem 1.19. Fix disturbance function realization $w^x \in \mathcal{X}$ and operators $\mathcal{A}, \mathcal{B}$.

- (I) If K_x of form (1.43) is given such that $\mathcal{R}(\mathcal{A}), \mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(K_x)$ and $\mathcal{R}(K_x) \subseteq \mathcal{D}$, then any trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ satisfying the closed-loop dynamics (1.44a)-(1.45a) also satisfies (1.47) with some causal CLM θ_x, θ_u satisfying $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\theta_x)$, $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_u)$ and

$$\begin{aligned} \langle f, \theta_x \hat{f} \rangle_{\mathcal{X}} &= \langle f, \theta_x \mathcal{A} \hat{f} \rangle_{\mathcal{X}} + \langle f, \theta_u \mathcal{B} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}}, \\ &\quad \forall f \in \mathcal{X}, \quad \forall \hat{f} \in \overline{\mathcal{D}}. \end{aligned} \quad (1.49)$$

(II) Let $\theta_x : \mathcal{D}(\theta_x) \rightarrow \mathcal{X}$, $\theta_u : \mathcal{D}(\theta_u) \rightarrow \mathcal{X}$ be arbitrary CLM satisfying (1.49) such that θ_x, θ_u are upper block-triangular, $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\theta_x)$, $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_u)$ and $\mathcal{R}(\theta_u) \subseteq \mathcal{D}(\theta_x^{-1})$. Then the corresponding trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ computed with (1.47) also satisfies (1.44a), (1.45a) with the controller K_x defined by $K_x := \theta_x^{-1}\theta_u$ and $\mathcal{D}(K_x) := \mathcal{D}(\theta_u)$.

This theorem shows that optimizing over convex closed-loop maps is equivalent to optimizing over linear feedback controllers, and we show a similar result for the output feedback setting, generalizing the finite-dimensional results in [4, 70] to infinite dimensions. This allows us to add constraints on θ_x and θ_u such as those discussed in Chapter 7, which studies constrained controllability and observability, opening many opportunities for constrained, distributed controllers.

For an integral equation, we show how orthonormal basis functions in the numerical implementation can reduce the problem to a finite-dimensional optimization of coefficients, with subproblems solved in parallel. We show that this also can result in better controllers than discretizing first, further illustrating motivation for the infinite-dimensional formulation.

INTRODUCTION BIBLIOGRAPHY

- [1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Online least squares estimation with self-normalized processes: an application to bandit problems. *arXiv preprint arXiv:1102.2670*, 2011.
- [2] C. D. Aliprantis and K. C. Border. Correspondences. en. In *Infinite Dimensional Analysis: A Hitchhiker's Guide*, pages 555–590. Springer, Berlin, Heidelberg, 1999. ISBN: 9783540295877. DOI: 10.1007/3-540-29587-9_17.
- [3] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, second edition, 2008, pages x+334. ISBN: 978-3-7643-8721-1.
- [4] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, 2019.
- [5] V. S. Borkar. Multiple Timescales. In *Stochastic Approximation*. Volume 48, pages 64–77. Hindustan Book Agency, Gurgaon, 2008. ISBN: 9788185931852. DOI: 10.1007/978-93-86279-38-5_6.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. DOI: 10.1561/22000000016.
- [7] M. Burger, L. He, and C.-B. Schönlieb. Cahn–hilliard inpainting and a generalization for grayvalue images. *SIAM Journal on Imaging Sciences*, 2(4):1129–1167, 2009. DOI: 10.1137/080728548.
- [8] A. Camacho and E. Conover. Manipulation of social program eligibility. *American Economic Journal: Economic Policy*, 3(2):41–65, 2011. ISSN: 1945-7731.
- [9] J. A. Carrillo, Y. Huang, and M. Schmidtchen. Zoology of a nonlocal cross-diffusion model for two species. *SIAM Journal on Applied Mathematics*, 78(2):1078–1104, Jan. 2018. ISSN: 0036-1399, 1095-712X. DOI: 10.1137/17M1128782.
- [10] J. A. Carrillo, R. J. McCann, and C. Villani. Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates. *Revista Matemática Iberoamericana*, 19(3):971–1018, Dec. 2003. ISSN: 0213-2230.
- [11] J. A. Carrillo, R. J. McCann, and C. Villani. Contractions in the 2-wasserstein length space and thermalization of granular media. *Archive for Rational Mechanics and Analysis*, 179(2):217–263, Feb. 1, 2006. ISSN: 1432-0673. DOI: 10.1007/s00205-005-0386-1.
- [12] G. Cavagnari, G. Savaré, and G. E. Sodini. Dissipative probability vector fields and generation of evolution semigroups in Wasserstein spaces. en. *Probability Theory and Related Fields*, 185(3):1087–1182, Apr. 2023. ISSN: 1432-2064. DOI: 10.1007/s00440-022-01148-7.

- [13] A. Cherukuri, B. Gharesifard, and J. Cortés. Saddle-Point Dynamics: Conditions for Asymptotic Stability of Saddle Points. en. *SIAM Journal on Control and Optimization*, 55(1):486–511, Jan. 2017. ISSN: 0363-0129, 1095-7138. DOI: 10.1137/15M1026924.
- [14] M. Chipot, S. Hastings, and D. Kinderlehrer. Transport in a molecular motor system. eng. *ESAIM: Mathematical Modelling and Numerical Analysis*, 38(6):1011–1034, Mar. 2010.
- [15] L. Chizat, A. Delalande, and T. Vaškevičius. Sharper exponential convergence rates for Sinkhorn’s algorithm in continuous settings. en. *Mathematical Programming*, 215(1):809–858, Jan. 2026. ISSN: 1436-4646. DOI: 10.1007/s10107-025-02242-z.
- [16] L. Conger, F. Hoffmann, E. Mazumdar, and L. Ratliff. Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows. en. *Advances in Neural Information Processing Systems*, 36:45971–46006, Dec. 2023.
- [17] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games, Nov. 2024. DOI: 10.48550/arXiv.2411.07403. arXiv:2411.07403 [math].
- [18] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Monotone Multispecies Flows, June 2025. DOI: 10.48550/arXiv.2506.22947. arXiv:2506.22947.
- [19] L. Conger, A. P. Leeman, and F. Hoffmann. Convex Constrained Controller Synthesis for Evolution Equations. In *2025 American Control Conference (ACC)*, pages 1169–1176, July 2025. DOI: 10.23919/ACC63710.2025.11108072. ISSN: 2378-5861.
- [20] L. Conger, Y. Lin, A. Wierman, and E. Mazumdar. Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pages 8491–8497, Dec. 2024. DOI: 10.1109/CDC56724.2024.10886635. ISSN: 2576-2370.
- [21] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *NeurIPS*, volume 26. Curran Associates, Inc., 2013.
- [22] J. M. Danskin. *The Theory of Max-Min and its Application to Weapons Allocation Problems*, volume 5 of *Econometrics and Operations Research*. Springer, Berlin, Heidelberg, 1967. ISBN: 9783642460944. DOI: 10.1007/978-3-642-46092-0.
- [23] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020.
- [24] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. *Advances in Neural Information Processing Systems*, 31:4188–4197, 2018.
- [25] T. Debiec, M. Mandal, and M. Schmidtchen. From finite to continuous phenotypes in (visco-)elastic tissue growth models. *Journal of Differential Equations*, 437:113375, Aug. 2025. DOI: 10.1016/j.jde.2025.113375.
- [26] M. Doumic, S. Hecht, B. Perthame, and D. Peurichard. Multispecies cross-diffusions: From a nonlocal mean-field to a porous medium system without self-diffusion. *Journal of Differential Equations*, 389:228–256, Apr. 2024. DOI: 10.1016/j.jde.2024.01.017.

- [27] Y. Du and J. S. Li. Distributed Continuous-Time Control via System Level Synthesis, Apr. 2026. DOI: 10.48550/arXiv.2410.08135. arXiv:2410.08135.
- [28] B. Düring, P. Markowich, J.-F. Pietschmann, and M.-T. Wolfram. Boltzmann and Fokker-Planck equations modelling opinion formation in the presence of strong leaders. en. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2112):3687–3708, Dec. 2009. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2009.0239.
- [29] A. Durmus and É. Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [30] E. E. Espejo, A. Stevens, and J. J. L. Velazquez. Simultaneous finite time blow-up in a two-species model for chemotaxis. *Analysis*, 29(3), Jan. 2009. ISSN: 0174-4747. DOI: 10.1524/anly.2009.1029.
- [31] J. S. Gibson. The Riccati Integral Equations for Optimal Control Problems on Hilbert Spaces. en. *SIAM Journal on Control and Optimization*, 17(4):537–565, 1979. ISSN: 0363-0129, 1095-7138. DOI: 10.1137/0317039.
- [32] V. Giunta, T. Hillen, M. Lewis, and J. R. Potts. Local and global existence for nonlocal multispecies advection-diffusion models. *SIAM Journal on Applied Dynamical Systems*, 21(3):1686–1708, Sept. 2022. ISSN: 1536-0040. DOI: 10.1137/21M1425992.
- [33] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. *International Journal of Computer Vision*, 60(3):225–240, 2004. ISSN: 1573-1405. DOI: 10.1023/B:VISI.0000036836.66311.97.
- [34] M. Hardt, E. Mazumdar, C. Mendler-Dünner, and T. Zrnic. Algorithmic Collective Action in Machine Learning. en. In *Proceedings of the 40th International Conference on Machine Learning*, pages 12570–12586. PMLR, July 2023.
- [35] S. Hastings, D. Kinderlehrer, and J. B. McLeod. Diffusion Mediated Transport in Multiple State Systems. en. *SIAM Journal on Mathematical Analysis*, 39(4):1208–1230, Jan. 2008. DOI: 10.1137/060650994.
- [36] E. Jensen. *Topics in Optimal Distributed Control*. Thesis, University of California Santa Barbara, 2020.
- [37] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck Equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [38] R. E. Kalman, Y. C. Ho, and K. C. Narendra. Controllability of Linear Dynamical Systems. *Contributions to Differential Equations*, 1(1), 1961.
- [39] R. E. Kalman. Contributions to the theory of optimal control. *Boletín de la Sociedad Matemática Mexicana*, 5(2):102–119, 1960.
- [40] D. Kinderlehrer and M. Kowalczyk. Diffusion-Mediated Transport and the Flashing Ratchet. *Archive for Rational Mechanics and Analysis*, 161(2):149–179, Feb. 2002. DOI: 10.1007/s002050100173.

- [41] A. Kiselev, F. Nazarov, L. Ryzhik, and Y. Yao. Chemotaxis and reactions in biology. en. *Journal of the European Mathematical Society*, 25(7):2641–2696, June 2022. DOI: 10.4171/jems/1247.
- [42] D. Kwon, Y. Fan, and K. Lee. Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance. en. *Advances in Neural Information Processing Systems*, 35:20205–20217, Dec. 2022.
- [43] G. Li and Y. Yao. Two-species competition model with chemotaxis: well-posedness, stability and dynamics. *Nonlinearity*, 35(3):1329, Feb. 3, 2022. ISSN: 0951-7715. DOI: 10.1088/1361-6544/ac4a8d.
- [44] J. S. Li and C. Amo Alonso. Global Performance Guarantees for Localized Model Predictive Control. *IEEE Open Journal of Control Systems*, 2:325–336, 2023. ISSN: 2694-085X. DOI: 10.1109/OJCSYS.2023.3308009.
- [45] S. Liu, H. Sun, and H. Zha. A Particle-Evolving Method for Approximating the Optimal Transport Plan. en. In F. Nielsen and F. Barbaresco, editors, *Geometric Science of Information*, pages 878–887, Cham. Springer International Publishing, 2021. ISBN: 9783030802097. DOI: 10.1007/978-3-030-80209-7_94.
- [46] Y.-A. Ma, Y. Chen, C. Jin, N. Flammarion, and M. I. Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885, 2019.
- [47] J. Maas, M. Rumpf, C. Schönlieb, and S. Simon. A generalized model for optimal transport of images including dissipation and density modulation. en. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1745–1769, 2015. DOI: 10.1051/m2an/2015043.
- [48] F. Marcellini. ODE-PDE models in traffic flow dynamics. en. *Bulletin of the Brazilian Mathematical Society, New Series*, 47(2):533–544, June 2016. ISSN: 1678-7714. DOI: 10.1007/s00574-016-0167-5.
- [49] R. J. McCann. A convexity principle for interacting gases. *Advances in Mathematics*, 128(1):153–179, June 1, 1997. ISSN: 0001-8708. DOI: 10.1006/aima.1997.1634.
- [50] A. Mielke and T. Roubíček. *Rate-Independent Systems: Theory and Application*. Applied Mathematical Sciences. Springer New York, 2015. ISBN: 9781493927067.
- [51] J. P. Miller, J. C. Perdomo, and T. Zrnic. Outside the echo chamber: optimizing the performative risk. In *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning, pages 7710–7720. PMLR, July 1, 2021.
- [52] R. H. Nochetto and G. Savaré. Nonlinear evolution governed by accretive operators in Banach spaces: error control and applications. *Math. Models Methods Appl. Sci.*, 16(3):439–477, 2006. ISSN: 0218-2025.
- [53] R. H. Nochetto, G. Savaré, and C. Verdi. A posteriori error estimates for variable time-step discretizations of nonlinear evolution equations. *Comm. Pure Appl. Math.*, 53(5):525–589, 2000. ISSN: 0010-3640.

- [54] F. Otto and C. Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, June 1, 2000. ISSN: 0022-1236. DOI: 10.1006/jfan.1999.3557.
- [55] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm. Partial Differential Equations*, 26(1-2):101–174, 2001. ISSN: 0360-5302.
- [56] B. Perthame. *Parabolic Equations in Biology: Growth, reaction, movement and diffusion*. Lecture Notes on Mathematical Modelling in the Life Sciences. Springer International Publishing, Cham, 2015. ISBN: 9783319194998. DOI: 10.1007/978-3-319-19500-1.
- [57] J. Peypouquet. *Convex Optimization in Normed Spaces: Theory, Methods and Examples*. en. SpringerBriefs in Optimization. Springer International Publishing, Cham, 2015. ISBN: 9783319137094. DOI: 10.1007/978-3-319-13710-0.
- [58] A. Quarteroni and A. Valli. *Numerical approximation of partial differential equations*, volume 23 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 1994.
- [59] J. Rabin and N. Papadakis. Convex color image segmentation with optimal transport distances. In J.-F. Aujol, M. Nikolova, and N. Papadakis, editors, *Scale Space and Variational Methods in Computer Vision*, pages 256–269, Cham. Springer International Publishing, 2015. ISBN: 978-3-319-18461-6.
- [60] F. Santambrogio. Dealing with Moment Measures via Entropy and Optimal Transport. *Journal of Functional Analysis*, 271(2):418–436, 2016. DOI: <https://doi.org/10.1016/j.jfa.2016.04.009>.
- [61] S. Shin, Y. Lin, G. Qu, A. Wierman, and M. Anitescu. Near-optimal distributed linear-quadratic regulator for networked systems. *SIAM Journal on Control and Optimization*, 61(3):1113–1135, 2023.
- [62] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht. Learning without mixing: towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- [63] S. Simons. Minimax Theorems and Their Proofs. en. In D.-Z. Du and P. M. Pardalos, editors, *Minimax and Applications*, pages 1–23. Springer US, Boston, MA, 1995. ISBN: 9781461335573. DOI: 10.1007/978-1-4613-3557-3_1.
- [64] M. Slemrod. A Note on Complete Controllability and Stabilizability for Linear Control Systems in Hilbert Space. en. *SIAM Journal on Control*, 12(3):500–508, 1974. ISSN: 0036-1402. DOI: 10.1137/0312038.
- [65] R. Vazquez and M. Krstic. Backstepping for pdes. In J. Baillieul and T. Samad, editors, *Encyclopedia of Systems and Control*, pages 1–4. Springer London, 2019. ISBN: 978-1-4471-5102-9. DOI: 10.1007/978-1-4471-5102-9_100023-1.
- [66] S. R. Venkatesh, A. Megretski, and M. A. Dahleh. On robust control synthesis and analysis in a Hilbert space. *Systems & Control Letters*, 39(1):1–12, 2000. ISSN: 0167-6911. DOI: 10.1016/S0167-6911(99)00078-X.

- [67] C. Villani. *Optimal transport, old and new*. Lecture Notes for the 2005 Saint-Flour summer school. Springer, 2007.
- [68] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003, pages xvi+370. ISBN: 0-8218-3312-X. DOI: 10.1090/gsm/058.
- [69] G. Wang and L. Chizat. Local convergence of mean-field Langevin dynamics: from gradient flows to linearly monotone games, Feb. 2026. DOI: 10.48550/arXiv.2602.11999. arXiv:2602.11999.
- [70] Y.-S. Wang, N. Matni, and J. C. Doyle. Separable and Localized System-Level Synthesis for Large-Scale Systems. *IEEE TAC*, 63(12):4234–4249, 2018. ISSN: 1558-2523. DOI: 10.1109/TAC.2018.2819246.
- [71] J. T.-Y. Wen and M. J. Balas. Robust adaptive control in Hilbert Space. *Journal of Mathematical Analysis and Applications*, 143(1):1–26, 1989. ISSN: 0022-247X. DOI: 10.1016/0022-247X(89)90025-5.
- [72] Y. Zhang, R. Cheng, and Y. Xia. Observability Blocking for Functional Privacy of Linear Dynamic Networks. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 7469–7474, Dec. 2023. DOI: 10.1109/CDC49753.2023.10384211. ISSN: 2576-2370.
- [73] T. Zrnic, E. Mazumdar, S. Sastry, and M. Jordan. Who leads and who follows in strategic classification? In *Advances in Neural Information Processing Systems*, volume 34, pages 15257–15269. Curran Associates, Inc., 2021.

Part I

MULTISPECIES GRADIENT FLOW ANALYSIS

*Chapter 2***COUPLED WASSERSTEIN GRADIENT FLOWS FOR MIN-MAX AND
COOPERATIVE GAMES**

L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games, Nov. 2024. DOI: 10.48550/arXiv.2411.07403. arXiv:2411.07403 [math]

ABSTRACT

We propose a framework for two-player infinite-dimensional games with cooperative or competitive structure. These games take the form of coupled partial differential equations in which players optimize over a space of measures, driven by either a gradient descent or gradient descent-ascent in Wasserstein-2 space. We characterize the properties of the Nash equilibrium of the system, and relate it to the steady state of the dynamics. In the min-max setting, we show, under sufficient convexity conditions, that solutions converge exponentially fast and with explicit rate to the unique Nash equilibrium. Similar results are obtained for the cooperative setting. We apply this framework to distribution shift induced by interactions among a strategic population of agents and an algorithm, proving additional convergence results in the timescale-separated setting. We illustrate the performance of our model on (i) real data from an economics study on Colombia census data, (ii) feature modification in loan applications, and (iii) performative prediction. The numerical experiments demonstrate the importance of distribution-level, rather than moment-level, modeling.

2.1 Introduction

Gradient flows generalize gradient descent to infinite dimensional spaces, for instance when probability distributions evolve in the direction of steepest descent for a given functional and metric. Recent years have seen a surge in the application of gradient flow theory including work in Monte Carlo sampling [96, 156], generative modeling [140], image processing (image registration, warping, shape classification, image segmentation and image restoration) [123, 186, 158, 40], modeling the behavior of plastic materials in material science [166], or biological systems [183, Ch 8]. The majority of the gradient flow literature considers a single species, that is, the evolution of a single distribution, yet in many applications settings multiple distributions evolve simultaneously. In these *multispecies* systems, each species still minimizes its own functional, but this functional can depend on the other species. This interdependence results in a system of coupled PDEs that globally can no longer be viewed as a gradient flow on a single functional. Due to the loss of the gradient flow structure, this class of PDEs can exhibit complex dynamics like cycling or chaos. Such systems model many types of phenomena, ranging from chemotaxis [100, 126, 210] to opinion dynamics [95]. More recently, such coupled structures have been shown to arise also in machine learning in the literature on distributionally robust optimization [151, 93, 83, 139, 157] and distribution shift in machine learning [62]. In such settings, natural questions that arise include the existence and characterization of solutions and steady states, and convergence of the coupled PDE system to these steady states (if they exist). Current results focus on proving existence of solutions for some of these systems [118, 131, 88, 54, 4], but less is known about their long-time behavior.

In this work we study multispecies systems that arise from competitive and cooperative games. In these games the two species either seek to maximize the same functional, or compete with one maximizing and the other minimizing the functional—i.e., they are infinite dimensional players in a cooperative or min-max game. We show that the steady states of the resulting system of partial differential equations (PDEs) coincide with the equilibria of the underlying games and characterize their rates of convergence under natural structural assumptions on the corresponding functionals. The convergence of Wasserstein-2 gradient flows for min-max problems over spaces of measures was recently posed as an open question in [208], and our analysis provides an answer to these questions for displacement convex-concave functionals over unbounded sets.

We use these results to investigate the long-term effects of strategic interactions in driving distribution shift in real-world machine learning contexts. In many machine learning systems, agents whose data is analyzed by the system are incentivized to manipulate their data to achieve a desired output. Additionally, distribution shift can occur naturally, or agents share information that causes other players to evolve. This behavior is not well-understood and has become a subject of recent interest; see for instance [170, 128, 209, 146, 169, 1, 182]. In settings where the objective of the learning

algorithm opposes that of the agents, the update process can be modeled as a min-max problem over a large number of agents, which in a mean-field limit can be analyzed as an optimization problem over measures. In particular, we incorporate intra-agent interactions in the model via an interaction potential, exogenous shifts, and strategic responses to the algorithm. We illustrate how these model components capture rich distributional behavior (see Section 2.5.1) and can show disparate effects of retraining among subpopulations (see Section 2.5.2). The implementation in Section 2.5.2 uses a particle method, highlighting how real-world economic settings of many agents can be analyzed via a mean-field description at the PDE level.

In this work, we consider two-species systems with an energy functional containing potential terms, self-interaction kernels, diffusion, and a coupling term which is linearly dependent on both species. Each species evolves according to a Wasserstein-2 gradient flow with respect to its own energy differential, in the direction of steepest ascent or descent. The resulting dynamical system is a joint gradient flow in the setting where both species descend, or a gradient descent-ascent flow, in the setting of opposing dynamics. In line with intuitive notions from game theory, we call the joint gradient flow setting the *cooperative* setting, because one can view the resulting dynamics as a game in which both players aim to achieve the same objective of minimizing the same energy. We name the gradient descent-ascent case the *competitive* setting, due to the zero-sum game structure in which one player aims to maximize a function which the other player aims to minimize. Let the energy in the cooperative setting be defined as $F_a : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \rightarrow \mathbb{R}$ and in the competitive setting as $F_c : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \rightarrow \mathbb{R}$, where $\mathcal{P}(\mathbb{R}^d)$ is the space of probability measures on $\mathbb{R}^d$,

$$F_a(\rho, \mu) = \iint f(z, x) d\rho(z) d\mu(x) + \mathcal{R}(\rho) + \mathcal{U}(\mu), \quad (2.1a)$$

$$F_c(\rho, \mu) = \iint f(z, x) d\rho(z) d\mu(x) - \mathcal{R}(\rho) + \mathcal{U}(\mu), \quad (2.1b)$$

where

$$\begin{aligned} \mathcal{R}(\rho) &= \alpha \mathcal{H}(\rho) + \frac{1}{2} \int W_1 * \rho(z) d\rho(z) + \int V_1(z) d\rho(z), \\ \mathcal{U}(\mu) &= \beta \mathcal{H}(\mu) + \frac{1}{2} \int W_2 * \mu(x) d\mu(x) + \int V_2(x) d\mu(x), \end{aligned}$$

with $\alpha, \beta \geq 0$. Here, we denote by $f : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$ the function governing coupling forces between the species ρ and μ , by $\mathcal{H}(\eta) : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ the entropy functional

$$\mathcal{H}(\eta) = \begin{cases} \int \eta \log \eta & \text{if } \eta \ll \mathcal{L}^d \\ +\infty & \text{otherwise} \end{cases},$$

for $\mathcal{L}^d$ the d -dimensional Lebesgue measure, by $V_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ external potentials, and by $W_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ interaction potentials. Further, $\mathcal{W}_2$ denotes the Wasserstein-2 metric, and $\nabla_{\mathcal{W}_2, \eta} F$ the

Wasserstein-2 gradient of F with respect to η . The dynamics in the cooperative setting are

$$\partial_t \rho = -\nabla_{\mathcal{W}_{2,\rho}} F_a(\rho, \mu), \quad \partial_t \mu = -\nabla_{\mathcal{W}_{2,\mu}} F_a(\rho, \mu). \quad (2.2)$$

In the competitive case, the dynamics are

$$\partial_t \rho = \nabla_{\mathcal{W}_{2,\rho}} F_c(\rho, \mu), \quad \partial_t \mu = -\nabla_{\mathcal{W}_{2,\mu}} F_c(\rho, \mu). \quad (2.3)$$

The analysis of the dynamics in the cooperative setting (2.2) proceeds similarly to the approach in [50], in which an HWI inequality is proven for a single species and log-Sobolev and Talagrand inequalities follow. However, because the dynamics in the competitive setting (2.3) no longer have a gradient flow structure, the classical gradient flow techniques no longer apply.

In this paper we show, under sufficient convexity conditions, that the competitive dynamics (2.3) converge exponentially fast in the joint Wasserstein-2 metric, with the rate dependent on the displacement convexity of F_c with respect to μ and displacement concavity of F_c with respect to ρ . We prove that any two solution pairs (ρ, μ) and $(\tilde{\rho}, \tilde{\mu})$ to (2.3) contract in $\mathcal{W}_2^2 \times \mathcal{W}_2^2$, the squared joint Wasserstein-2 metric. Based on this result, we then show existence of a unique steady state for the dynamics. In order to show uniform boundedness of the second moments for both species, we show that they converge exponentially fast to a ball, and then remain inside that ball for all time. Finally, we show that the steady state is in fact a critical point of F_c , and the unique Nash equilibrium. While the convexity and smoothness assumptions can be generalized, even mild relaxations on the lower-bounds in finite dimensions do not give the same guarantees. For example, in Euclidean space, assuming that the energy satisfies a Polyak Łojasiewicz (PL) condition instead of convexity results in non-uniqueness of Nash equilibria. With respect to convexity, our results mirror the state-of-the-art guarantees existing for finite-dimensional games. However, relaxing the regularity assumptions on the functionals is likely possible, and an interesting direction of future research.

2.1.1 Related Literature

The existence and convergence results utilize tools from long-time behavior analysis of PDEs and game theory, and apply to multispecies PDEs with a cooperative or competitive structure. Our results extend what is known in min-max problems with applications in machine learning, DRO, and strategic distribution shifts.

Coupled PDEs The structure of the models (2.2) and (2.3) is closely related to multispecies PDEs in a variety of application areas, including chemotaxis [100, 126, 211, 210], opinion formation [95], pedestrian dynamics [12], population biology [58, 87] and cell-cell adhesion [105]. Recent progress focuses on well-posedness questions, connections to related models via limiting procedures, and asymptotic pattern formation; see for instance [112, 39, 91, 19, 119, 88, 86, 54, 4] and references

therein. The mathematical theory for multispecies PDEs is still nascent, and even if the equations exhibit a gradient flow structure, general results characterizing the asymptotic behavior of solutions (especially results achieving exponential convergence with explicit rates) are rare and concern rather special modeling choices. In [49], a numerical method for computing solutions to two-species non-local cross-diffusion models is used to analyze the steady states of these systems; we build upon this by presenting convergence analysis. The two-species model in [94] considers potential terms in addition to the cross-diffusion and self-interaction terms, which is closer to our model, and proves a mean-field limit from the particle stochastic differential equation system to the PDE limit. In [91], the authors prove conditions under which measures converge to Diracs under cross-diffusion, without any self-diffusion. We allow for the possibility of diffusion which we show leads to Lebesgue-measurable steady states with support over the entire space.

Game Theory Although the dynamical system we study is a system of PDEs, the steady state of the system can be analyzed through the lens of game theory and optimization. In game theory, existence of equilibria in min-max problems has been studied over spaces of (1) deterministic strategies or (2) probabilistic strategies over compact sets [120, 199]. Recent progress in developing algorithms to solve infinite-dimensional min-max optimization problems includes [144], which proposes a mirror ascent-descent scheme to compute a solution to a min-max problem over convex sets of measures and assumes existence of an equilibrium, while [207, 116, 154, 152] propose other gradient ascent-descent schemes either under the assumption that a Nash equilibrium exists or that the optimization is over a convex set. We build upon these works by proving existence of a unique Nash equilibrium for a general class of energy functionals over unbounded sets of measures, and prove that gradient ascent-descent in Wasserstein-2 converges exponentially to this equilibrium. This addresses the strongly displacement convex-concave setting of the open problem posed in [208].

Applications in Machine Learning Solving min-max games is an important problem for many applications in machine learning which can be formulated as min-max problems over the space of distributions, such as distributionally robust optimization (DRO) [151], strategic distribution shift [62, 182, 215], and generative adversarial networks (GANs) [2, 121]. Existing results in DRO [151, 93, 83, 139, 157] for machine learning considers optimization over a *bounded* set of measures. Our setting builds upon this by considering optimization over an *unbounded* set of measures. Other applications of DRO include portfolio selection [212, 111] and train freight optimization [11]. In support of developing numerical methods, duality structures of distributionally robust optimization problems are studied [115, 214] as well as sensitivity analysis [17]. By connecting the gradient-flow structure of min-max problems to multispecies PDEs, analysis tools and numerical methods from the PDE literature become available for these applications.

2.1.2 Contributions

Our contributions in this work are twofold; first on the theoretical side, our theorems extend what is known in the PDE literature and optimization literature. Secondly, we apply our framework to illustrate the importance of modeling distribution shift in strategic populations in machine learning.

Theory This framework sits at the intersection of PDE analysis and optimization, providing contributions to each of these fields. From an optimization perspective, the existence of a Nash equilibrium over measures on unbounded sets has been an open question. Since existence is unknown, there are no systematic tools for computing equilibria and in particular, convergence of gradient descent-ascent to an equilibrium is an open problem [208]. We expand this area of game theory in two key ways.

1. Classical proofs for existence of Nash equilibria assume optimization over compact spaces of measures; we prove results without this assumption by showing contraction in $\mathcal{P}_2 \times \mathcal{P}_2$.
2. By analyzing distributions over action spaces rather than deterministic actions, the achieved equilibrium is not restricted to a pure Nash equilibrium; it can be a mixed Nash equilibrium. Outside of specific games, such as ones with a finite number of actions or structure that allows direct computation via calculus of variations, computing mixed Nash equilibria over continuous action spaces is difficult to solve in the general setting. Our results suggest that the gradient ascent-descent structure in Wasserstein-2 offers a solution.

From a PDE perspective, this setting opens the door for using techniques from calculus of variations and gradient flows in metric spaces. In particular, it can be framed in the language of multi-species systems, a field for which only very few and recent results exist on asymptotics via entropy methods. We show the existence of a unique steady state and exponential convergence to it with explicit rates in four different two-species settings. This extends what is known about long-time asymptotics for systems of coupled PDEs; in particular, the technical contributions include the following.

1. In the cooperative setting, classical functional inequalities are extended to the case of multiple species.
2. In the competitive setting, convergence is proven without the use of timescale separation; this requires entirely different proof techniques both for existence of the steady state and convergence. Direct differentiation of $\mathcal{W}_2$ results in convergence, and existence of a unique Nash equilibrium is shown via contraction. A dynamical systems-type argument is used for uniform estimates of the second moments.

3. We demonstrate in a particular application setting how a Danskin-type result (also known as an envelope theorem in analysis) can be obtained from basic assumptions using a Γ -convergence argument (see Proposition 2.55). This removes a key assumption in [62] on the differentiability of the best response (see [62, Lemma 29]). Such a Γ -convergence approach is expected to generalize to other choices of functionals.

Application One particular setting in which models of type (2.2) and (2.3) appear is when machine learning algorithms interact with strategic populations [62]. In many real-world settings, populations dynamically adapt their strategy based on algorithm behavior. Optimization methods for algorithms do not usually account for this data manipulation, and we provide examples illustrating why modeling distribution shift in the face of learning is critical for improved performance.

1. We illustrate our model on real data from an economics study (see Section 2.5.1), a setting in which agents manipulated data in response to the action of an algorithm, showing that our model is able to accurately capture such behavior.
2. We show the importance of modeling distribution shift in detail. A state-of-the-art performative prediction method, based on mean shift, is outperformed when the classifier follows a simple gradient descent scheme. We also illustrate how modeling population interactions can be overlooked when looking at classifier accuracy, but these interaction terms matter when considering classifier performance on certain subpopulations.

2.1.3 Paper Structure

In Section 2.2, we provide relevant definitions and notation. Section 2.3 contains the key assumptions and main results for the cooperative and competitive settings. In Section 2.4, we discuss an application of the model to strategic distribution shift in machine learning, with timescale-separated convergence results. Numeric examples and insights are shown in Section 2.5. The proof of the key convergence results in the cooperative and competitive settings are postponed to Sections 2.6 and 2.7 respectively. Appendices 2.A and 2.B contain proofs for the timescale-separated settings and in Appendix 2.C we collect the supporting technical lemmas.

2.2 Preliminaries

This section provides definitions and notation used throughout the paper. I_d denotes the $d \times d$ identity matrix, and id denotes the identity map. $\text{Hess}(f)$ denotes the Hessian of f in all variables, while $\nabla_x^2 f$ denotes the Hessian of f in the variable x only. The notation $\mathbb{I}\{B\}$ is an indicator function for the set B . Unless otherwise specified, $\|\cdot\|$ notes the Euclidean norm for vectors and $\|\cdot\|_2$ is the induced 2-norm when the argument is a matrix. Let $L_+^1(\mathbb{R}^d) = \{\mu \in L^1(\mathbb{R}^d) : \mu \geq 0 \text{ a.e.}\}$. The

narrow topology is defined as convergence in duality with continuous bounded functions and the weak-* topology is defined in duality with continuous functions vanishing at infinity. Throughout the manuscript, we will use the following related notion of convergence, which we refer to as *weak topology*.

Definition 2.1 (Weak Convergence). *A sequence of measures (ρ_n) converges in the weak topology, denoted by $\rho_n \rightharpoonup \rho$, when (ρ_n) converges narrowly, that is, in duality with all continuous bounded functions, and there exists a uniform second moment bound for (ρ_n) .*

Note that weak convergence implies narrow convergence, which implies weak-* convergence; the converse however does not hold. If a sequence converging weak-* also has a uniform second moment bound, then tightness follows from Markov's inequality [117] and so the masses also converge; consequently, the sequence also converges narrowly and weakly.

The energy functionals we are considering are usually defined on the set of probability measures on $\mathbb{R}^d$, denoted by $\mathcal{P}(\mathbb{R}^d)$. At times we abbreviate this as $\mathcal{P}$ if the underlying space is clear from context. The set $\mathcal{P}^{ac}(\mathbb{R}^d)$ denotes the set of probability measures on $\mathbb{R}^d$ that are absolutely continuous with respect to the Lebesgue measure. Throughout, we use the same notation for measures in $\mathcal{P}(\mathbb{R}^d)$ and their densities with respect to the Lebesgue measure. We also use $\tilde{\mathcal{P}}(\mathbb{R}^d) := \mathcal{P}^{ac}(\mathbb{R}^d) \cup \{\rho \in \mathcal{P}(\mathbb{R}^d) : \rho = \delta_z \text{ for some } z \in \mathbb{R}^d\}$. If we consider the subset $\mathcal{P}_2(\mathbb{R}^d)$ of probability measures with bounded second moment,

$$\mathcal{P}_2(\mathbb{R}^d) := \left\{ \rho \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|z\|^2 d\rho(z) < \infty \right\},$$

then we can endow this space with the Wasserstein-2 metric,

$$\mathcal{W}_2(\mu, \nu)^2 = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|z - z'\|^2 d\gamma(z, z')$$

where $\Gamma(\mu, \nu) \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$ is the set of all joint probability distributions with bounded second moments and marginals μ, ν , i.e. $\mu(dz) = \int \gamma(dz, z') dz'$ and $\nu(dz') = \int \gamma(z, dz') dz$. Throughout this paper, we set $z \in \mathbb{R}^{d_1}$ and $x \in \mathbb{R}^{d_2}$, and denote by $\overline{W}$ the joint Wasserstein metric.

Definition 2.2 (Joint Wasserstein Metric). *Denote by $\overline{W}$ the metric over $\mathcal{P}_2(\mathbb{R}^{d_1}) \times \mathcal{P}_2(\mathbb{R}^{d_2})$ given by*

$$\overline{W}((\rho, \mu), (\tilde{\rho}, \tilde{\mu}))^2 = \mathcal{W}_2(\rho, \tilde{\rho})^2 + \mathcal{W}_2(\mu, \tilde{\mu})^2$$

for all pairs $(\rho, \mu), (\tilde{\rho}, \tilde{\mu}) \in \mathcal{P}_2(\mathbb{R}^{d_1}) \times \mathcal{P}_2(\mathbb{R}^{d_2})$.

Geodesic convexity in the Wasserstein-2 space $(\mathcal{P}_2(\mathbb{R}^{d_i}), \mathcal{W}_2)$ is known as displacement convexity.

Definition 2.3 (Displacement Convexity [162]). *A functional $G : \mathcal{P}(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ is displacement convex if for all $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^{d_i})$ that are atomless we have*

$$G(\rho_s) \leq (1-s)G(\rho_0) + sG(\rho_1),$$

where $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ is the displacement interpolant between ρ_0 and ρ_1 for all $s \in [0, 1]$. Further, $G : \mathcal{P}(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ is uniformly displacement convex with constant $\lambda > 0$ if

$$G(\rho_s) \leq (1-s)G(\rho_0) + sG(\rho_1) - s(1-s)\frac{\lambda}{2}\mathcal{W}_2(\rho_0, \rho_1)^2,$$

where $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ is the displacement interpolant between ρ_0 and ρ_1 .

Remark 2.4. *In other words, G is displacement convex (concave) if the function $G(\rho_s)$ is convex (concave) in $s \in [0, 1]$ with $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ being the displacement interpolant between ρ_0 and ρ_1 . Contrast this with the classical notion of convexity (concavity) for G , where we require that the function $G((1-s)\rho_0 + s\rho_1)$ is convex (concave). One can think of displacement convexity for an energy functional defined on $\mathcal{P}_2$ as convexity along the shortest path in the Wasserstein-2 metric (linear interpolation in the Wasserstein-2 space) between any two given probability distributions.*

We will use s to denote the interpolation parameter for geodesics, and t to denote time related to solutions of (2.2)-(2.3). In fact, if the energy G is twice differentiable along geodesics, then the condition $\frac{d^2}{ds^2}G(\rho_s) \geq 0$ along any geodesic $(\rho_s)_{s \in [0,1]}$ between ρ_0 and ρ_1 is sufficient to obtain displacement convexity. Similarly, when $\frac{d^2}{ds^2}G(\rho_s) \geq \lambda\mathcal{W}_2(\rho_0, \rho_1)^2$, then G is uniformly displacement convex with constant $\lambda > 0$. For more details, see [162] and [204, Chapter 5.2].

Definition 2.5 (Relative Energy). *The relative energy of a functional G is given by $G(\gamma|\gamma_\infty) = G(\gamma) - G(\gamma_\infty)$, where $G(\gamma_\infty)$ is the energy at some reference measure γ_∞ .*

Using the first variation, we can express the gradient in Wasserstein-2 space, see for example [204, Exercise 8.8]. More precisely, the gradient of an energy $G : \mathcal{P}_2(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ in the Wasserstein-2 space is given by

$$\nabla_{\mathcal{W}_2}G(\rho) = -\text{div}(\rho\nabla\delta_\rho G[\rho](x)),$$

where $\delta_\rho G[\rho](x)$ denotes the first variation of G at ρ (if it exists). As a consequence, the infinite dimensional steepest descent in Wasserstein-2 space of a given energy $G : \mathcal{P}_2 \rightarrow \mathbb{R} \cup \{+\infty\}$ can be expressed as the PDE

$$\partial_t \rho = -\nabla_{\mathcal{W}_2}G(\rho) = \text{div}(\rho\nabla\delta_\rho G[\rho]) . \quad (2.4)$$

All the coupled gradient flows considered in this work have this Wasserstein-2 structure.

Steady states The main goal in our theoretical analysis is to characterize the asymptotic behavior for the models (2.2) and (2.3) as time goes to infinity. The steady states of these equations are the natural candidates to be asymptotic profiles for the corresponding dynamics. Thanks to the gradient flow structure, we expect to be able to make a connection between critical points of the energy functionals and the steady states of the corresponding gradient ascent/descent dynamics. More precisely, any minimizer or maximizer is in particular a critical point of the energy, and therefore satisfies that the first variation is constant on disconnected components of its support. If this ground state also has enough regularity (weak differentiability) to be a solution to the equation, it immediately follows that it is in fact a steady state. To make this connection precise, we first introduce what exactly we mean by a steady state.

Definition 2.6 (Steady states). *For $\rho_\infty, \mu_\infty \in \mathcal{P}_2$, the pair $(\rho_\infty, \mu_\infty)$ is a steady state for the systems (2.2)-(2.3) if*

- (i) $\nabla W_1 * \rho_\infty \in L^1_{loc}(\mathbb{R}^{d_1})$, $\nabla W_2 * \mu_\infty \in L^1_{loc}(\mathbb{R}^{d_2})$,
- (ii) if additionally, $\alpha > 0$, then $\rho_\infty \in W^{1,2}_{loc}(\mathbb{R}^{d_1}) \cap L^1_+(\mathbb{R}^{d_1}) \cap L^\infty_{loc}(\mathbb{R}^{d_1})$ and $\|\rho_\infty\|_1 = 1$,
- (iii) if additionally, $\beta > 0$, then $\mu_\infty \in W^{1,2}_{loc}(\mathbb{R}^{d_2}) \cap L^1_+(\mathbb{R}^{d_2}) \cap L^\infty_{loc}(\mathbb{R}^{d_2})$, $\|\mu_\infty\|_1 = 1$,
- (iv) $(\rho_\infty, \mu_\infty)$ satisfy (2.5a) for dynamics (2.2) or (2.5b) for dynamics (2.3).

The conditions (2.5) are given by

$$\nabla_z \delta_\rho F_a[\rho_\infty, \mu_\infty](z) = 0, \quad \nabla_x \delta_\mu F_a[\rho_\infty, \mu_\infty](x) = 0 \quad \forall z \in \text{supp}(\rho_\infty), x \in \text{supp}(\mu_\infty) \quad (2.5a)$$

$$\nabla_z \delta_\rho F_c[\rho_\infty, \mu_\infty](z) = 0, \quad \nabla_x \delta_\mu F_c[\rho_\infty, \mu_\infty](x) = 0 \quad \forall z \in \text{supp}(\rho_\infty), x \in \text{supp}(\mu_\infty) \quad (2.5b)$$

in the sense of distributions.

We define the Nash equilibrium of a game, and later show that the steady state of the dynamics in the zero-sum setting is in fact a Nash equilibrium.

Definition 2.7 (Nash Equilibrium). *A pair of measures $\gamma_* = (\rho_*, \mu_*) \in \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2})$ is a Nash equilibrium for the competitive setting if it satisfies*

$$F_c(\rho_*, \mu_*) \geq F_c(\rho, \mu_*) \quad \forall \rho \in \mathcal{P}(\mathbb{R}^{d_1}) \quad (2.6a)$$

$$F_c(\rho_*, \mu_*) \leq F_c(\rho_*, \mu) \quad \forall \mu \in \mathcal{P}(\mathbb{R}^{d_2}). \quad (2.6b)$$

2.3 Main Results

The convergence analysis of these systems allows us to understand and predict the long-time behavior of the dynamics. The asymptotics are given by the ground and saddle states of the energy functionals F_a and F_c respectively. We prove existence and uniqueness of the critical points of the functionals and, under sufficient convexity criteria, convergence with explicit rates.

Remark 2.8 (Cauchy-Problem). *To execute the arguments on convergence to equilibrium, we require sufficient regularity of solutions to the PDEs under consideration. In fact, it is sufficient if we can show that equations (2.2) - (2.3) can be approximated by equations with smooth solutions. Albeit tedious, these are standard techniques in the regularity theory for partial differential equations; see for example [50, Proposition 2.1 and Appendix A], [176], [204, Chapter 9], and the references therein. Similar arguments as in [85] are expected to apply to the coupled gradient flows considered here to guarantee existence of smooth solutions with fast enough decay at infinity. In this work, we do not focus on the existence and regularity of solutions.*

2.3.1 Assumptions

The key results on existence and uniqueness of a ground state or saddle point, as well as the convergence behavior of solutions, depend on convexity (concavity) of the corresponding functionals. The notion of convexity that we will employ for energy functionals in the Wasserstein-2 geometry is (uniform) displacement convexity, which is analogous to (strong) convexity in Euclidean spaces; see Definition 2.3.

We use subsets of the following assumptions in the cooperative and competitive cases.

Assumption 2.9. *The coupling potential f satisfies $f \in C^2(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}, \mathbb{R})$, and for all $(z, x) \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$,*

- (i) *Cooperative Setting: There exists $\lambda_f \in \mathbb{R}$ such that $\text{Hess}(f)(z, x) \succeq \lambda_f \mathbf{I}_{d_1 \times d_2}$. That is, f is λ_f -convex. Additionally, $f \geq 0$.*
- (ii) *Competitive Setting: There exists $\lambda_{f,1} \in \mathbb{R}$ such that $-\nabla_z^2 f(z, x) \succeq \lambda_{f,1} \mathbf{I}_{d_1}$ and $\lambda_{f,2} \in \mathbb{R}$ such that $\nabla_x^2 f(z, x) \succeq \lambda_{f,2} \mathbf{I}_{d_2}$. That is, f is $\lambda_{f,1}$ -concave in z and $\lambda_{f,2}$ -convex in x .*

Assumption 2.10. *The external potentials $V_i : \mathbb{R}^{d_i} \rightarrow [0, \infty)$ are in C^2 and satisfy lower Hessian bounds: there exists $\lambda_{V,i} \in \mathbb{R}$ such that $\nabla^2 V_i \succeq \lambda_{V,i} \mathbf{I}_{d_i}$.*

Assumption 2.11. *The interaction potentials $W_i : \mathbb{R}^{d_i} \rightarrow [0, \infty)$ are in C^2 , are symmetric and satisfy lower Hessian bounds: there exists $\lambda_{W,i} \geq 0$ such that $\nabla^2 W_i \succeq \lambda_{W,i} \mathbf{I}_{d_i}$.*

In the timescale-separated competitive setting, we use additional assumptions which provide upper bounds on the Hessian terms; see Section 2.4.

Remark 2.12. *The assumptions above and in Section 2.4 are not intended to be optimal; rather, they provide conditions analogous to conditions in the finite-dimensional setting under which convergence is guaranteed. For details on how some convexity assumptions can be weakened in combination with stronger assumptions on other terms, see [50]. The assumptions that the above potentials are in C^2 is strong and can likely be weakened without losing the main convergence guarantees. For the application settings considered in Sections 2.4 and 2.5 and in [62], all potentials are in C^2 ; however, more singular potentials are common in other settings such as interacting species in math-biology.*

2.3.2 The Cooperative Setting

The cooperative setting can be viewed as a class that includes potential games. From the PDE perspective, the system has a joint gradient flow structure which is utilized to prove convergence.

Theorem 2.13. *Suppose that Assumptions 2.9(i), 2.10, and 2.11 are satisfied with*

$$\lambda_a := \lambda_f + \min\{\lambda_{V,1}, \lambda_{V,2}\} > 0.$$

Consider solutions $\gamma_t := (\rho_t, \mu_t)$ to the dynamics (2.2) with initial condition satisfying $\gamma_0 \in \mathcal{P}_2(\mathbb{R}^{d_1}) \times \mathcal{P}_2(\mathbb{R}^{d_2})$, $F_a(\gamma_0) < \infty$, and

$$\int \|\nabla_z \delta_\rho F_a[\gamma_0](z)\|^2 d\rho_0(z) + \int \|\nabla_x \delta_\mu F_a[\gamma_0](x)\|^2 d\mu_0(x) < \infty.$$

Then the following hold:

(a) *There exists a unique minimizer $\gamma_\infty = (\rho_\infty, \mu_\infty)$ of F_a in $\mathcal{P} \times \mathcal{P}$, which is also a steady state for equation (2.2). Further, $\gamma_\infty \in \mathcal{P}_2 \times \mathcal{P}_2$.*

(i) *If $\alpha > 0$, then $\rho_\infty \in L_+^1(\mathbb{R}^{d_1}) \cap C^2(\mathbb{R}^{d_1})$ and $\text{supp}(\rho_\infty) = \mathbb{R}^{d_1}$.*

(ii) *If $\beta > 0$, then $\mu_\infty \in L_+^1(\mathbb{R}^{d_2}) \cap C^2(\mathbb{R}^{d_2})$ and $\text{supp}(\mu_\infty) = \mathbb{R}^{d_2}$.*

(b) *The solution γ_t converges exponentially fast in $F_a(\cdot | \gamma_\infty)$ and $\bar{W}$,*

$$F_a(\gamma_t | \gamma_\infty) \leq e^{-2\lambda_a t} F_a(\gamma_0 | \gamma_\infty) \quad \text{and} \quad \bar{W}(\gamma_t, \gamma_\infty) \leq c e^{-\lambda_a t} \quad \text{for all } t \geq 0,$$

where $c > 0$ is a constant only depending on γ_0 , γ_∞ and the parameter λ_a .

To prove existence and uniqueness, we leverage classical techniques in the calculus of variations. To obtain convergence to equilibrium in energy, our key result is an HWI-type inequality, providing as a consequence generalizations of the log-Sobolev inequality and the Talagrand inequality. Together,

these inequalities relate the energy (classically denoted by H in the case of the Boltzmann entropy), the metric (classically denoted by W in the case of the Wasserstein-2 metric) and the energy dissipation (classically denoted by I in the case of the Fisher information)¹. Combining these inequalities with Grönwall's inequality allows us to deduce convergence both in energy and in the metric $\overline{W}$. See Section 2.6 for a detailed proof.

2.3.3 The Competitive Setting

In the competitive setting, gradient descent by each player results in convergence to the unique Nash equilibrium.

Theorem 2.14. *Suppose Assumptions 2.9(ii), 2.10 and 2.11 are satisfied with*

$$\lambda_c := \min\{\lambda_{f,1} + \lambda_{V,1}, \lambda_{f,2} + \lambda_{V,2}\} > 0.$$

Consider solutions to (2.3) with initial condition $\gamma_0 \in \mathcal{P}_2(\mathbb{R}^{d_1}) \times \mathcal{P}_2(\mathbb{R}^{d_2})$ satisfying

$$\int \|\nabla_z \delta_\rho F_c[\gamma_0](z)\|^2 d\rho_0(z) + \int \|\nabla_x \delta_\mu F_c[\gamma_0](x)\|^2 d\mu_0(x) < \infty.$$

If $\alpha = 0$, assume $\rho_0 = \delta_{z_0}$ for some $z_0 \in \mathbb{R}^{d_1}$. If $\beta = 0$, assume $\mu_0 = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^{d_2}$. Then the following hold:

(a) *There exists a unique critical point $\gamma_* \in \tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$ for F_c which is also a steady state for equation (2.3) and the unique Nash equilibrium.*

(i) *If $\alpha > 0$, then $\rho_* \in L_+^1(\mathbb{R}^{d_1}) \cap C^2(\mathbb{R}^{d_1})$ and $\text{supp}(\rho_*) = \mathbb{R}^{d_1}$.*

(ii) *If $\beta > 0$, then $\mu_* \in L_+^1(\mathbb{R}^{d_2}) \cap C^2(\mathbb{R}^{d_2})$ and $\text{supp}(\mu_*) = \mathbb{R}^{d_2}$.*

(b) *The solution $\gamma_t := (\rho_t, \mu_t)$ to the dynamics (2.3) is in $\tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$, has uniformly bounded second moments,*

$$\exists K > 0 : \quad \int \|z\|^2 d\rho_t(z) + \int \|x\|^2 d\mu_t(x) \leq K \quad \forall t \geq 0,$$

and converges exponentially fast in $\overline{W}$ with rate λ_c ,

$$\overline{W}(\gamma_t, \gamma_*) \leq e^{-\lambda_c t} \overline{W}(\gamma_0, \gamma_*).$$

This theorem tells us that the distributions ρ_t and μ_t converge at a rate corresponding to the displacement concavity-convexity of the energy functional; we expect this to be true from similar

¹Hence the name HWI inequalities.

analysis of finite-dimensional zero-sum games. To prove the existence of a steady state of the dynamics (2.3), we first show that the second moments of ρ_t and μ_t are uniformly upper-bounded; then we show contraction in the joint metric, which allows us to conclude the existence of a steady state. Convergence follows from differentiating the Wasserstein-2 distance between the current pair of distributions and the steady state. Then we prove that the steady state is also a critical point and the unique Nash equilibrium using properties of the Euler-Lagrange conditions.

2.4 Application: Distribution Shift in Machine Learning

Machine learning algorithms in real-world settings often update their parameters over time to improve performance; as more data is collected, it is natural for the algorithm to update based on more recent information. However, in common applications the distribution of data on which the algorithm is training may not be stationary over time. This phenomenon is known as *distribution shift*, and is induced from a variety of causes, including mis-aligned incentives, interactions with other agents, and natural causes.

The model in [62] proposes an energy functional which the algorithm seeks to minimize and which the population aims to maximize; some terms have dependence on both distributions while others model energies specific to the evolution of agents and algorithms respectively. The model in [62] can be seen as a special case of the competitive setting considered here by choosing

$$\begin{aligned} f(z, x) &= f_1(z, x), \quad W_1(z) = W(z), \quad \alpha > 0, \quad \beta = 0, \\ V_1(z) &= -\alpha \log \tilde{\rho}(z), \quad V_2(x) = \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2, \end{aligned}$$

with all other terms set to zero and $d_1 = d_2 = d$. Here, $\tilde{\rho} \in \mathcal{P}(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ and $\pi \in \mathcal{P}(\mathbb{R}^d)$ are fixed reference measures. Let us denote

$$\begin{aligned} \mathcal{V}(\rho, \mu) &= \iint f_1(z, x) d\rho(z) d\mu(x) + \iint f_2(z, x) d\pi(z) d\mu(x), \\ \mathcal{R}(\rho) &= \frac{1}{2} \int W * \rho(z) \rho(z) + \alpha KL(\rho | \tilde{\rho}), \quad \mathcal{Q}(\mu) = \frac{\kappa}{2} \int \|x - x_0\|^2 d\mu(x). \end{aligned}$$

The functional $\mathcal{V}(\rho, \mu)$ models the cost which the algorithm seeks to minimize, and the population minimizes or maximizes depending on the setting. For example, when μ represents a binary classifier, the distribution π models individuals carrying the true label 1, and the distribution $\rho(t)$ model individuals carrying a true label 0, where 0 and 1 denote the labels of two classes of interest. The term $\int f_1(z, x) \mu(t, dx)$ represents a penalty for incorrectly classifying an individual with features z with true label 0 when using the classifier $\mu(t, x)$. Analogously, the term $\int f_2(z, x) \pi(dz)$ is large if x incorrectly classifies the population π that carries the true label 1. The functions f_1 and f_2 can be chosen according to the application; a standard choice for classification problems is the logarithmic loss [170].

The functional $Q(\mu)$ is a regularizer for the algorithm; this penalizes the classifier for selecting extreme learning parameters, and provides convexity for the loss function. The coefficient $\kappa > 0$ parameterizes the strength of the regularizer.

The functional $\mathcal{R}(\rho)$ contains two terms, the Kulback-Leibler divergence (denoted KL), also called the relative entropy, and the interaction term driven by the potential W . The term $\alpha KL(\rho|\tilde{\rho})$ forces the evolution of $\rho(t)$ to approach $\tilde{\rho}$. In other words, it penalizes (in energy) deviations from a given reference measure $\tilde{\rho}$. In many application settings, we take $\tilde{\rho}$ to be the initial distribution $\rho(t = 0)$. The solution $\rho(t)$ then evolves away from $\tilde{\rho}$ over time due to the other forces that are present. Therefore, the term $KL(\rho|\tilde{\rho})$ in the energy both provides smoothing of the flow and a penalization for deviations away from the reference measure $\tilde{\rho}$.

The self-interaction term $W * \rho$ introduces non-locality into the dynamics, as the decision for any given individual to move in a certain direction is influenced by the behavior of all other individuals in the population. The choice of W is application dependent. Very often, the interaction between two individuals only depends on the distance between them. This suggests a choice of W as a radial function, i.e. $W(z) = \omega(|z|)$. A choice of $\omega : [0, \infty) \rightarrow \mathbb{R}$ such that $\omega'(r) > 0$ corresponds to an *attractive* force between individuals, whereas $\omega'(r) < 0$ corresponds to a *repulsive* force.

However, in many real-world applications of this model, either the population or the algorithm may update much faster than the other. For example, many government policies are updated on a much slower timescale compared with how quickly individuals can adjust their response. In settings like advertising algorithms, companies can update their algorithms faster than users will adjust search strategies or viewing patterns. We provide convergence results for these two settings for competitive objectives, taking the timescale separations to be large enough that one entity instantly minimizes or maximizes its objective while the other entity evolves over time. Here we consider the case $\mu(t) = \delta_{x(t)}$ where $x(t)$ solves $\dot{x}(t) = -\nabla_x L(\rho, x)$ with loss function $L(\rho, x) = \int f_1(z, x) d\rho(z) + V_2(x)$. The energy in these two setting is denoted $G(\rho, x)$. The energy functional is given by

$$G(\rho, x) := \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2 - \mathcal{R}(\rho).$$

When the algorithm updates quickly relative to the population, we consider dynamics given by

$$\partial_t \rho = \nabla_{\mathcal{H}_2, \rho} G(\rho, x)|_{x=b(\rho)}, \quad b(\rho) := \underset{\bar{x} \in \mathbb{R}^d}{\operatorname{argmin}} G(\rho, \bar{x}). \quad (2.8)$$

In the opposite setting where the population is fast relative the algorithm, we can consider the population immediately responding to the algorithm, which results in the dynamics

$$\frac{d}{dt} x = -\nabla_x G(\rho, x)|_{\rho=r(x)}, \quad r(x) := \underset{\hat{\rho} \in \mathcal{P}}{\operatorname{argmax}} G(\hat{\rho}, x). \quad (2.9)$$

In this time-scale separated setting, model (2.8) is a dynamic maximization of G with respect to ρ in Wasserstein-2 space, and an instantaneous minimization of G with respect to the algorithm parameters x . Model (2.9) is an instantaneous maximization of G with respect to ρ and a dynamic minimization of G with respect to the algorithm parameters x .

2.4.1 Assumptions under Timescale Separation

We begin by rewriting Assumptions 2.9(ii), 2.10 and 2.11 in the timescale-separated setting (competitive objectives). From Assumption 2.9(ii), we have $f_1 \in C^2(\mathbb{R}^d \times \mathbb{R}^d)$ and (renaming parameters) assume

$$-\nabla_z^2 f_1(z, x) \succeq -\Lambda_1 \mathbf{I}_d, \quad \nabla_x^2 f_1(z, x) \succeq \lambda_1 \mathbf{I}_d$$

for some $\Lambda_1, \lambda_1 \in \mathbb{R}$. Assumption 2.10 is guaranteed by imposing

$$-\nabla_z^2 \log \tilde{\rho}(z) \succeq \tilde{\lambda} \mathbf{I}_d, \quad \nabla_x^2 f_2(z, x) \succeq \lambda_2 \mathbf{I}_d$$

for some $\lambda_2, \tilde{\lambda} \in \mathbb{R}$, together with $\log \tilde{\rho}(\cdot), f_2(z, \cdot) \in C^2(\mathbb{R}^d)$, $0 < \tilde{\rho} \leq 1$ and $0 \leq \int f_2(z, \cdot) d\pi(z) < \infty$. It follows that for any $R > 0$ there exists a constant $c_2 = c_2(R) \geq 0$ such that

$$\sup_{x \in B_R(0)} \int f_2(z, x) d\pi(z) < c_2. \quad (2.10)$$

For Assumption 2.11, we denote $\lambda_{W,1}$ by λ_W and assume $W \in C^2(\mathbb{R}^d)$ is symmetric with $W \geq 0$. In summary,

$$\lambda_{f,1} = -\Lambda_1, \quad \lambda_{f,2} = \lambda_1, \quad \lambda_{V,1} = \alpha \tilde{\lambda}, \quad \lambda_{V,2} = \lambda_2 + \kappa, \quad \lambda_{W,1} = \lambda_W, \quad \lambda_{W,2} = 0.$$

To analyze timescale-separation, we use the following additional assumptions.

Assumption 2.15. *The coupling potential f_1 satisfies for all $(z, x) \in \mathbb{R}^d \times \mathbb{R}^d$,*

(a) *Upper Hessian bound: There exists $\ell_1 \in \mathbb{R}$ such that $-\nabla_z^2 f_1(z, x) \preceq -\ell_1 \mathbf{I}_d$.*

(b) *Cross-terms: There exists $L \geq 0$ such that $\|\nabla_{xz}^2 f_1(z, x)\|_2 \leq L$.*

Note that in [62], $\lambda_1 = \ell_1$. In the results below, the value of ℓ_1 does not actually play a role and can be chosen independently of λ_1 .

Assumption 2.16. *Upper Hessian bound for external potential: There exists $\tilde{\Lambda} \in \mathbb{R}$ such that*

$$-\nabla_z^2 \log \tilde{\rho}(z) \preceq \tilde{\Lambda} \mathbf{I}_d.$$

Assumption 2.17. *Upper Hessian bound for the interaction potential: there exists $\Lambda_W \in \mathbb{R}$ such that*

$$\nabla_z^2 W \preceq \Lambda_W \mathbf{I}_d.$$

Assumption 2.18. The functions f_1, f_2 satisfy for all $(z, x) \in \mathbb{R}^d \times \mathbb{R}^d$: There exist constants $a_i > 0$ such that

$$x \cdot \nabla_x f_i(z, x) \geq -a_i \quad \text{for } i = 1, 2.$$

2.4.2 Analysis of Competitive Objectives with Timescale Separation

In the timescale separated cases, the dynamics are

$$\begin{aligned} \partial_t \rho &= -\operatorname{div}(\rho [\nabla(f_1(\cdot, b(\rho)) - \alpha \log(\rho/\tilde{\rho}) - W * \rho)]), \\ b(\rho) &:= \operatorname{argmin}_{\bar{x} \in \mathbb{R}^d} \int f_1(z, \bar{x}) d\rho(z) + \int f_2(z', \bar{x}) d\pi(z') + \frac{\kappa}{2} \|\bar{x} - x_0\|^2 \end{aligned}$$

for (2.8), and

$$\begin{aligned} \frac{d}{dt} x &= -\nabla_x \left(\int f_1(z, x) r(x) (dz) + \int f_2(z', x) d\pi(z') + \frac{\kappa}{2} \|x - x_0\|^2 \right), \\ r(x) &:= \operatorname{argmax}_{\hat{\rho} \in \mathcal{P}} \int f_1(z, x) d\hat{\rho}(z) - \alpha KL(\hat{\rho}|\tilde{\rho}) - \frac{1}{2} \int \hat{\rho} W * \hat{\rho}. \end{aligned}$$

for (2.9). Our results are summarized in the following theorems.

Theorem 2.19 (Fast Algorithm). Suppose Assumptions 2.9(ii), 2.10, 2.11, and 2.18 are satisfied with $\lambda_b := \alpha\tilde{\lambda} - \Lambda_1 > 0$ and $\lambda_d := \kappa + \lambda_1 + \lambda_2 > 0$. Define $G_b(\rho) := G(\rho, b(\rho))$. Consider a solution ρ_t to the dynamics (2.8) with initial condition $\rho_0 \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho_0) < \infty$. Then the following hold:

(a) There exists a unique maximizer ρ_∞ of $G_b(\rho)$, which is also a steady state for equation (2.8).

Moreover, $\rho_\infty \in L_+^1(\mathbb{R}^d) \cap C(\mathbb{R}^d)$ with the same support as $\tilde{\rho}$.

(b) The solution ρ_t converges exponentially fast to ρ_∞ in $G_b(\cdot | \rho_\infty)$ and $\mathcal{W}_2$,

$$G_b(\rho_t | \rho_\infty) \leq e^{-2\lambda_b t} G_b(\rho_0 | \rho_\infty) \quad \text{and} \quad \mathcal{W}_2(\rho_t, \rho_\infty) \leq c e^{-\lambda_b t} \quad \text{for all } t \geq 0,$$

where $c > 0$ is a constant only depending on ρ_0, ρ_∞ and the parameter λ_b .

Theorem 2.20 (Fast Population). Suppose Assumptions 2.9(ii), 2.10-2.17 are satisfied with $\lambda_b := \alpha\tilde{\lambda} - \Lambda_1 > 0$ and $\lambda_d := \kappa + \lambda_1 + \lambda_2 > 0$. Define $G_d(x) := G(r(x), x)$. Then it holds:

(a) There exists a unique minimizer x_∞ of $G_d(x)$ which is also a steady state for (2.9).

(b) The vector $x(t)$ solving the dynamics (2.9) with initial condition $x(0) \in \mathbb{R}^d$ converges exponentially fast to x_∞ in G_d and in the Euclidean norm:

$$\begin{aligned} \|x(t) - x_\infty\| &\leq e^{-\lambda_d t} \|x(0) - x_\infty\|, \\ G_d(x(t)) - G_d(x_\infty) &\leq e^{-2\lambda_d t} (G_d(x(0)) - G_d(x_\infty)) \end{aligned}$$

for all $t \geq 0$. Moreover, $G_d \in C^1(\mathbb{R}^d)$.

Theorems 2.19 and 2.20 state that we will observe exponential convergence in the timescale-separated cases. The proof of Theorem 2.19 hinges on proving a generalized version of Danskin’s Theorem, that is, showing that $\delta_\rho G_b(\rho) = (\delta_\rho G(\rho, x))|_{x=b(\rho)}$. Displacement convexity of G_b follows from this, and then standard HWI techniques apply. In order to prove a similar Danskin’s result as an ingredient in the proof of Theorem 2.20, we first show that $G_d(x)$ is differentiable. The best response function $r(x)$ is defined implicitly as the result of a minimization problem, and we employ Gamma convergence tools to obtain the regularity result. For the proofs, see Appendices 2.A and 2.B.

For the time-scale separated setting, we only require additional assumptions in the variable that optimizes instantaneously: Assumption 2.18 concerning the x -variable is used in Theorem 2.19, and Assumptions 2.15-2.17 concerning the z -variable appear in Theorem 2.20. In particular, Assumption 2.18 is used in Theorem 2.19 to prove that the norm of the best response $\|b(\rho)\|$ is uniformly bounded, while Assumptions 2.15-2.17 are used in Theorem 2.20 to prove convergence of the second moment of a sequence (ρ_n) in order to obtain the Γ -convergence result.

2.5 Insights from Numerical Experiments

In this section, we provide examples of numerical experiments that illustrate how our framework models real-world data and implications of using various algorithm learning strategies while interacting with a dynamic, strategic population. First, we model data from an economic study of how local government officials are incentivized to misreport census data. Then we show how individuals modify loan application data to achieve a more desirable outcome from a classifier algorithm, comparing different interaction models among agents and showing how subpopulations are affected differently. Finally, we illustrate our strategic population model under a state-of-the-art performative prediction algorithm, showing that it is critical to consider distribution shift beyond only mean shift when optimizing performance. The numerical experiments are implemented using the finite volume method from [47].

2.5.1 Census Data in Colombia

A study of Colombia census data [42] from 1995 to 2003 investigates how local officials misreported data in order to obtain lower poverty index scores for their constituents. Households with a poverty index score below a given threshold receive government aid, a desired outcome of the census data collection. The algorithm for the poverty index score and threshold was release in 1997, and the distribution of scores shifts from a Gaussian-like shape to having a sharp drop-off above the threshold.

We model this dynamical system as a classification problem in the competitive setting with a suitable energy G ; an algorithm with parameter x (government aid threshold) aims to separate poverty index scores into ones which qualify for aid and ones that do not. In this setting, we keep the algorithm




Figure 2.1: After the criteria for government aid was released in 1997, local officials misreported income data to increase the number of constituents qualifying for aid. The PDE (2.3) is able to capture the sharp drop at the classifier threshold. The convergence rate for the loss of the population and algorithm are 0.00995 and 0.0102; the convergence rate for $\mathcal{W}_2(\rho_t, \rho^{(98)})$, where $\rho^{(98)}$ is the steady state distribution, is also 0.0114 which is similar to the expected rate of 0.01. The expected rate is computed using convexity properties of the KL term.

fixed as the population adjusts, that is, $\frac{d}{dt}x = 0$. Each family aims to be classified as qualifying for aid, regardless of their true label. The distribution of poverty index scores for families whose true poverty index is qualifying them for aid is assumed stationary and given by π ; families whose true poverty index does not qualify them for aid is given by ρ , which evolves according to (2.3) and represents a strategic population. The threshold is given by $x(t) = 47$ for all time.

The initial condition for the strategic population is set to $\rho_0 = \mathcal{N}(54, 10)$. The stationary distribution of families that should qualify is given by the data distribution in the year 1995, denoted $\rho^{(95)}$, minus the strategic population, which we assume is half of the total population: $\pi = 2\rho^{(95)} - \mathcal{N}(54, 10)$. The utility functions are $f_1(z, x) = 1 - q(z, x) - lz$ and $f_2(z, x) = q(z, x)$, where $q(z, x) = (1 + \exp(-a(z - x)))^{-1}$. We use $\alpha = 0.1$, $a = 2$, $l = 0.06$, $\tilde{\rho} = \rho_0$ and $W = 0$. Here, $1 - q(z, x)$ is the probability that the classifier assigns a label of "qualified" to a family with attributes z and classifier parameters x . Families aim to maximize their probability of such a classification. The lz term models a preference for a lower poverty index score, regardless of the classifier parameters.

The functions appearing in the functional G are fit so that the global minimizer is $\rho^{(98)}$. In Figure 2.1 (left), the model distribution plotted is $(\rho + \pi)/2$, with $1e6$ samples drawn to generate the plot. We observe that the model is able to capture the sharp drop on the right side of the qualifying threshold due to the steep classifier f_1 , as well as the curvature of the distribution close to the threshold. In Figure 2.1, we also plot the loss of the classifier and population (top right), and the




Figure 2.2: While the accuracy of the classifier is similar under both interaction models, the precision differs; this indicates that understanding the intra-agent interactions is important for understanding how errors impact subpopulations, in this case, those agents with algorithm label “qualified.”

Wasserstein distance between the data from 1998, denoted $\rho^{(98)}$, and ρ_t (bottom right), and fit exponential functions to estimate the rate of decay. The expected rate of decay is specified by the rate given in Theorem 2.13 since the classifier is stationary. Because f_1 is such that $\lambda_{f,1} = -\Lambda_1 = 0$, the rate is generated by the convexity of $\tilde{\rho}$, which here we set to ρ_0 ; the expected rate is 0.01 since $\rho_0 \sim \mathcal{N}(54, 10)$. The convergence rate for the loss of the population and algorithm and for $\mathcal{W}_2(\rho_t, \rho^{(98)})$ are close to the theoretical value; see Figure 2.1 for details.

2.5.2 Loan Applications: Feature Modification

In settings such as loan applications, agents not eligible for a loan (label 0) aim to be misclassified to receive a more desirable outcome, such as qualifying for a loan (algorithm predicts label 1). In this numerical experiment, we consider real loan application data from [72] and allow label-0 agents to manipulate two out of their eleven features. We selected the two features that, pairwise, gave the lowest classification loss, which are “age” and “number of times the borrower has been 90 days or more past due.” The agents have a penalty for deviating from the initial condition, as enforced by the KL divergence term with $\tilde{\rho} = \rho_0$, and the potential function is the negative of the probability of being classified as a label-1 agent. While the agents can only manipulate two features, we allow the classifier to update based on all eleven features. The agents with true label 1 do not manipulate any features. For this application, we consider interactions between agents assuming that people exchange information about their loan applications and application outcomes. We compare the evolution of the population for $N = 1000$ agents under two different interaction kernels:

$$\begin{aligned} W_r(z) &= -|z| && \text{repulsive,} \\ W_a(z) &= 4e^{-5|z|} - 2e^{-|z|/4} && \text{attractive-repulsive,} \end{aligned}$$

implemented via the underlying particle system given by

$$dz_t^{(i)} = -\nabla_z f_1(z_t^{(i)}, x)dt - \sum_{j \neq i} \nabla_z W(|z_t^{(i)} - z_t^{(j)}|)dt + \alpha \nabla \log \rho_0(z_t^{(i)}) + \sqrt{2\alpha} dB_t, \quad z_t^{(i)} \in \mathbb{R}^2,$$




Figure 2.3: Agent densities are split based on their true label (eligible vs ineligible). The color denotes the algorithm output regarding loan qualification (qualified vs unqualified). The repulsive kernel causes the agents to spread apart along attributes 1 (normalized age) and 2 (normalized past due), while the attractive-repulsive kernel causes more swarm-like behavior. This is not evident from classifier performance only, indicating the importance of understanding the population dynamics explicitly.

with $W = W_a$ and $W = W_r$ for each setting. The repulsive interaction kernel encourages agents to move away from their neighbors, while the attractive-repulsive kernel encourages swarm-like behavior by pushing neighboring agents apart while attracting agents that are far from each other. We update the agent data and linear classifier using code adapted from [168], which uses a particle-based gradient descent scheme in which the agents take 500 update steps per 1 update step for the algorithm. In Figure 2.2, we observe that under the same training scheme, the accuracy for the classifier is about the same for both interaction kernels. However, the precision, which is the number of true positives over the number of all positives, is not the same. We compute the accuracy for a subpopulation of the label-0 population: those agents initially classified as 0 and those initially (mis)classified as 1. Under the repulsive kernel, these agents are classified with 96% and 20% accuracy at the steady state; in contrast, the agents are classified with 90% and 9% accuracy under the attractive-repulsive kernel. Although we do not have convergence guarantees in this setting, this example highlights the importance of modeling intra-species interactions for applications in which non-asymptotic behavior is relevant.

This indicates that the initially-mislabeled subpopulation benefits from attractive-repulsive interactions more than repulsive interactions. In Figure 2.3, we plot the density estimate and observe that the kernels induce different feature distributions, even though the accuracy is similar, indicating that observing only the performance of the classifier fails to indicate important details about specific subpopulations.

2.5.3 Performative Prediction

A current state-of-the-art model for updating an algorithm in the face of distribution shift is given in [170]; the algorithm perturbs the x parameter randomly and records the mean of the population distribution after it responds to the perturbation. Then a linear model is fit between the perturbation

direction and the distribution mean, and the algorithm is set to the minimizer of the linear model. We simulate the strategic distribution using our PDE model in the competitive setting, comparing




Figure 2.4: A gradient descent approach outperforms a state-of-the-art technique of learning and using a linear mapping between the classifier parameters and the mean of the strategic distribution, illustrating the importance of having more detailed population models.

classifier performance under two update strategies: (1) the perturbation method described above, and (2) naïve gradient descent, specifically Wasserstein gradient descent in the dynamics (2.3). The resulting distribution shifts and algorithm updates are shown in Figure 2.4. In Figure 2.4a, we observe that the optimization problem from the mean shift model [170] sets the classifier to $x(t) \approx 0$ for all time, while in Figure 2.4b, the classifier moves to the right, resulting in better performance. This difference occurs because the perturbation method detects only a mean shift, not that two modes are appearing, and underestimates the impact of the population mass that splits from the main population. This illustrates that modeling the population with more fine-grained detail than just mean information is critical for selecting algorithm update strategies that are most effective. In Figure 2.4c, the convergence of the loss for the population ρ and classifier x are shown, where we subtracted the value at steady state from the energy functional so that the energy converges to zero. The rates are given by 0.192 and 0.192 for the population and classifier (resp.). The theoretical rate from Theorem 2.14 is $\lambda_c = \min\{0.05, 1 - 1.483\} = -0.483$, which indicates that the theorem conditions are not satisfied, since $\lambda_c < 0$. However, the condition is not tight (for example, when the mass is not concentrated around the worst-case convexity given by Λ_1), and we observe convergence empirically. The loss for ρ is increasing because ρ is maximizing G , and although x is minimizing G , the evolution of ρ causes the loss of x to increase slightly.

2.6 Cooperative Setting (Proof of Theorem 2.13)

Denote

$$\begin{aligned}\mathcal{V}(\rho, \mu) &= \iint f(z, x) d\rho(z) d\mu(x) + \int V_1(z) d\rho(z) + \int V_2(x) d\mu(x), \\ \mathcal{W}(\rho, \mu) &= \frac{1}{2} \int (W_1 * \rho)(z) d\rho(z) + \frac{1}{2} \int (W_2 * \mu)(x) d\mu(x),\end{aligned}$$

so that the functional F_a is given by

$$F_a(\rho, \mu) = \mathcal{V}(\rho, \mu) + \alpha H(\rho) + \beta H(\mu) + \mathcal{W}(\rho, \mu).$$

In order to prove the existence of a unique ground state for F_a , a natural approach is to consider the corresponding Euler-Lagrange equations

$$\alpha \log \rho(z) + \int f(z, x) d\mu(x) + V_1(z) + (W_1 * \rho)(z) = c_1[\rho, \mu] \quad \text{for all } z \in \text{supp}(\rho), \quad (2.11a)$$

$$\beta \log \mu(x) + \int f(z, x) d\rho(z) + V_2(x) + (W_2 * \mu)(x) = c_2[\rho, \mu] \quad \text{for all } x \in \text{supp}(\mu), \quad (2.11b)$$

where c_1, c_2 are constants that may differ on different connected components of $\text{supp}(\rho)$ and $\text{supp}(\mu)$. These equations are not easy to solve explicitly, and we are therefore using general non-constructive techniques from calculus of variations. We first show continuity and λ_a -convexity properties for the functional F_a (Lemma 2.21 and Proposition 2.22), where

$$\lambda_a := \lambda_f + \min\{\lambda_{V_1}, \lambda_{V_2}\};$$

essential properties that will allow us to deduce existence and uniqueness of ground states using the direct method in the calculus of variations (Proposition 2.25). Using the Euler-Lagrange equation (2.11), we then prove properties on the support of the ground state (Corollary 2.26). To obtain convergence results, we apply the HWI method: we first show a general 'interpolation' inequality among the energy, the energy dissipation and the metric (Proposition 2.27); this fundamental inequality will then imply a generalized logarithmic Sobolev inequality (Corollary 2.29) relating the energy to the energy dissipation, and a generalized Talagrand inequality (Corollary 2.30) that translates convergence in energy into convergence in metric. Putting all these ingredients together will then allow us to conclude the statements in Theorem 2.13. Throughout this section, we assume $\lambda_a > 0$ and that Assumptions 2.9(i), 2.10 and 2.11 hold.

2.6.1 Ground States and Steady States

Lemma 2.21 (Lower semi-continuity). *The functional $F_a : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ is lower semi-continuous with respect to the weak topology.*

Proof. We split the energy F_a into three parts: (i) $\alpha\mathcal{H}(\rho) + \beta\mathcal{H}(\mu)$, (ii) $\mathcal{W}(\rho, \mu)$, and (iii) the joint potential energy $\mathcal{V}(\rho, \mu)$. For (i), Lemma 2.64 gives lower-semicontinuity for $\rho \in \mathcal{P}$ with respect to the weak topology. [194, Proposition 7.2] provides lower-semicontinuity for (ii) because W_i is continuous. For (iii), note that f is lower semi-continuous and bounded below thanks to Assumptions 2.9(i) and 2.10, and so the result follows from [194, Proposition 7.1]. $\square$

Proposition 2.22 (Uniform displacement convexity). *Fix $\gamma_0, \gamma_1 \in \mathcal{P}_2 \times \mathcal{P}_2$. Along any geodesic $(\gamma_s)_{s \in [0,1]} \in \mathcal{P}_2 \times \mathcal{P}_2$ connecting γ_0 to γ_1 , we have for all $s \in [0, 1]$*

$$\frac{d^2}{ds^2} F_a(\gamma_s) \geq \lambda_a \overline{W}(\gamma_0, \gamma_1)^2. \quad (2.12)$$

As a result, the functional $F_a : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ is uniformly displacement convex with constant $\lambda_a > 0$.

Proof. Let γ_0 and γ_1 be two absolutely continuous probability measures with bounded second moments. The general case can be recovered using approximation arguments. Denote by $\phi, \psi : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ the optimal Kantorovich potentials pushing ρ_0 onto ρ_1 , and μ_0 onto μ_1 , respectively:

$$\begin{aligned} \rho_1 &= \nabla \phi_{\#} \rho_0 \quad \text{such that} \quad \mathcal{W}_2(\rho_0, \rho_1)^2 = \int_{\mathbb{R}^{d_1}} \|z - \nabla \phi(z)\|^2 d\rho_0(z), \\ \mu_1 &= \nabla \psi_{\#} \mu_0 \quad \text{such that} \quad \mathcal{W}_2(\mu_0, \mu_1)^2 = \int_{\mathbb{R}^{d_2}} \|x - \nabla \psi(x)\|^2 d\mu_0(x). \end{aligned}$$

The now-classical results in [20] guarantee that there exist convex functions ϕ, ψ that satisfy the conditions above. Then the path $(\gamma_s)_{s \in [0,1]} = (\rho_s, \mu_s)_{s \in [0,1]}$ defined by

$$\rho_s = [(1-s) \text{id} + s \nabla \phi]_{\#} \rho_0, \quad \mu_s = [(1-s) \text{id} + s \nabla \psi]_{\#} \mu_0$$

is a $\overline{W}$ -geodesic from γ_0 to γ_1 .

The first derivative of $\mathcal{V}$ along geodesics in the Wasserstein metric is given by

$$\begin{aligned} \frac{d}{ds} \mathcal{V}(\gamma_s) &= \frac{d}{ds} \left[\iint f((1-s)z + s \nabla \phi(z), (1-s)x + s \nabla \psi(x)) d\rho_0(z) d\mu_0(x) \right. \\ &\quad \left. + \int V_1((1-s)z + s \nabla \phi(z)) d\rho_0(z) + \int V_2((1-s)x + s \nabla \psi(x)) d\mu_0(x) \right] \\ &= \iint \nabla_x f((1-s)z + s \nabla \phi(z), (1-s)x + s \nabla \psi(x)) \cdot (\nabla \psi(x) - x) d\rho_0(z) d\mu_0(x) \\ &\quad + \iint \nabla_z f((1-s)z + s \nabla \phi(z), (1-s)x + s \nabla \psi(x)) \cdot (\nabla \phi(z) - z) d\rho_0(z) d\mu_0(x) \\ &\quad + \int \nabla_z V_1((1-s)z + s \nabla \phi(z)) \cdot (\nabla \phi(z) - z) d\rho_0(z) \\ &\quad + \int \nabla_x V_2((1-s)x + s \nabla \psi(x)) \cdot (\nabla \psi(x) - x) d\mu_0(x), \end{aligned}$$

and taking another derivative we have

$$\begin{aligned}
\frac{d^2}{ds^2} \mathcal{V}(\gamma_s) &= \iint \left[\frac{(\nabla \psi(x) - x)}{(\nabla \phi(z) - z)} \right]^\top \cdot D_s(z, x) \cdot \left[\frac{(\nabla \psi(x) - x)}{(\nabla \phi(z) - z)} \right] d\rho_0(z) d\mu_0(x) \\
&\quad + \int (\nabla \varphi(z) - z)^\top \cdot \nabla_z^2 V_1((1-s)z + s\nabla \psi(z)) \cdot (\nabla \varphi(z) - z) d\rho_0(x) \\
&\quad + \int (\nabla \psi(x) - x)^\top \cdot \nabla_x^2 V_2((1-s)x + s\nabla \psi(x)) \cdot (\nabla \psi(x) - x) d\mu_0(x) \\
&\geq \lambda_f \bar{W}(\gamma_0, \gamma_1)^2 + \lambda_{f,1} \mathcal{W}_2(\rho_0, \rho_1)^2 + \lambda_{f,2} \mathcal{W}_2(\mu_0, \mu_1)^2,
\end{aligned}$$

where we denoted $D_s(z, x) := \text{Hess}(f)((1-s)z + s\nabla \phi(z), (1-s)x + s\nabla \psi(x))$, and the last inequality follows from Assumptions 2.9(i) and 2.10 and the optimality of the potentials ϕ and ψ .

Following [50, 204] and using Assumption 2.11, the second derivatives of the diffusion and self-interaction terms are given by

$$\frac{d^2}{ds^2} \mathcal{H}(\rho_s) \geq 0, \quad \frac{d^2}{ds^2} \mathcal{H}(\mu_s) \geq 0, \quad \frac{d^2}{ds^2} \mathcal{W}(\gamma_s) \geq 0. \quad (2.13)$$

Putting the above estimates together, we obtain (2.12). $\square$

Remark 2.23. *If the dynamics are such that the center of mass of ρ_t or μ_t are preserved for all time, then the convexity of W_i contributes to the rate of convergence. This occurs, for example, when V_1 and W_1 are radially symmetric and the initial condition ρ_0 is radially symmetric; then the rate λ_a would be $\lambda_a = \lambda_f + \min\{\lambda_{V,1} + \lambda_{W,1}, \lambda_{V,2}\}$. For details, see [50, Theorems 2.2, 2.4, 2.5].*

Lemma 2.24 (Lower bound). *We have $F_a \geq 0$ over $\mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2})$.*

Proof. By Assumption 2.9(i), 2.10 and 2.11, $f, V_i, W_i \geq 0$. When $\alpha = \beta = 0$, then $F_a \geq 0$. If $\alpha > 0$ or $\beta > 0$, we will write the log term as a KL divergence to show that F_a is non-negative. Since $\lambda_a > 0$, either $f \neq 0$ or $V_1, V_2 \neq 0$. We will show the setting in which $f \neq 0$ and $V_1, V_2 = 0$ as all other cases follow analogously. If either ρ (with $\alpha > 0$) or μ (with $\beta > 0$) is a singular measure, then $F_a = +\infty$ according to the definition of the entropy functional, and so the claim holds trivially true. In all other cases, the functional F_a can be rewritten as

$$\begin{aligned}
F_a[\rho, \mu] &= \frac{1}{2} \int \rho W_1 * \rho + \frac{1}{2} \int \mu W_2 * \mu \\
&\quad + \alpha \int \left(\int \rho(z) \log \frac{\rho(z)}{\tilde{f}_1(z, x)} dz \right) d\mu(x) + \beta \int \left(\int \mu(x) \log \frac{\mu(x)}{\tilde{f}_2(z, x)} dx \right) d\rho(z),
\end{aligned}$$

where $\tilde{f}_1(z, x) := \exp\left(-\frac{1}{2\alpha} f(z, x)\right)$ and $\tilde{f}_2(z, x) := \exp\left(-\frac{1}{2\beta} f(z, x)\right)$. Since ρ, μ are absolutely continuous with respect to $\tilde{f}_1, \tilde{f}_2$ respectively, we have $F_a \geq 0$ by Pinsker's inequality and using $W_i \geq 0$. $\square$

Proposition 2.25. (Ground state) The functional $F_a : \mathcal{P}(\mathbb{R}^{d_1}) \times \mathcal{P}(\mathbb{R}^{d_2}) \rightarrow [0, \infty]$ admits a unique minimizer $\gamma_* = (\rho_*, \mu_*)$ which satisfies $\gamma_* \in \mathcal{P}_2 \times \mathcal{P}_2$. Moreover, if $\alpha > 0$ it satisfies $\rho_* \in L_+^1(\mathbb{R}^{d_1})$, and if $\beta > 0$ it satisfies $\mu_* \in L_+^1(\mathbb{R}^{d_2})$.

Proof. We show existence of a minimizer of F_a using the direct method in the calculus of variations. Denote by $\gamma = (\rho, \mu) \in \mathcal{P} \times \mathcal{P} \subset \mathcal{M} \times \mathcal{M}$ a pair of probability measures as a point in the product space of Radon measures. Since $F_a \geq 0$ on $\mathcal{P} \times \mathcal{P}$ by Lemma 2.24 and not identically $+\infty$ everywhere, there exists a minimizing sequence $(\gamma_n) \in \mathcal{P} \times \mathcal{P}$. Note that (γ_n) is in the closed unit ball of the dual space of continuous functions vanishing at infinity $(C_0(\mathbb{R}^{d_1}) \times C_0(\mathbb{R}^{d_2}))^*$ endowed with the dual norm $\|\gamma_n\|_* = \sup \frac{|\int f d\rho_n + \int g d\mu_n|}{\|(f, g)\|_\infty}$ over $f, g \in C_0(\mathbb{R}^{d_i})$ with $\|(f, g)\|_\infty := \|f\|_\infty + \|g\|_\infty \neq 0$. By the Banach-Alaoglu theorem [192, Thm 3.15] there exists a limit $\gamma_* = (\rho_*, \mu_*) \in \mathcal{M} \times \mathcal{M} = (C_0 \times C_0)^*$ and a convergent subsequence (not relabelled) such that $\gamma_n \xrightarrow{*} \gamma_*$.

It remains to show that $\int d\rho_* = \int d\mu_* = 1$ to conclude that $\gamma_* \in \mathcal{P} \times \mathcal{P}$. To this aim, it is sufficient to show tightness of (ρ_n) and (μ_n) , preventing the escape of mass to infinity as we have $\int d\rho_n = \int d\mu_n = 1$ for all $n \geq 1$. Tightness follows from Markov's inequality [117] if we can establish uniform bounds on the second moments, i.e. we want to show that there exists a constant $C > 0$ independent of n such that

$$\int \|z\|^2 d\rho_n(z) + \int \|x\|^2 d\mu_n(x) < C \quad \forall n \in \mathbb{N}. \quad (2.14)$$

To establish (2.14), observe that thanks to Assumption 2.9, there exist a constant $c_0 \in \mathbb{R}$ and vector $c_1 \in \mathbb{R}^{d_1+d_2}$ such that $f(z, x) \geq c_0 \left\| \begin{bmatrix} z \\ x \end{bmatrix} + c_1 \right\|^2$ for all $[z, x] \in \mathbb{R}^{d_1+d_2}$. Additionally, from Lemma 2.63 applied with $\varepsilon = c_0/(2 \max\{\alpha, \beta\}(\|c_1\| + 1))$, we have that $\mathcal{H}(\rho) \geq -\varepsilon \int \|z\|^2 d\rho(z) - c_\varepsilon$ for some $c_\varepsilon \geq 0$, with the analogous bound for $\mathcal{H}(\mu)$. Then

$$\begin{aligned} c_0 \iint \left\| \begin{bmatrix} z \\ x \end{bmatrix} + c_1 \right\|^2 d\gamma_n(z, x) &\leq \iint f(z, x) d\gamma_n(z, x) \\ &\leq F_a(\gamma_n) + \varepsilon \left(\alpha \int \|z\|^2 d\rho_n(z) + \beta \int \|x\|^2 d\mu_n(x) \right) + (\alpha + \beta)c_\varepsilon. \end{aligned}$$

Since $\frac{\| [z, x] \|_2^2}{\|c_1\| + 1} - \|c_1\| \leq \left\| \begin{bmatrix} z \\ x \end{bmatrix} + c_1 \right\|^2$, the estimate can be rearranged to

$$\frac{c_0}{2(\|c_1\| + 1)} \left(\int \|z\|^2 d\rho_n(z) + \int \|x\|^2 d\mu_n(x) \right) \leq F_a(\gamma_n) + \hat{c}_\varepsilon \leq F_a(\gamma_1) + \hat{c}_\varepsilon < \infty,$$

where $\hat{c}_\varepsilon := c_0 \|c_1\| + (\alpha + \beta)c_\varepsilon$. Hence, the second moment is uniformly bounded. This concludes the proof that the limit γ_* satisfies $\gamma_* \in \mathcal{P} \times \mathcal{P}$, and indeed $\rho_* \in \mathcal{P}_2(\mathbb{R}^{d_1})$, $\mu_* \in \mathcal{P}_2(\mathbb{R}^{d_2})$ as well. Further, note that the above second moment bound implies that (γ_n) also converges weakly according to Definition 2.1. Finally, γ_* is a minimizer of F_a thanks to weak lower-semicontinuity of F_a following Lemma 2.21.

If $\alpha > 0$, the ground state ρ_* satisfies

$$\rho_*(z) = c_3 \exp \left(-\frac{1}{\alpha} \left(\int f(z, x) d\mu_*(x) + (W_1 * \rho_*)(z) + V_1(z) \right) \right) \quad \text{on } \text{supp } \rho_*.$$

and so $\rho_* \in L_+^1(\mathbb{R}^{d_1})$. Similarly, if $\beta > 0$, the ground state μ_* satisfies

$$\mu_*(x) = c_4 \exp \left(-\frac{1}{\beta} \left(\int f(z, x) d\rho_*(z) + (W_2 * \mu_*)(x) + V_2(x) \right) \right) \quad \text{on } \text{supp } \mu_*,$$

and we have that $\mu_* \in L_+^1(\mathbb{R}^{d_2})$.

Next we show uniqueness using a contradiction argument. Suppose $\gamma_* = (\rho_*, \mu_*)$ and $\gamma'_* = (\rho'_*, \mu'_*)$ are minimizers of F_a . For $s \in [0, 1]$, define $\gamma_s := ((1-s)\text{id} + sT, (1-s)\text{id} + sS)_\# \gamma_*$, where $T, S : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^{d_i}$ are the optimal transport maps such that $\rho'_* = T_\# \rho_*$ and $\mu'_* = S_\# \mu_*$. By Proposition 2.22 the energy F_a is uniformly displacement convex, and so we have

$$F_a(\gamma_s) \leq (1-s)F_a(\gamma_*) + sF_a(\gamma'_*) = F_a(\gamma_*).$$

If $\gamma_* \neq \gamma'_*$ and $s \in (0, 1)$, then strict inequality holds by applying similar arguments as in [162, Proposition 1.2]. However, the strict inequality $F_a(\gamma_s) < F_a(\gamma_*)$ for $\gamma_* \neq \gamma'_*$ is a contradiction to the minimality of γ_* . Hence, the minimizer is unique. $\square$

Corollary 2.26. *Any minimizer $\gamma_* = (\rho_*, \mu_*)$ of F_a is a steady state for equation (2.2) according to Definition 2.5a. If $\alpha > 0$, then $\text{supp}(\rho_*) = \mathbb{R}^{d_1}$ and $\rho_* \in C^2(\mathbb{R}^{d_1})$. If $\beta > 0$, then $\text{supp}(\mu_*) = \mathbb{R}^{d_2}$ and $\mu_* \in C^2(\mathbb{R}^{d_2})$.*

Proof. By Proposition 2.25, we have $\rho_*, \mu_* \in \mathcal{P}_2$. As γ_* is a minimizer, it is in particular a critical point, and therefore satisfies equations (2.11). In order to show that γ_* is a steady state for equation (2.2), we need to show that $\nabla W_1 * \rho_* \in L_{loc}^1$ and $\nabla W_2 * \mu_* \in L_{loc}^1$, and that if $\alpha > 0$, then $\rho_* \in W_{loc}^{1,2} \cap L_+^1 \cap L_{loc}^\infty$ and if $\beta > 0$, then $\mu_* \in W_{loc}^{1,2} \cap L_+^1 \cap L_{loc}^\infty$.

We claim that $\int f(z, \tilde{x}) d\mu_*(\tilde{x}) < \infty$, $\int f(\tilde{z}, x) d\rho_*(\tilde{z}) < \infty$, $W_1 * \rho_*(z) < \infty$ and $W_2 * \mu_*(x) < \infty$ for any fixed $(z, x) \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$. Indeed, each term in the energy without the interaction and coupling potentials can be lower-bounded by a constant separately using positivity of the potentials and Lemma 2.63 together with the second moment bound for γ_* . Hence, $\iint f(z, x) d\mu_*(x) d\rho_*(z) < \infty$, $\int (W_1 * \rho_*)(z) d\rho_*(z) < \infty$ and $\int (W_2 * \mu_*)(x) d\mu_*(x) < \infty$ as γ_* is a minimizer. This concludes the bounds, and so we also obtain $\nabla W_1 * \rho_* \in L_{loc}^1$ and $\nabla W_2 * \mu_* \in L_{loc}^1$ since $W_1, W_2 \in C^2$.

If $\alpha = \beta = 0$, we can differentiate (2.11) directly to obtain (2.5). Now, consider the case $\alpha > 0, \beta > 0$. Rearranging (2.11), we obtain (for possibly different constants $c_1[\rho_*, \mu_*], c_2[\rho_*, \mu_*] \neq 0$) that

$$\begin{aligned} \rho_*(z) &= c_1[\rho_*, \mu_*] \exp \left[-\frac{1}{\alpha} \left(\int f(z, x) d\mu_*(x) + W_1 * \rho_*(z) + V_1(z) \right) \right] \quad \text{on } \text{supp}(\rho_*), \\ \mu_*(x) &= c_2[\rho_*, \mu_*] \exp \left[-\frac{1}{\beta} \left(\int f(z, x) d\rho_*(z) + W_2 * \mu_*(x) + V_2(x) \right) \right] \quad \text{on } \text{supp}(\mu_*), \end{aligned} \tag{2.15}$$

and so $\rho_*, \mu_* \in L_+^1$. Then for any compact set $K \subset \mathbb{R}^{d_1}$,

$$\sup_{z \in K} \rho_*(z) \leq c_1 [\rho_*, \mu_*] \sup_{z \in K} \exp \left(-\frac{1}{\alpha} \left(\int f(z, x) d\mu_*(x) + V_1(z) \right) \right) \sup_{z \in K} \exp \left(-\frac{1}{\alpha} W_1 * \rho_* \right).$$

As $f, V_1, W_1 \geq 0$, the exponential terms on the right-hand side are finite. Therefore $\rho_* \in L_{loc}^\infty$. To show that $\rho_* \in W_{loc}^{1,2}$, note that for any compact set $K \subset \mathbb{R}^{d_1}$, we have $\int_K |\rho_*(z)|^2 dz < \infty$ as a consequence of $\rho_* \in L_{loc}^\infty$. Moreover, defining $T[\gamma](z) := -\frac{1}{\alpha} \left(\int f(z, x) d\mu(x) + W_1 * \rho(z) + V_1(z) \right) \leq 0$, we have

$$\int_K |\nabla \rho_*|^2 dz = c_1 [\rho_*, \mu_*]^2 \int_K |\nabla T[\gamma_*]|^2 \exp(2T[\gamma_*]) dz,$$

which is bounded noting that $\exp(2T[\gamma_*]) \leq 1$ and that $\nabla T[\gamma_*](\cdot)$ is in L_{loc}^∞ , since $f(\cdot, x), W_1(\cdot), V_1(\cdot) \in C^1(\mathbb{R}^{d_1})$ by Assumptions 2.9-2.11. We conclude that $\rho_* \in W_{loc}^{1,2}$, and use an identical argument for the case when $\beta > 0$, and indeed (ρ_*, μ_*) solves (2.5a) in the sense of distributions as a consequence of (2.11).

Next, we show that if $\alpha > 0$ then $\text{supp}(\rho_*) = \mathbb{R}^{d_1}$ using again the relation (2.15). Indeed, $\exp \left[-\frac{1}{\alpha} \left(\int f(z, x) d\mu_*(x) + W_1 * \rho_*(z) + V_1(z) \right) \right] > 0$ for all $z \in \mathbb{R}^{d_1}$ since $\int f(z, x) d\mu_*(x) < \infty$, $W_1 * \rho_*(z) < \infty$ and $V_1(z) < \infty$. Then, thanks to continuity of $f \in C^2$, $V_1 \in C^2$, and $W_1 * \rho_*(z) \in C^2$, we conclude $\rho_* \in C^2(\mathbb{R}^{d_1})$. The same argument is applied to μ_* when $\beta > 0$ to obtain $\text{supp} \mu_* = \mathbb{R}^{d_2}$ and $\mu_* \in C^2(\mathbb{R}^{d_2})$. $\square$

2.6.2 Functional Inequalities

The following inequality is referred to as HWI inequality and represents the key result to obtain convergence to equilibrium.

Proposition 2.27 (HWI inequality). *Define the dissipation functional*

$$D_a(\gamma) := \iint \left\| \begin{bmatrix} \nabla_z \delta_\rho F_a[\rho, \mu](z) \\ \nabla_x \delta_\mu F_a[\rho, \mu](x) \end{bmatrix} \right\|^2 d\gamma(z, x).$$

Let $\gamma_0, \gamma_1 \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $F_a(\gamma_0), F_a(\gamma_1), D_a(\gamma_0) < \infty$, and let Assumptions 2.9(i), 2.10 and 2.11 hold with $\lambda_a > 0$. Then

$$F_a(\gamma_0) - F_a(\gamma_1) \leq \overline{W}(\gamma_0, \gamma_1) \sqrt{D_a(\gamma_0)} - \frac{\lambda_a}{2} \overline{W}(\gamma_0, \gamma_1)^2. \quad (2.16)$$

Proof. For simplicity, consider γ_0, γ_1 that have smooth Lebesgue densities of compact support. The general case can be recovered using approximation arguments. Let $(\gamma_s)_{s \in [0,1]}$ denote a $\overline{W}$ -geodesic between γ_0, γ_1 . Following similar arguments as in [50] and [176, Section 5] and making use of the

calculations in the proof of Proposition 2.22, we have

$$\left. \frac{d}{ds} F_a(\gamma_s) \right|_{s=0} = \iint \begin{bmatrix} \xi_1(z) \\ \xi_2(x) \end{bmatrix} \cdot \begin{bmatrix} (\nabla \phi(z) - z) \\ (\nabla \psi(x) - x) \end{bmatrix} d\gamma_0(z, x),$$

where

$$\begin{aligned} \xi_1[\gamma_0](z) &:= \int \nabla_z f(z, x) d\mu_0(x) + \alpha \nabla_z \log \rho_0(z) + \nabla W_1 * \rho_0(z) + \nabla_z V_1(z) = \nabla_z \delta_\rho F_a[\gamma_0](z), \\ \xi_2[\gamma_0](x) &:= \int \nabla_x f(z, x) d\rho_0(z) + \beta \nabla_x \log \mu_0(x) + \nabla W_2 * \mu_0(x) + \nabla_x V_2(x) = \nabla_x \delta_\mu F_a[\gamma_0](x). \end{aligned}$$

Note that the dissipation functional can then be written as

$$D_a(\gamma_0) = \iint \left(\|\xi_1(z)\|^2 + \|\xi_2(x)\|^2 \right) d\gamma_0(z, x).$$

Using the double integral Cauchy-Schwarz inequality [200], we obtain

$$\begin{aligned} \left. \frac{d}{ds} F_a(\gamma_s) \right|_{s=0} &\geq - \left(\sqrt{\iint \left\| \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} \right\|^2 d\gamma_0} \right) \left(\sqrt{\iint \left\| \begin{bmatrix} \nabla \phi(z) - z \\ \nabla \psi(x) - x \end{bmatrix} \right\|^2 d\gamma_0} \right) \\ &= -\sqrt{D_a(\gamma_0)} \sqrt{\int \|\nabla \phi(z) - z\|^2 d\rho_0 + \int \|\nabla \psi(x) - x\|^2 d\mu_0} \\ &= -\sqrt{D_a(\gamma_0)} \overline{W}(\gamma_0, \gamma_1). \end{aligned}$$

Next, we compute a Taylor expansion of $F_a(\gamma_s)$ when considered as a function in s and use the bound on $\frac{d^2}{ds^2} F_a$ from (2.12):

$$\begin{aligned} F_a(\gamma_1) &= F_a(\gamma_0) + \left. \frac{d}{ds} F_a(\gamma_s) \right|_{s=0} + \int_0^1 (1-t) \left(\left. \frac{d^2}{ds^2} F_a(\gamma_s) \right|_{s=t} \right) dt \\ &\geq F_a(\gamma_0) - \sqrt{D_a(\gamma_0)} \overline{W}(\gamma_0, \gamma_1) + \frac{\lambda_a}{2} \overline{W}(\gamma_0, \gamma_1)^2. \end{aligned}$$

□

Remark 2.28. *The HWI inequality in Proposition 2.27 immediately implies uniqueness of minimizers for F_a in the set $\{\gamma \in \mathcal{P} \times \mathcal{P} : D_a(\gamma) < +\infty\}$. Indeed, if γ_0 is such that $D_a(\gamma_0) = 0$, then for any other minimizer γ_1 in the above set we have $F_a(\gamma_0) \leq F_a(\gamma_1)$ with equality if and only if $\overline{W}(\gamma_0, \gamma_1) = 0$.*

Corollary 2.29 (Generalized Log Sobolev Inequality). *Denote by γ_* the unique minimizer of F_a . With Assumptions 2.9(i), 2.10 and 2.11, any $\gamma \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $F_a(\gamma), D_a(\gamma) < \infty$ satisfies*

$$D_a(\gamma) \geq 2\lambda_a F_a(\gamma | \gamma_*). \quad (2.17)$$

Proof. This statement follows immediately from Proposition 2.27. Indeed, let $\gamma_1 = \gamma_*$ and $\gamma_0 = \gamma$ in (2.16). Then

$$\begin{aligned} F_a(\gamma | \gamma_*) &\leq \overline{W}(\gamma, \gamma_*) \sqrt{D_a(\gamma)} - \frac{\lambda_a}{2} \overline{W}(\gamma, \gamma_*)^2 \\ &\leq \max_{t \geq 0} \left(\sqrt{D_a(\gamma)} t - \frac{\lambda_a}{2} t^2 \right) = \frac{D_a(\gamma)}{2\lambda_a}. \end{aligned}$$

□

Corollary 2.30 (Talagrand inequality). *Denote by γ_* the unique minimizer of F_a . With Assumptions 2.9(i), 2.10 and 2.11 and $\lambda_a > 0$, it holds*

$$\overline{W}(\gamma, \gamma_*)^2 \leq \frac{2}{\lambda_a} F_a(\gamma | \gamma_*)$$

for any $\gamma \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $F_a(\gamma) < \infty$.

Proof. This is also a direct consequence of Proposition 2.27 by setting $\gamma_0 = \gamma_*$ and $\gamma_1 = \gamma$. Then $F_a(\gamma_*) < \infty$ and $D_a(\gamma_*) = 0$, and the result follows. □

Proof of Theorem 2.13. The entropy terms $\mathcal{H}(\rho)$ and $\mathcal{H}(\mu)$ produce diffusion in ρ and μ for the corresponding PDEs in (2.2). As a consequence, if $\alpha > 0$ (resp. $\beta > 0$) a solution ρ_t (resp. μ_t) to (2.2) and minimizer ρ^* (resp. μ^*) for F_a has to be an L^1 function. Result (a) corresponds to the statements in Proposition 2.25 and Corollary 2.26. To obtain (b), we differentiate the energy F_a along solutions γ_t to the equation (2.2):

$$\begin{aligned} \frac{d}{dt} F_a(\gamma_t) &= \int \delta_\rho F_a[\gamma_t](z) \partial_t \rho_t dz + \int \delta_\mu F_a[\gamma_t](x) \partial_t \mu_t dx \\ &= - \int \|\nabla_z \delta_\rho F_a[\gamma_t](z)\|^2 d\rho_t(z) - \int \|\nabla_x \delta_\mu F_a[\gamma_t](x)\|^2 d\mu_t(x) \\ &= -D_a(\gamma_t) \leq -2\lambda_a F_a(\gamma_t | \gamma_*), \end{aligned}$$

where the last bound follows from Corollary 2.29. Applying Grönwall's inequality, we immediately obtain decay in energy,

$$F_a(\gamma_t | \gamma_*) \leq e^{-2\lambda_a t} F_a(\gamma_0 | \gamma_*).$$

Finally, applying Talagrand's inequality (Corollary 2.30), the decay in energy implies decay in the product Wasserstein metric,

$$\overline{W}(\gamma_t, \gamma_*) \leq c e^{-\lambda_a t},$$

where $c > 0$ is a constant only depending on γ_0 , γ_* and the parameter λ_a . It follows immediately that γ_* is a steady state; see Corollary 2.26. □

2.7 Competitive Setting (Proof of Theorem 2.14)

In the competitive setting, the dynamics are parameterized by the energy functional

$$F_c(\rho, \mu) = \iint f(z, x) d\rho(z) d\mu(x) - \mathcal{R}(\rho) + \mathcal{U}(\mu)$$

where $\mathcal{R}(\rho) = \alpha \mathcal{H}(\rho) + \frac{1}{2} \int (W_1 * \rho)(z) d\rho(z) + \int V_1(z) d\rho(z)$ and $\mathcal{U}(\mu) = \beta \mathcal{H}(\mu) + \frac{1}{2} \int (W_2 * \mu)(x) d\mu(x) + \int V_2(x) d\mu(x)$. For all results in this section, let Assumptions 2.9(ii), 2.10, and 2.11 hold. We define two convexity coefficients,

$$\lambda_{c,1} := \lambda_{f,1} + \lambda_{V,1}, \quad \lambda_{c,2} := \lambda_{f,2} + \lambda_{V,2}$$

where $\lambda_{c,1}$ is the displacement concavity coefficient of F_c with respect to ρ , and $\lambda_{c,2}$ is the displacement convexity coefficient of F_c with respect to μ . We assume that $\lambda_{c,1}, \lambda_{c,2} > 0$ in order to obtain convergence. The rate of convergence depends on the species with weaker convexity; the convergence rate is given by

$$\lambda_c = \min\{\lambda_{c,1}, \lambda_{c,2}\} > 0.$$

Lemma 2.31 (Concavity-Convexity of F_c). *The functional F_c is uniformly displacement $\lambda_{c,2}$ -convex in μ for any fixed $\rho \in \mathcal{P}_2$ and uniformly displacement $\lambda_{c,1}$ -concave in ρ for any fixed $\mu \in \mathcal{P}_2$.*

Proof. For a fixed μ , the $\lambda_{c,1}$ concavity of the functional $F_c(\rho, \mu)$ can be computed as in Proposition 2.22. For a fixed ρ , the $\lambda_{c,2}$ convexity of the functional $F_c(\rho, \mu)$ can be computed as in Proposition 2.22. $\square$

To show contraction, we apply [203, Theorem 23.9], which provides an expression for the time derivative of $\overline{W}(\gamma_t, \gamma'_t)^2$. This theorem requires that the velocities of the trajectories are in L^2 , which we show in the following lemma using the dissipation functional $D_c : \tilde{\mathcal{P}} \times \tilde{\mathcal{P}} \rightarrow \mathbb{R} \cup \{+\infty\}$,

$$D_c(\gamma) := \iint \left\| \begin{bmatrix} \nabla_z \delta_\rho F_c[\gamma](z) \\ \nabla_x \delta_\mu F_c[\gamma](x) \end{bmatrix} \right\|^2 d\gamma(z, x).$$

Lemma 2.32. *Let γ_t be a solution of the dynamics (2.3), with initial condition $\gamma_0 \in \mathcal{P}_2(\mathbb{R}^{d_1}) \times \mathcal{P}_2(\mathbb{R}^{d_2})$ such that $D_c(\gamma_0) < \infty$. Then*

$$D_c(\gamma_t) \leq e^{-2\lambda_c t} D_c(\gamma_0) \quad \forall t \geq 0.$$

Proof. Denote

$$h(t) := \int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 d\rho_t(z) + \int \|\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x)\|^2 d\mu_t(x) = D_c(\gamma_t).$$

By direct differentiation, we can write $h(t)$ as the difference of two dissipations,

$$h(t) = \left[\frac{d}{dt} F_c[\rho_t, \mu_t] - \frac{d}{d\tau} F_c[\rho_\tau, \mu_\tau] \right] \Big|_{\tau=t}.$$

We use that the coupling term cancels when differentiating $h(t)$ so that

$$\frac{d}{dt} h(t) = \left[\frac{d^2}{dt^2} F_c[\rho_t, \mu_t] \right] \Big|_{\tau=t} - \left[\frac{d^2}{d\tau^2} F_c[\rho_t, \mu_t] \right] \Big|_{\tau=t}. \quad (2.18)$$

To show (2.18), recall from the definition of the energy that

$$-\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z) = \nabla_z \left(- \int f(z, x) \mu_t(x) dx + \alpha \log \rho_t(z) + V_1(z) + W_1 * \rho_t(z) \right).$$

Computing the time derivative of the weighted L^2 norm of this velocity gives

$$\begin{aligned} \frac{d}{dt} \int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 d\rho_t(z) &= \int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 \partial_t \rho_t dz \\ &\quad - 2 \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla W_1(z - z') \partial_t \rho_t(z') + \alpha \nabla(\partial_t \rho_t / \rho_t) \rangle d\rho_t(z) dz' \\ &\quad + 2 \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_z f(z, x) \partial_t \mu_t(x) \rangle d\rho_t(z) dx. \end{aligned}$$

For the diffusion term,

$$\begin{aligned} -2\alpha \int \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \alpha \nabla(\partial_t \rho_t / \rho_t) \rangle d\rho_t(z) &= 2\alpha \int \operatorname{div}(\rho_t \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)) \partial_t \rho_t(z) \\ &= -2\alpha \int |\operatorname{div}(\rho_t \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z))|^2 dz. \end{aligned}$$

Considering each of the three remaining terms individually, we start with

$$\begin{aligned} \int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 \partial_t \rho_t dz &= \int \nabla_z \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z) d\rho_t(z) \\ &= 2 \int \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_z^2 \delta_\rho F_c[\rho_t, \mu_t](z) \cdot \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z) \rangle d\rho_t(z). \end{aligned}$$

For the second term, we obtain

$$\begin{aligned} -2 \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla W_1(z - z') \partial_t \rho_t(z') \rangle d\rho_t(z) dz' \\ = -2 \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_{zz'}^2 W_1(z - z') \cdot \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z') \rangle d\rho_t(z) d\rho_t(z'). \end{aligned}$$

The third term is

$$\begin{aligned} & \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_z f(z, x) \partial_t \mu_t(x) \rangle d\rho_t(z) dx \\ &= - \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_{zx}^2 f(z, x) \cdot \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x) \rangle d\rho_t(z) d\mu_t(x). \end{aligned}$$

Likewise, we compute the time derivative of the term $\int \|\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x)\|^2 d\mu_t(x)$ which is nearly identical to that of the ρ_t velocity term, where the diffusion term again can be bounded above by zero. Further, note that the expression for the coupling term is exactly the same as for ρ_t , just with the opposite sign. Due to the zero-sum structure, when both velocities are summed, this term cancels. Thus, we obtain

$$\begin{aligned} \frac{d}{dt} h(t) &= 2 \int \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_z^2 \delta_\rho F_c[\rho_t, \mu_t](z) \cdot \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z) \rangle d\rho_t(z) \\ &\quad - 2 \int \langle \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x), \nabla_x^2 \delta_\mu F_c[\rho_t, \mu_t](x) \cdot \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x) \rangle d\mu_t(x) \\ &\quad - 2 \iint \langle \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z), \nabla_{zz'}^2 W_1(z - z') \cdot \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z') \rangle d\rho_t(z) d\rho_t(z') \\ &\quad - 2 \iint \langle \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x), \nabla_{xx'}^2 W_2(x - x') \cdot \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x') \rangle d\mu_t(x) d\mu_t(x') \\ &\quad - 2\alpha \int |\operatorname{div}(\rho_t \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z))|^2 dz - 2\beta \int |\operatorname{div}(\mu_t \nabla_x \delta_\mu F_c[\rho_t, \mu_t](x))|^2 dx. \quad (2.19) \end{aligned}$$

This expression is equivalent to

$$\left[\frac{d}{dt} \int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 d\rho_t(z) + \frac{d}{d\tau} \int \|\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x)\|^2 d\mu_t(x) \right] \Big|_{\tau=t},$$

which proves (2.18).

By Lemma 2.31, we have $F_c[\rho, \mu]$ with fixed μ is $\lambda_{c,1}$ -concave in ρ and $F_c[\rho, \mu]$ with fixed ρ is $\lambda_{c,2}$ -displacement convex in μ . Therefore, we have for all $t, \tau \geq 0$,

$$\frac{d^2}{dt^2} F_c[\rho_t, \mu_\tau] \leq -2\lambda_{c,1} \frac{d}{dt} F_c[\rho_t, \mu_\tau], \quad \frac{d^2}{d\tau^2} F_c[\rho_t, \mu_\tau] \geq -2\lambda_{c,2} \frac{d}{d\tau} F_c[\rho_t, \mu_\tau].$$

This allows us to use a Bakry-Émery type approach for deriving a decay estimate for $h(t)$. In particular, using (2.18), the fact that $\frac{d}{dt} F_c[\rho_t, \mu_\tau] \geq 0$, $\frac{d}{d\tau} F_c[\rho_t, \mu_\tau] \leq 0$ and $\lambda_c := \min\{\lambda_{c,1}, \lambda_{c,2}\}$ we have

$$\begin{aligned} \frac{d}{dt} h(t) &= \left[\frac{d^2}{dt^2} F_c[\rho_t, \mu_\tau] \right] \Big|_{\tau=t} - \left[\frac{d^2}{d\tau^2} F_c[\rho_t, \mu_\tau] \right] \Big|_{\tau=t} \\ &\leq \left[-2\lambda_{c,1} \frac{d}{dt} F_c[\rho_t, \mu_\tau] + 2\lambda_{c,2} \frac{d}{d\tau} F_c[\rho_t, \mu_\tau] \right] \Big|_{\tau=t} \leq -2\lambda_c h(t). \end{aligned}$$

We conclude using Grönwall's estimate. $\square$

Proposition 2.33 (Contraction). *Fix $T > 0$. Let γ_t and γ'_t be any two solutions of the dynamics (2.3), with initial conditions $\gamma_0, \gamma'_0 \in \mathcal{P}_2^{ac}(\mathbb{R}^{d_1}) \times \mathcal{P}_2^{ac}(\mathbb{R}^{d_2})$ such that $D_c(\gamma_0) < \infty$ and $D_c(\gamma'_0) < \infty$. Assume $\gamma_t, \gamma'_t \in \mathcal{P}_2^{ac}(\mathbb{R}^{d_1}) \times \mathcal{P}_2^{ac}(\mathbb{R}^{d_2})$ for all $t \in [0, T]$ and $\nabla_z \delta_\rho F_c[\gamma_t](z), \nabla_z \delta_\rho F_c[\gamma'_t](z), \nabla_x \delta_\mu F_c[\gamma_t](x), \nabla_x \delta_\mu F_c[\gamma'_t](x)$ are locally Lipschitz in z, x for all $t \in [0, T]$. Then γ_t and γ'_t satisfy*

$$\overline{W}(\gamma_t, \gamma'_t) \leq e^{-\lambda_c t} \overline{W}(\gamma_0, \gamma'_0) \quad \text{for all } t \in [0, T].$$

Remark 2.34. *Although the contraction theorem is stated for measures that are absolutely continuous, the result [203, Theorem 23.9] on the time derivative of $\overline{W}$ can be generalized to the case where $\rho_t = \delta_{z(t)}, \rho'_t = \delta_{z'(t)}$ or $\mu_t = \delta_{x(t)}, \mu'_t = \delta_{x'(t)}$ for all times $t \geq 0$. For details, see Lemma 2.62 for the setting in which $\beta = 0$ with $\mu_0 = \delta_{x(0)}$, and $\alpha > 0$ with $\rho_0 \in \mathcal{P}_2^{ac}$.*

Proof. Define $\nabla \varphi_t(z)$ and $\nabla \psi_t(x)$ so that $\rho_t = \nabla \varphi_{t\#} \rho'_t$ and $\mu_t = \nabla \psi_{t\#} \mu'_t$. Due to Lemma 2.32,

$$\int \|\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)\|^2 d\rho_t(z) + \int \|\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x)\|^2 d\mu_t(x) < \infty \quad \forall t \geq 0,$$

with the same holding for γ' . From [203, Theorem 23.9], the time derivative of the joint metric along solutions to (2.3) can be controlled by

$$\begin{aligned} \frac{d}{dt} \overline{W}(\gamma_t, \gamma'_t)^2 &= 2 \iint \left[\nabla \varphi_t^{-1}(z) - z \right] \cdot \left[\begin{array}{c} -\nabla \delta_\rho F_c[\gamma_t](z) \\ \nabla \delta_\mu F_c[\gamma_t](x) \end{array} \right] d\gamma_t(z, x) \\ &\quad + 2 \iint \left[\nabla \varphi_t(z) - z \right] \cdot \left[\begin{array}{c} -\nabla \delta_\rho F_c[\gamma'_t](z) \\ \nabla \delta_\mu F_c[\gamma'_t](x) \end{array} \right] d\gamma'_t(z, x) \\ &= 2 \iint \left[\nabla \varphi_t(z) - z \right] \cdot \left[\begin{array}{c} \nabla \delta_\rho F_c[\gamma_t](\nabla \varphi_t(z)) - \nabla \delta_\rho F_c[\gamma'_t](z) \\ \nabla \delta_\mu F_c[\gamma'_t](x) - \nabla \delta_\mu F_c[\gamma_t](\nabla \psi_t(x)) \end{array} \right] d\gamma'_t(z, x) \\ &\leq -2\lambda_c \overline{W}(\gamma_t, \gamma'_t)^2 \quad \forall t \in (0, T), \end{aligned}$$

where the last inequality follows from Lemma 2.61. By Grönwall's lemma, exponential convergence follows. For time $t = 0$, the inequality holds by definition since $e^{-\lambda_c t}|_{t=0} = 1$. $\square$

Next, our goal is to show that the semigroup for the dynamics (2.3) maps to $\mathcal{P}_2 \times \mathcal{P}_2$. The key ingredient is control of second moments along the evolution. More precisely, we will show that the second moments converge exponentially to a ball and remain in that ball for all time. This result will then allow us to apply a contractive inequality to prove the existence of a steady state for the dynamics (2.3).

Proposition 2.35 (Uniformly Bounded Second Moments). *Let γ_t be a solution to (2.3) with $\gamma_0 \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $D_c(\gamma_0) < \infty$. If $\alpha = 0$, assume $\rho_0 = \delta_{z_0}$ for some $z_0 \in \mathbb{R}^{d_1}$. If $\beta = 0$, assume $\mu_0 = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^{d_2}$. Then $\bar{W}(\gamma_t, \bar{\delta})^2$ satisfies*

$$\frac{d}{dt} \bar{W}(\gamma_t, \bar{\delta})^2 \leq -\lambda_c \bar{W}(\gamma_t, \bar{\delta})^2 + 2\hat{c},$$

for some $\hat{c} \geq 0$, where $\bar{\delta}(z, x) = \delta_{(0,0)}(z, x)$. For any time $t \geq 0$, it holds that

$$\int \|z\|^2 d\rho_t(z) + \int \|x\|^2 d\mu_t(x) \leq K := \max \left\{ \bar{W}(\gamma_0, \bar{\delta}), \frac{2\hat{c}}{\lambda_c} \right\}.$$

Proof. If $\alpha = 0$ ($\beta = 0$), then ρ_t (μ_t) remains a Dirac Delta for all times, and so its second moment vanishes. Let us assume $\alpha, \beta > 0$. Thanks to the diffusion, $\gamma_t \in \mathcal{P}^{ac} \times \mathcal{P}^{ac}$ for all $t > 0$, and thanks to Lemma 2.32 we have $D_c(\gamma_t) < \infty$ for all $t \geq 0$. The sum of the squared second moments can be written as $\bar{W}(\gamma_t, \bar{\delta})^2$. Directly differentiating the second moments along solutions to (2.3), we have

$$\frac{d}{dt} \bar{W}(\gamma_t, \bar{\delta})^2 \leq 2 \iint \begin{bmatrix} z \\ x \end{bmatrix} \cdot \begin{bmatrix} \nabla_z \delta_\rho F_c[\gamma_t](z) \\ -\nabla_x \delta_\mu F_c[\gamma_t](x) \end{bmatrix} d\rho_t(z) d\mu_t(x). \quad (2.20)$$

Our goal is to upper-bound the right-hand side in terms of $\bar{W}(\gamma_t, \bar{\delta})^2$. We will compute the terms in F_c separately, starting with the entropy terms. We have

$$\begin{aligned} - \int z \cdot \nabla_z \delta_\rho \mathcal{H}(\rho_t) d\rho_t(z) &= - \int z \cdot \nabla \log \rho_t d\rho_t(z) = - \int z \cdot \nabla_z \rho_t(z) dz = d_1 \int d\rho_t(z) = d_1, \\ - \int x \cdot \nabla_x \delta_\mu \mathcal{H}(\mu_t) d\mu_t(x) &= - \int x \cdot \nabla \log \mu_t d\mu_t(x) = - \int x \cdot \nabla_x \mu_t(x) dx = d_2 \int d\mu_t(x) = d_2. \end{aligned}$$

For the remaining terms, we will use the convexity inequality

$$f(y) \geq f(y') + \nabla f(y') \cdot (y - y') + \frac{\lambda}{2} \|y - y'\|^2 \quad \forall y, y' \in \mathbb{R}^d.$$

Applying this to W_1 , we use a change of variables on half of the integral and use the symmetry of W_1 , in particular that $\nabla W_1(z - z') = -\nabla W_1(z' - z)$, to rewrite

$$- \iint z \cdot \nabla W_1(z - z') d\rho_t(z) d\rho_t(z') = -\frac{1}{2} \int (z - z') \cdot \nabla W_1(z - z') d\rho_t(z) d\rho_t(z').$$

Selecting $y' = z - z'$ and $y = 0$ results in

$$\begin{aligned} -\frac{1}{2} \iint (z - z') \cdot \nabla W_1(z - z') d\rho_t(z) d\rho_t(z') &\leq \iint -\frac{\lambda_{W,1}}{4} \|z - z'\|^2 d\rho_t(z) d\rho_t(z') \\ &\quad + \iint [W_1(0) - W_1(z - z')] d\rho_t(z) d\rho_t(z') \leq 0, \end{aligned}$$

since $W_1(0) \leq W_1(z)$ for all z by symmetry and convexity of W_1 . A similar estimate can be computed for W_2 . For the cross-term and potential terms, we will use the convexity inequality for each species, with $y = 0$ and $y' = z$ for $z \cdot \nabla_z f(z, x)$ and $y = 0$ and $y' = x$ for $-x \cdot \nabla_x f(z, x)$,

$$\begin{aligned} \iint \begin{bmatrix} z \\ x \end{bmatrix} \cdot \begin{bmatrix} \nabla_z(f(z, x) - V_1(z)) \\ -\nabla_x(f(z, x) - V_2(x)) \end{bmatrix} d\rho_t(z) d\mu_t(x) &\leq - \int \frac{\lambda_{c,1}}{2} \|z\|^2 d\rho_t(z) - \int \frac{\lambda_{c,2}}{2} \|x\|^2 d\mu_t(x) \\ &+ \iint (-f(0, x) + f(z, 0) + V_1(0) - V_1(z) + f(z, x) - f(z, x) + V_2(0) - V_2(x)) d\gamma_t(z, x). \end{aligned}$$

Next, we use that $f(z, 0) - V_1(z)$ is $\lambda_{c,1}$ concave in z , $f(0, x) + V_2(x)$ is $\lambda_{c,2}$ convex in x , and continuity to define

$$c := V_1(0) + V_2(0) + \max_{z \in \mathbb{R}^{d_1}} f(z, 0) - V_1(z) + \max_{x \in \mathbb{R}^{d_2}} -f(0, x) - V_2(x) < \infty,$$

which gives an upper-bound for the cross term

$$\iint \begin{bmatrix} z \\ x \end{bmatrix} \cdot \begin{bmatrix} \nabla_z f(z, x) \\ -\nabla_x f(z, x) \end{bmatrix} d\rho_t(z) d\mu_t(x) \leq - \int \frac{\Lambda_{f,1}}{2} \|z\|^2 d\rho_t(z) - \int \frac{\lambda_{f,2}}{2} \|x\|^2 d\mu_t(x) + c.$$

Note that $c \geq 0$ because $\max_{z \in \mathbb{R}^{d_1}} f(z, 0) - V_1(z) \geq f(0, 0) - V_1(0)$ and $\max_{x \in \mathbb{R}^{d_2}} -f(0, x) - V_2(x) \geq -f(0, 0) - V_2(0)$. Combining all terms gives

$$\frac{1}{2} \frac{d}{dt} \overline{W}(\gamma_t, \bar{\delta})^2 \leq -\frac{\lambda_c}{2} \overline{W}(\gamma_t, \bar{\delta})^2 + \hat{c},$$

where $\hat{c} = c + \alpha d_1 + \beta d_2$. It holds that $\hat{c} \geq 0$ because all terms are non-negative. The solution to this ODE satisfies

$$\overline{W}(\gamma_t, \bar{\delta})^2 \leq \max \left\{ \overline{W}(\gamma_0, \bar{\delta})^2, \frac{2\hat{c}}{\lambda_c} \right\} \quad \forall t \geq 0.$$

□

Proposition 2.36 (Existence and Uniqueness of Steady States). *There exists a unique steady state $\gamma_\infty = (\rho_\infty, \mu_\infty) \in \tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$ of equation (2.3) according to Definition 2.6 with $\int \rho_\infty = 1, \int \mu_\infty = 1$. This steady state is a Nash equilibrium for F_c . Additionally,*

- if $\alpha > 0$, the steady state γ_∞ satisfies $\rho_\infty \in L_+^1(\mathbb{R}^{d_1}) \cap C^2(\mathbb{R}^{d_1})$ with $\|\rho_\infty\|_1 = 1$ and $\text{supp}(\rho_\infty) = \mathbb{R}^{d_1}$;
- if $\beta > 0$, the steady state γ_∞ satisfies $\mu_\infty \in L_+^1(\mathbb{R}^{d_2}) \cap C^2(\mathbb{R}^{d_2})$ with $\|\mu_\infty\|_1 = 1$ and $\text{supp}(\mu_\infty) = \mathbb{R}^{d_2}$.

Proof. We split the proof into three steps. First, (i) we prove there exists $\gamma_\infty \in \mathcal{P}_2 \times \mathcal{P}_2$ satisfying $\overline{W}(\gamma_\infty, \gamma(t)) = 0$ for all $t \geq 0$ with $\gamma(t)$ the solution to (2.3) with initial condition γ_∞ . Then (ii) we will show that γ_∞ is also a critical point; specifically it is a Nash equilibrium. Lastly, (iii) we show γ_∞ has the required regularity properties to satisfy Definition 2.6.

(i) Fix $\gamma_0 \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $D_c(\gamma_0) < \infty$. If $\alpha = 0$, assume $\rho_0 = \delta_{z_0}$ for some $z_0 \in \mathbb{R}^{d_1}$. If $\beta = 0$, assume $\mu_0 = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^{d_2}$. By Proposition 2.35, we know that $\gamma_t \in \mathcal{P}_2 \times \mathcal{P}_2$ for all $t \geq 0$. If $\alpha, \beta > 0$, we have $\gamma(t) \in \mathcal{P}_2^{ac} \times \mathcal{P}_2^{ac}$ thanks to the diffusion. Contraction in $\mathcal{P}_2^{ac} \times \mathcal{P}_2^{ac}$ follows from Proposition 2.33. If $\alpha = 0$ or $\beta = 0$, contraction is in $\mathcal{P}_2$ instead of $\mathcal{P}_2^{ac}$ (see Remark 2.34). Let $T(t) : \mathcal{P}_2 \times \mathcal{P}_2 \rightarrow \mathcal{P}_2 \times \mathcal{P}_2$ be the semigroup for (2.3). Using a contraction result, [52, Lemma 7.3], there exists a unique steady state $\gamma_\infty \in \mathcal{P}_2^{ac} \times \mathcal{P}_2^{ac}$ (or $\mathcal{P}_2 \times \mathcal{P}_2$ resp.), that is,

$$\overline{W}(\gamma_\infty, T(t)\gamma_\infty) = 0.$$

(ii) To show that γ_∞ is a critical point, we consider the optimization problems:

$$c_1 = \sup_{\rho \in \mathcal{P}_2} F_c(\rho, \mu_\infty), \quad c_2 = \inf_{\mu \in \mathcal{P}_2} F_c(\rho_\infty, \mu).$$

First, note that the supremum c_1 and infimum c_2 are in fact attained by a unique maximizer $\rho_\dagger \in \mathcal{P}_2$ and unique minimizer $\mu_\dagger \in \mathcal{P}_2$. This can be shown using similar arguments as those in Proposition 2.25; concavity of $F_c(\cdot, \mu_\infty)$ and convexity of $F_c(\rho_\infty, \cdot)$ follow from Lemma 2.31. The functional $F_c(\cdot, \mu_\infty)$ has an upper-bound due to concavity via a similar argument as in Lemma 2.24 and Lemma 2.65, and $F_c(\rho_\infty, \cdot)$ has a lower-bound using the same argument with convexity. Upper and lower-semicontinuity complete the set of required ingredients for the proof of Proposition 2.25. We can then write the well-defined optimization problem

$$\rho_\dagger = \operatorname{argmax}_{\rho \in \mathcal{P}_2} F_c(\rho, \mu_\infty), \quad \mu_\dagger = \operatorname{argmin}_{\mu \in \mathcal{P}_2} F_c(\rho_\infty, \mu). \quad (2.21)$$

As a maximizer and minimizer respectively, they satisfy the Euler-Lagrange (EL) conditions

$$\delta_\rho F_c[\rho_\dagger, \mu_\infty](z) = c_1 \quad \forall z \in \operatorname{supp} \rho_\dagger, \quad \delta_\mu F_c[\rho_\infty, \mu_\dagger](x) = c_2 \quad \forall x \in \operatorname{supp} \mu_\dagger.$$

The EL condition for $\rho_\dagger$ is

$$\begin{aligned} \delta_\rho F_c[\rho_\dagger, \mu_\infty](z) &= -\alpha \log \rho_\dagger(z) + \int f(z, x) d\mu_\infty(x) - V_1(z) - (W_1 * \rho_\dagger)(z) \\ &= \tilde{c}_1 \quad \forall z \in \operatorname{supp} \rho_\dagger. \end{aligned}$$

Note that the left-hand side of the EL condition is finite for any fixed z . Rearranging, in the case when $\alpha > 0$, we have

$$\rho_\dagger(z) = \exp \left(\frac{1}{\alpha} \left[-\tilde{c}_1 + \int f(z, x) d\mu_\infty(x) - W_1 * \rho_\dagger(z) - V_1(z) \right] \right). \quad (2.22)$$

We have $W_1 * \rho_{\dagger} \in C^2$ due to $W_1 \in C^2$, $\int f(z, x) d\mu_{\infty}(x) \in C^2$ because $f(\cdot, x) \in C^2$, and $V_1 \in C^2$ from Assumptions 2.9(ii), 2.10, and 2.11. Therefore $\rho_{\dagger} \in C^2$. If $\alpha = 0$, then $\delta_{\rho} F_c[\rho_{\dagger}, \mu_{\infty}](\cdot) \in C^2$ follows immediately. We conclude that $\delta_{\rho} F_c[\rho_{\dagger}, \mu_{\infty}](\cdot) \in C^2$ and $\operatorname{div}(\rho_{\dagger} \nabla \delta_{\rho} F_c[\rho_{\dagger}, \mu_{\infty}](z)) = 0$ pointwise a.e. Similarly, $\operatorname{div}(\mu_{\dagger} \nabla \delta_{\mu} F_c[\rho_{\infty}, \mu_{\dagger}](x)) = 0$ pointwise a.e., and if $\beta > 0$, then $\mu_{\dagger} \in C^2$.

Define $S_1(t) : \mathcal{P}_2 \rightarrow \mathcal{P}_2$ and $S_2(t) : \mathcal{P}_2 \rightarrow \mathcal{P}_2$ as the semigroups for the uncoupled dynamics given by

$$\partial_t \hat{\rho} = -\operatorname{div}(\hat{\rho} \nabla \delta_{\rho} F_c[\hat{\rho}, \mu_{\infty}](z)) , \quad \partial_t \hat{\mu} = \operatorname{div}(\hat{\mu} \nabla \delta_{\mu} F_c[\rho_{\infty}, \hat{\mu}](x)) , \quad (2.23)$$

respectively. Note that $S_i(t)$ maps to $\mathcal{P}_2$ by the same argument as in Proposition 2.35. Due to the regularity of $\rho_{\dagger}$ and $\mu_{\dagger}$ and the fact that they are critical points, they are also steady states of (2.23)

$$\mathcal{W}_2(\rho_{\dagger}, S_1(t)\rho_{\dagger}) = 0 , \quad \mathcal{W}_2(\mu_{\dagger}, S_2(t)\mu_{\dagger}) = 0 .$$

Since γ_{∞} is a steady state of (2.3),

$$\mathcal{W}_2(\rho_{\infty}, S_1(t)\rho_{\infty}) = 0 , \quad \mathcal{W}_2(\mu_{\infty}, S_2(t)\mu_{\infty}) = 0 .$$

Next, we again apply the contraction theorem [52, Lemma 7.3] to (2.23), which states that there is a unique steady state for $S_1(t)$ and for $S_2(t)$. Therefore,

$$\overline{W}(\gamma_{\dagger}, \gamma_{\infty}) = 0 .$$

(iii) If $\alpha, \beta > 0$, then $\gamma_{\dagger} \in C^2 \times C^2$. Since equality of $\gamma_{\dagger}$ with γ_{∞} holds almost everywhere with respect to the Lebesgue measure, it holds that γ_{∞} is a weak solution to (2.3), such that $\gamma_{\infty} \in (W_{loc}^{1,2} \cap L_{loc}^{\infty}) \times (W_{loc}^{1,2} \cap L_{loc}^{\infty})$. When $\alpha > 0$, the argument for $\operatorname{supp}(\rho_{\infty}) = \mathbb{R}^{d_1}$ and $\rho_{\infty} \in C^2(\mathbb{R}^{d_1})$ follows as in Corollary 2.26, where instead of showing that $\int f(z, x) d\mu_{\infty}(x) + V_1(z) < \infty$ for any fixed $z \in \mathbb{R}^{d_1}$, showing $-\int f(z, x) d\mu_{\infty}(x) + V_1(z) < \infty$ for any fixed $z \in \mathbb{R}^{d_1}$ gives the same result. This holds because $-\int f(z, x) d\mu_{\infty}(x) + V_1(z)$ is continuous in z , just as $\int f(z, x) d\mu_{\infty}(x) + V_1(z)$ is continuous in z in the setting of Corollary 2.26. The setting when $\beta > 0$ follows similarly to $\alpha > 0$. $\square$

Proposition 2.37 (Uniqueness of Nash Equilibrium). *There exists at most one Nash equilibrium in $\tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$ for the energy functional F_c .*

Proof. Suppose there exists two equilibria, $\gamma_* \in \tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$ and $\gamma_{\infty} \in \tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$ such that $\overline{W}(\gamma_{\infty}, \gamma_*) > 0$. In the weak sense, the following hold:

$$\begin{aligned} \nabla_z \delta_{\rho} F_c[\gamma_{\infty}](z) &= 0 \quad \forall z \in \operatorname{supp} \rho_{\infty} , & \nabla_z \delta_{\rho} F_c[\gamma_*](z) &= 0 \quad \forall z \in \operatorname{supp} \rho_* , \\ \nabla_x \delta_{\mu} F_c[\gamma_{\infty}](x) &= 0 \quad \forall x \in \operatorname{supp} \mu_{\infty} , & \nabla_x \delta_{\mu} F_c[\gamma_*](x) &= 0 \quad \forall x \in \operatorname{supp} \mu_* . \end{aligned}$$

We plug γ_* and γ_∞ into the inequality proved in Lemma 2.61, with $\rho_* = \nabla\varphi_\# \rho_\infty$ and $\mu_* = \nabla\psi_\# \mu_\infty$. Note that both $\nabla\varphi$ and $\nabla\psi$ exist and are invertible (on the support of ρ_∞ and μ_∞ resp.) because the measures are either absolutely continuous or propagating as Diracs, and so the left-hand side in Lemma 2.61 is well-defined. Hence,

$$\begin{aligned} & \iint \left[\frac{z - \nabla\varphi(z)}{x - \nabla\psi(x)} \right] \cdot \left[\frac{\nabla_z \delta_\rho F_c[\gamma_*](\nabla\varphi(z))}{-\nabla_x \delta_\mu F_c[\gamma_*](\nabla\psi(x))} \right] d\rho_\infty(z) d\mu_\infty(x) \\ &= \iint \left[\frac{(\nabla\varphi)^{-1}(z) - z}{(\nabla\psi)^{-1}(x) - x} \right] \cdot \left[\frac{\nabla_z \delta_\rho F_c[\gamma_*](z)}{-\nabla_x \delta_\mu F_c[\gamma_*](x)} \right] d\rho_*(z) d\mu_*(x) = 0 \geq \lambda_c \overline{W}(\gamma_\infty, \gamma_*)^2, \end{aligned}$$

which is a contradiction and therefore the Nash equilibrium is unique. $\square$

Now we present the proof of Theorem 2.14.

Proof of Theorem 2.14. From Proposition 2.36 we have existence and uniqueness of steady states of (2.3) in $\tilde{\mathcal{P}}_2 \times \tilde{\mathcal{P}}_2$, and that the steady state is a Nash equilibrium for F_c . This Nash equilibrium is unique from Proposition 2.37. This concludes the proof of (a).

For (b), the uniform bound on the second moment follows from Proposition 2.35. Further, let $(\rho_0, \mu_0) \in \mathcal{P}_2^{ac}(\mathbb{R}^{d_1}) \times \mathcal{P}_2^{ac}(\mathbb{R}^{d_2})$, $\rho_\infty = \nabla\varphi_\# \rho_t$ and $\mu_\infty = \nabla\psi_\# \mu_t$. We compute the time derivative of the joint Wasserstein-2 metric using [203, Theorem 23.9]. We have

$$\begin{aligned} \frac{d}{dt} \overline{W}(\gamma_t, \gamma_\infty)^2 &= 2 \iint \left[\frac{\nabla\varphi(z) - z}{\nabla\psi(x) - x} \right] \cdot \left[\frac{-\nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)}{\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x)} \right] d\gamma_t(z, x) \\ &= 2 \iint \left[\frac{\nabla\varphi(z) - z}{\nabla\psi(x) - x} \right] \cdot \left[\frac{\nabla_z \delta_\rho F_c[\rho_\infty, \mu_\infty](\nabla\varphi(z)) - \nabla_z \delta_\rho F_c[\rho_t, \mu_t](z)}{\nabla_x \delta_\mu F_c[\rho_t, \mu_t](x) - \nabla_x \delta_\mu F_c[\rho_\infty, \mu_\infty](\nabla\psi(x))} \right] d\gamma_t(z, x), \end{aligned}$$

where we can add the second term because γ_∞ is a steady state, so

$$\int \nabla_z \delta_\rho F_c[\rho_\infty, \mu_\infty](\nabla\varphi(z)) d\rho_t(z) = 0.$$

This can be shown by rewriting the support of the steady state using the pushforward from ρ_∞ to ρ_t . A similar argument applies to the first variation in μ , and in the setting in which the measure propagating is a Dirac. Now we can apply Lemma 2.61 to obtain $\frac{d}{dt} \overline{W}(\gamma_t, \gamma_\infty)^2 \leq -2\lambda_c \overline{W}(\gamma_t, \gamma_\infty)^2$ which gives exponential convergence with rate λ_c using Grönwall's Lemma, completing the proof of (b). $\square$

Appendix

The appendices contain proofs in the timescale-separated settings as well as supporting results. The first two appendices focus on the competitive setting: Appendix 2.A treats the case of a fast

algorithm (proof of Theorem 2.19), and Appendix 2.B of a fast population (proof of Theorem 2.20). In both cases, the corresponding objective functional is given by $G(\rho, x) : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, \infty]$ as follows:

$$G(\rho, x) = \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2 - \mathcal{R}(\rho),$$

where

$$\mathcal{R}(\rho) := \begin{cases} \alpha KL(\rho | \tilde{\rho}) + \frac{1}{2} \int W * \rho d\rho, & \text{if } \rho \ll \tilde{\rho}, \\ +\infty & \text{else,} \end{cases}$$

where $\alpha > 0$ and the reference measure $\tilde{\rho}$ satisfies $\tilde{\rho} \in \mathcal{P}(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$, $\log \tilde{\rho} \in C^2(\mathbb{R}^d)$ and $\text{supp } \tilde{\rho} = \mathbb{R}^d$. Throughout Appendices 2.A and 2.B, we will assume that Assumptions 2.9(ii), 2.10 and 2.11 hold such that

$$\lambda_b := \alpha \tilde{\lambda} - \Lambda_1 > 0 \quad \text{and} \quad \lambda_d := \lambda_1 + \lambda_2 + \kappa = \lambda_{V,2} + \lambda_{f,2} > 0.$$

Appendix 2.C contains auxiliary lemmas used throughout the paper.

2.A Competitive Objective, Fast Algorithm (Proof of Theorem 2.19)

Define

$$G_b(\rho) := G(\rho, b(\rho)),$$

where the best response $b(\rho)$ of the classifier for any $\rho \in \mathcal{P}(\mathbb{R}^d)$ with finite energy is given by

$$b(\rho) := \underset{\bar{x} \in \mathbb{R}^d}{\operatorname{argmin}} G(\rho, \bar{x}). \quad (2.24)$$

Since for such fixed $\rho \in \mathcal{P}(\mathbb{R}^d)$, the energy $G(\rho, \cdot)$ is strictly λ_d -convex in x , it has a unique minimizer, and so the best response is well defined. If $\rho \in \mathcal{P}(\mathbb{R}^d)$ is such that $G(\rho, x_1) = \pm\infty$ for some $x_1 \in \mathbb{R}^d$, then $G(\rho, x) = \pm\infty$ for any $x \in \mathbb{R}^d$ and we define $b(\rho) := 0$ in that case. We begin with auxiliary results computing the first variations of the best response b and then the different terms in the first variation of $G_b(\rho)$. The following result can be deduced directly from Definition 2.6.

Lemma 2.38 (Steady states for (2.8)). *Let $\rho_\infty \in L^1_+(\mathbb{R}^d) \cap L^\infty_{loc}(\mathbb{R}^d)$ with $\|\rho_\infty\|_1 = 1$. Then ρ_∞ is a steady state for the system (2.8) if $\rho_\infty \in W^{1,2}_{loc}(\mathbb{R}^d)$, $\nabla W * \rho_\infty \in L^1_{loc}(\mathbb{R}^d)$, ρ_∞ is absolutely continuous with respect to $\tilde{\rho}$, and ρ_∞ satisfies*

$$\nabla_z \left(f_1(z, b(\rho_\infty)) - \alpha \log \left(\frac{\rho_\infty(z)}{\tilde{\rho}(z)} \right) - W * \rho_\infty(z) \right) = 0, \quad \text{on } \text{supp}(\rho_\infty), \quad (2.25)$$

in the sense of distributions, where $b(\rho_\infty) := \operatorname{argmin}_x G(\rho_\infty, x)$.

Lemma 2.39 (First variation of the best response). *The first variation of the best response of the classifier at ρ (if it exists) is*

$$\delta_\rho b[\rho](z) = -Q(\rho)^{-1} \nabla_x f_1(z, b(\rho)) \quad \text{for almost every } z \in \mathbb{R}^d,$$

where $Q(\rho) \succeq (\kappa + \lambda_1 + \lambda_2) \mathbf{I}_d$ is a symmetric matrix, constant in z and x , defined as

$$Q(\rho) := \kappa \mathbf{I}_d + \int \nabla_x^2 f_1(z, b(\rho)) d\rho(z) + \int \nabla_x^2 f_2(z, b(\rho)) d\pi(z).$$

In particular, we then have for any $\psi \in C_c^\infty(\mathbb{R}^d)$ with $\int \psi dz = 0$ that

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left\| b[\rho + \varepsilon\psi] - b[\rho] - \varepsilon \int \delta_\rho b[\rho](z) \psi(z) dz \right\| = 0.$$

Proof. Let $\psi \in C_c^\infty(\mathbb{R}^d)$ with $\int \psi dz = 0$ and fix $\varepsilon > 0$. Any minimizer of $G(\rho + \varepsilon\psi, x)$ for fixed ρ must satisfy

$$\nabla_x G(\rho + \varepsilon\psi, b(\rho + \varepsilon\psi)) = 0.$$

Differentiating in ε , we obtain

$$\begin{aligned} & \int \delta_\rho \nabla_x G[\rho + \varepsilon\psi, b(\rho + \varepsilon\psi)] \psi(z) dz \\ & + \nabla_x^2 G(\rho + \varepsilon\psi, b(\rho + \varepsilon\psi)) \int \delta_\rho b[\rho + \varepsilon\psi](z) \psi(z) dz = 0. \end{aligned} \tag{2.26}$$

Next, we explicitly compute all terms involved in (2.26). Computing the derivatives yields

$$\begin{aligned} \nabla_x G(\rho, x) &= \int \nabla_x f_1(z, x) d\rho(z) + \int \nabla_x f_2(z, x) d\pi(z) + \kappa(x - x_0) \\ \delta_\rho \nabla_x G[\rho, x](z) &= \nabla_x f_1(z, x) \\ \nabla_x^2 G(\rho, x) &= \int \nabla_x^2 f_1(z, x) d\rho(z) + \int \nabla_x^2 f_2(z, x) d\pi(z) + \kappa \mathbf{I}_d. \end{aligned}$$

Note that $\nabla_x^2 G$ is invertible because $\lambda_1 + \lambda_2 + \kappa > 0$. Inverting this term and substituting these expressions into (2.26) for $\varepsilon = 0$ gives

$$\begin{aligned} \int \delta_\rho b[\rho](z) \psi(z) dz &= - \left[\kappa \mathbf{I}_d + \int \nabla_x^2 f_1(z, b(\rho)) d\rho(z) \right. \\ &\quad \left. + \int \nabla_x^2 f_2(z, b(\rho)) d\pi(z) \right]^{-1} \int \nabla_x f_1(z, b(\rho)) \psi(z) dz \\ &= - \int Q(\rho)^{-1} \nabla_x f_1(z, b(\rho)) \psi(z) dz. \end{aligned}$$

Finally, the lower bound on $Q(\rho)$ follows thanks to Assumptions 2.9(ii) and 2.10. $\square$

Lemma 2.40 (First variation of G_b). *The first variation of G_b is given by*

$$\delta_\rho G_b[\rho](z) = h_1(z) + h_2(z) + \kappa h_3(z) - \delta_\rho \mathcal{R}[\rho](z),$$

where

$$\begin{aligned} h_1(z) &:= \frac{\delta}{\delta \rho} \left(\int f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}) \right) (z) = \left\langle \int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle + f_1(z, b(\rho)), \\ h_2(z) &:= \frac{\delta}{\delta \rho} \left(\int f_2(\tilde{z}, b(\rho)) d\pi(\tilde{z}) \right) (z) = \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\pi(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle, \\ h_3(z) &:= \frac{1}{2} \frac{\delta}{\delta \rho} \|b(\rho) - x_0\|^2 = \left\langle b(\rho) - x_0, \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle, \end{aligned}$$

and

$$\delta_\rho \mathcal{R}[\rho](z) = \alpha \log(\rho(z)/\tilde{\rho}(z)) + (W * \rho)(z).$$

Proof. We begin with general expressions for the Taylor expansions of $b : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ and $f_i(z, b(\cdot)) : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ for $i = 1, 2$ around ρ . Let $\psi \in \mathcal{T}$ with $\mathcal{T} = \{\psi : \int \psi(z) dz = 0\}$. Then

$$b(\rho + \varepsilon \psi) = b(\rho) + \varepsilon \int \frac{\delta b}{\delta \rho}[\rho](z') \psi(z') dz' + O(\varepsilon^2) \quad (2.27)$$

and

$$f_i(z, b(\rho + \varepsilon \psi)) = f_i(z, b(\rho)) + \varepsilon \left\langle \nabla_x f_i(z, b(\rho)), \int \frac{\delta b}{\delta \rho}[\rho](z') \psi(z') dz' \right\rangle + O(\varepsilon^2). \quad (2.28)$$

We compute explicitly each of the first variations:

(i) Using (2.28), we have

$$\begin{aligned} \int \psi(z) h_1(z) dz &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left[\int f_1(z, b(\rho + \varepsilon \psi)) (\rho(z) + \varepsilon \psi(z)) dz - \int f_1(z, b(\rho)) \rho(z) dz \right] \\ &= \left\langle \int \nabla_x f_1(z, b(\rho)) d\rho(z), \int \frac{\delta b(\rho)}{\delta \rho}[\rho](z') \psi(z') dz' \right\rangle + \int f_1(z, b(\rho)) \psi(z) dz \\ &= \int \left\langle \int \nabla_x f_1(z, b(\rho)) d\rho(z), \frac{\delta b(\rho)}{\delta \rho}[\rho](z') \right\rangle \psi(z') dz' + \int f_1(z, b(\rho)) \psi(z) dz \\ \Rightarrow h_1(z) &= \left\langle \int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle + f_1(z, b(\rho)). \end{aligned}$$

(ii) Similarly, using again (2.28),

$$\begin{aligned} \int \psi(z) h_2(z) dz &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left[\int f_2(z, b(\rho + \varepsilon \psi)) d\pi(z) - \int f_2(z, b(\rho)) \pi(z) dz \right] \\ &= \int \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\pi(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle \psi(z) dz \\ \Rightarrow h_2(z) &= \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\pi(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle. \end{aligned}$$

(iii) Finally, from (2.27) it follows that

$$\begin{aligned} \int \psi(z) h_3(z) dz &= \lim_{\varepsilon \rightarrow 0} \frac{1}{2\varepsilon} \left[\langle b(\rho + \varepsilon\psi) - x_0, b(\rho + \varepsilon\psi) - x_0 \rangle - \langle b(\rho) - x_0, b(\rho) - x_0 \rangle \right] \\ &= \int \left\langle b(\rho) - x_0, \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle \psi(z) dz \\ \Rightarrow h_3(z) &= \left\langle b(\rho) - x_0, \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle. \end{aligned}$$

Finally, the expression for $\delta_\rho \mathcal{R}[\rho]$ follows by direct computation (2.24). $\square$

Proposition 2.41 (Danskin-Type Result). *Denote $G_b(\rho) := G(\rho, b(\rho))$ with $b(\rho)$ given by (2.24). Then*

$$\delta_\rho G_b[\rho] = \delta_\rho G[\rho] \Big|_{x=b(\rho)}.$$

Proof. We start by computing $\delta_\rho G(\cdot, x)[\rho](z)$ for any $z, x \in \mathbb{R}^d$:

$$\delta_\rho G(\cdot, x)[\rho](z) = f_1(z, x) - \delta_\rho \mathcal{R}[\rho](z). \quad (2.29)$$

Next, we compute $\delta_\rho G_b$. Using Lemma 2.40, the first variation of G_b is given by

$$\begin{aligned} \delta_\rho G_b[\rho](z) &= h_1(z) + h_2(z) + \kappa h_3(z) - \delta_\rho \mathcal{R}[\rho](z) \\ &= - \left\langle \left[\int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}) + \int \nabla_x f_2(\tilde{z}, b(\rho)) d\pi(\tilde{z}) + \kappa(b(\rho) - x_0) \right], \delta_\rho b[\rho](z) \right\rangle \\ &\quad + f_1(z, b(\rho)) - \delta_\rho \mathcal{R}[\rho](z). \end{aligned}$$

Note that

$$\nabla_x G(\rho, x) = \int \nabla_x f_1(\tilde{z}, x) d\rho(\tilde{z}) + \int \nabla_x f_2(\tilde{z}, x) d\pi(\tilde{z}) + \kappa(x - x_0), \quad (2.30)$$

and by the definition of the best response $b(\rho)$, we have $\nabla_x G(\rho, x)|_{x=b(\rho)} = 0$. Substituting into the expression for $\delta_\rho G_b$ and using (2.29), we obtain

$$\delta_\rho G_b[\rho](z) = f_1(z, b(\rho)) - \delta_\rho \mathcal{R}[\rho](z) = \delta_\rho G(\cdot, x)[\rho](z) \Big|_{x=b(\rho)}.$$

This concludes the proof. $\square$

Proposition 2.42 (Displacement concavity). *Fix $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$. Along any geodesic $(\rho_s)_{s \in [0,1]} \in \mathcal{P}_2(\mathbb{R}^d)$ connecting ρ_0 to ρ_1 , we have for all $s \in [0, 1]$*

$$\frac{d^2}{ds^2} G_b(\rho_s) \leq -\lambda_b \mathcal{W}_2(\rho_0, \rho_1)^2, \quad \lambda_b := \alpha \tilde{\lambda} - \Lambda_1. \quad (2.31)$$

As a result, the functional $G_b : \mathcal{P}_2(\mathbb{R}^d) \rightarrow [-\infty, +\infty]$ is uniformly displacement concave with constant $\lambda_b > 0$.

Proof. Consider any $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$. Then any $\mathcal{W}_2$ -geodesic $(\rho_s)_{s \in [0,1]}$ connecting ρ_0 with ρ_1 solves the following system of geodesic equations:

$$\begin{cases} \partial_s \rho_s + \operatorname{div}(\rho_s v_s) = 0, \\ \partial_s(\rho_s v_s) + \operatorname{div}(\rho_s v_s \otimes v_s) = 0, \end{cases} \quad (2.32)$$

where $\rho_s : \mathbb{R}^d \rightarrow \mathbb{R}$ and $v_s : \mathbb{R}^d \rightarrow \mathbb{R}^d$. The first derivative of G_b along geodesics can be computed explicitly as

$$\begin{aligned} \frac{d}{ds} G_b(\rho_s) &= \int \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz - \frac{d}{ds} \mathcal{R}(\rho_s) \\ &\quad + \left\langle \left[\int \nabla_x f_1(z, x) d\rho_s(z) + \int \nabla_x f_2(z, x) d\pi(z) + \kappa(x - x_0) \right] \Big|_{x=b(\rho_s)}, \frac{d}{ds} b(\rho_s) \right\rangle. \end{aligned}$$

For the last term, the left-hand side of the inner product is zero by definition of the best response $b(\rho_s)$ to ρ_s , see (2.30). Therefore

$$\frac{d}{ds} G_b(\rho_s) = \int \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz - \frac{d}{ds} \mathcal{R}(\rho_s).$$

Differentiating a second time, using (2.32) and integration by parts, we obtain

$$\frac{d^2}{ds^2} G_b(\rho_s) = L_1(\rho_s) + L_2(\rho_s) - \frac{d^2}{ds^2} \mathcal{R}(\rho_s),$$

where

$$\begin{aligned} L_1(\rho_s) &:= \int \nabla_z^2 f_1(z, b(\rho_s)) \cdot (v_s \otimes v_s) \rho_s dz = \int \langle v_s, \nabla_z^2 f_1(z, b(\rho_s)) \cdot v_s \rangle \rho_s dz, \\ L_2(\rho_s) &:= \int \frac{d}{ds} b(\rho_s) \cdot \nabla_{xz}^2 f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz. \end{aligned}$$

From (2.13), we have that

$$\frac{d^2}{ds^2} \mathcal{R}(\rho_s) \geq \alpha \tilde{\lambda} \mathcal{W}_2(\rho_0, \rho_1)^2,$$

and thanks to Assumption 2.9(ii) it follows that

$$L_1(s) \leq \Lambda_1 \mathcal{W}_2(\rho_0, \rho_1)^2.$$

This leaves L_2 to bound; we first consider the term $\frac{d}{ds} b(\rho_s)$:

$$\begin{aligned} \frac{d}{ds} b(\rho_s) &= \int \delta_\rho b[\rho_s](\tilde{z}) \partial_s \rho_s(\tilde{z}) d\tilde{z} = - \int \delta_\rho b[\rho_s](\tilde{z}) \operatorname{div}(\rho_s v_s)(\tilde{z}) d\tilde{z} \\ &= \int \nabla_z \delta_\rho b[\rho_s](\tilde{z}) \cdot v_s(\tilde{z}) d\rho_s(\tilde{z}). \end{aligned}$$

Defining $u(\rho_s) \in \mathbb{R}^d$ by

$$u(\rho_s) := \int \nabla_{xz}^2 f_1(z, b(\rho_s)) \cdot v_s(z) d\rho_s(z),$$

using the results from Lemma 2.39 for $\nabla_z \delta_\rho b[\rho_s]$ and the fact that $Q(\rho)$ is constant in z and x , we have

$$\begin{aligned} L_2(\rho_s) &= - \iint [Q(\rho_s)^{-1} \nabla_{xz}^2 f_1(\tilde{z}, b(\rho_s)) \cdot v_s(\tilde{z})] \cdot \nabla_{xz}^2 f_1(z, b(\rho_s)) \cdot v_s(z) d\rho_s(z) d\rho_s(\tilde{z}) \\ &= - \langle u(\rho_s), Q(\rho_s)^{-1} u(\rho_s) \rangle \leq 0 \end{aligned}$$

Combining all terms together, we obtain

$$\frac{d^2}{ds^2} G_b(\rho_s) \leq -(\alpha \tilde{\lambda} - \Lambda_1) \mathcal{W}_2(\rho_0, \rho_1)^2.$$

□

Remark 2.43. Under some additional assumptions on the functions f_1 and f_2 , we can obtain an improved convergence rate. In particular, assume that for all $z, x \in \mathbb{R}^d$,

- there exists a constant $\Lambda_{1u} \geq \lambda_1$ such that $\nabla_x^2 f_1(z, x) \preceq \Lambda_{1u} \mathbf{I}_d$;
- there exists a constant $\Lambda_{2u} \geq \lambda_2$ such that $\nabla_x^2 f_2(z, x) \preceq \Lambda_{2u} \mathbf{I}_d$;
- there exists a constant $\sigma \geq 0$ such that $\|\nabla_{xz}^2 f_1(z, x)\|_2 \geq \sigma$.

Then we have $-Q(\rho_s)^{-1} \preceq -1/(\kappa + \Lambda_{1u} + \Lambda_{2u}) \mathbf{I}_d$. Using Lemma 2.39, we then obtain a stronger bound on L_2 as follows:

$$\begin{aligned} L_2(\rho_s) &\leq -\frac{1}{\kappa + \Lambda_{1u} + \Lambda_{2u}} \|u(\rho_s)\|^2 \\ &\leq -\frac{1}{\kappa + \Lambda_{1u} + \Lambda_{2u}} \int \|\nabla_{xz}^2 f_1(z, b(\rho_s))\|_2^2 d\rho_s(z) \int \|v_s(z)\|^2 d\rho_s(z) \\ &\leq -\frac{\sigma^2}{\kappa + \Lambda_{1u} + \Lambda_{2u}} \mathcal{W}_2(\rho_0, \rho_1)^2. \end{aligned}$$

This means we can improve the convergence rate in (2.31) to $\lambda_b := \alpha \tilde{\lambda} + \frac{\sigma^2}{\kappa + \Lambda_{1u} + \Lambda_{2u}} - \Lambda_1$.

Lemma 2.44 (Uniform boundedness of the best response). *Let Assumption 2.18 hold. Then for any $\rho \in \mathcal{P}(\mathbb{R}^d)$, we have*

$$\|b(\rho)\|^2 \leq \|x_0\|^2 + \frac{2(a_1 + a_2)}{\kappa}.$$

Proof. The bound trivially holds if $\rho \in \mathcal{P}(\mathbb{R}^d)$ is such that $G(\rho, x) = \pm\infty$ for some $x \in \mathbb{R}^d$ since then we have $b(\rho) = 0$. Else, by definition of the best response $b(\rho)$, we have

$$\int \nabla_x f_1(z, b(\rho)) d\rho(z) + \int \nabla_x f_2(z, b(\rho)) d\pi(z) + \kappa(b(\rho) - x_0) = 0.$$

To show that $b(\rho)$ is uniformly bounded, we take the inner product of the above expression with $b(\rho)$ itself

$$\kappa \|b(\rho)\|^2 = \kappa x_0 \cdot b(\rho) - \int \nabla_x f_1(z, b(\rho)) \cdot b(\rho) d\rho(z) - \int \nabla_x f_2(z, b(\rho)) \cdot b(\rho) d\pi(z).$$

Using Assumption 2.18 to bound the two integrals, together with using Young's inequality to bound the first term on the right-hand side, we obtain

$$\kappa \|b(\rho)\|^2 \leq \frac{\kappa}{2} \|x_0\|^2 + \frac{\kappa}{2} \|b(\rho)\|^2 + a_1 + a_2.$$

□

Lemma 2.45 (Upper semi-continuity). *Let Assumption 2.18 hold. The functional $G : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, +\infty]$ is upper semi-continuous when $\mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d$ is endowed with the product topology of the weak topology and the Euclidean topology. Moreover, the functional $G_b : \mathcal{P}(\mathbb{R}^d) \rightarrow [-\infty, +\infty]$ is upper semi-continuous with respect to the weak topology.*

Proof. The functional $G : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, +\infty]$ is continuous in the second variable thanks to Assumption 2.9(ii). Similarly, $\int f_1(z, x) d\rho(z) + \int f_2(z, x) d\pi(z)$ is continuous in ρ thanks to [194, Proposition 7.1] using the continuity of f_1 and f_2 . Further, $-\mathcal{R}$ is upper semi-continuous using Lemma 2.64 and [194, Proposition 7.2] thanks to Assumptions 2.10 and 2.11. This concludes the continuity properties for G .

The upper semi-continuity of G_b then follows from a direct application of a version of Berge's maximum theorem [5, Lemma 16.30]. Let $R := \|x_0\|^2 + \frac{2(a_1+a_2)}{\kappa} > 0$. We define $\varphi : (\mathcal{P}(\mathbb{R}^d), \mathcal{W}_2) \rightarrow \mathbb{R}^d$ as the correspondence that maps any $\rho \in \mathcal{P}(\mathbb{R}^d)$ to the closed ball $\overline{B_R(0)} \subset \mathbb{R}^d$. Then the graph of φ is $\text{Gr } \varphi = \mathcal{P}(\mathbb{R}^d) \times \{\overline{B_R(0)}\}$. With this definition of φ , the range of φ is compact and φ is continuous with respect to weak convergence, and so it is in particular upper hemicontinuous. Thanks to Lemma 2.44, the best response function $b(\rho)$ is always contained in $\overline{B_R(0)}$ for any choice of $\rho \in \mathcal{P}(\mathbb{R}^d)$. As a result, maximizing $-G(\rho, x)$ in x over $\mathbb{R}^d$ for a fixed $\rho \in \mathcal{P}(\mathbb{R}^d)$ reduces to maximizing it over $\overline{B_R(0)}$. Using the notation introduced above, we can restrict G to $G : \text{Gr } \varphi \rightarrow \mathbb{R}$ and write

$$G_b(\rho) := \max_{\hat{x} \in \varphi(\rho)} -G(\rho, \hat{x}).$$

Because $G(\rho, x)$ is upper semi-continuous when $\mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d$ is endowed with the product topology of the weak topology and the Euclidean topology, [5, Lemma 16.30] guarantees that $G_b(\cdot)$ is upper semi-continuous in the weak topology. $\square$

Proposition 2.46 (Ground state). *Let Assumption 2.18 hold. There exists a unique maximizer ρ_* for the functional G_b over $\mathcal{P}(\mathbb{R}^d)$, it satisfies $\rho_* \in \mathcal{P}_2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ and is absolutely continuous with respect to $\tilde{\rho}$.*

Proof. Uniqueness of the maximizer (if it exists) is guaranteed by the uniform concavity provided by Lemma 2.42. To show existence of a maximizer, we use the direct method in the calculus of variations, requiring the following key properties for G_b : (1) boundedness from above, (2) upper semi-continuity, and (3) tightness of any maximizing sequence. To show (1), note that $\nabla_z^2(f_1(z, x) + \alpha \log \tilde{\rho}(z)) \preceq -(\alpha \tilde{\lambda} - \Lambda_1) \text{Id}$ for all $z, x \in \mathbb{R}^d \times \mathbb{R}^d$ by Assumptions 2.9(ii) and 2.10, and so

$$f_1(z, x) + \alpha \log \tilde{\rho}(z) \leq c_0(x) - \frac{(\alpha \tilde{\lambda} - \Lambda_1)}{4} \|z\|^2 \quad \forall (z, x) \in \mathbb{R}^d \times \mathbb{R}^d \quad (2.33)$$

with $c_0(x) := f_1(0, x) + \alpha \log \tilde{\rho}(0) + \frac{1}{\alpha \tilde{\lambda} - \Lambda_1} \|\nabla_z [f_1(0, x) + \alpha \log \tilde{\rho}(0)]\|^2$. Therefore,

$$\begin{aligned} G_b(\rho) &= \int [f_1(z, b(\rho)) + \alpha \log \tilde{\rho}(z)] d\rho(z) + \int f_2(z, b(\rho)) d\pi(z) + \frac{\kappa}{2} \|b(\rho) - x_0\|^2 \\ &\quad - \alpha \int \rho \log \rho - \frac{1}{2} \int \rho W * \rho \\ &\leq c_0(b(\rho)) + \int f_2(z, b(\rho)) d\pi(z) + \frac{\kappa}{2} \|b(\rho) - x_0\|^2. \end{aligned}$$

To estimate each of the remaining terms on the right-hand side, denote $R := \|x_0\|^2 + \frac{2(a_1+a_2)}{\kappa}$ and recall that $\|b(\rho)\| \leq R$ for any $\rho \in \mathcal{P}(\mathbb{R}^d)$ thanks to Lemma 2.44. By continuity of f_1 and $\log \tilde{\rho}$, there exists a constant $c_1 \in \mathbb{R}$ such that

$$\sup_{x \in B_R(0)} c_0(x) = \sup_{x \in B_R(0)} \left[f_1(0, x) + \alpha \log \tilde{\rho}(0) + \frac{1}{\alpha \tilde{\lambda} - \Lambda_1} \|\nabla_z (f_1(0, x) + \alpha \log \tilde{\rho}(0))\|^2 \right] \leq c_1. \quad (2.34)$$

The term $\int f_2(z, b(\rho)) d\pi(z)$ is controlled by c_2 thanks to (2.10). The third term can be bounded directly to obtain

$$G_b(\rho) \leq c_1 + c_2 + \kappa(R^2 + \|x_0\|^2).$$

This concludes the proof of (1). Statement (2) was shown in Lemma 2.45 using Assumption 2.18. Then we obtain a maximizing sequence $(\rho_n) \in \mathcal{P}(\mathbb{R}^d)$ which is in the closed unit ball of $C_0(\mathbb{R}^d)^*$ and so the Banach-Alaoglu theorem [192, Theorem 3.15] there exists a limit ρ_* in the Radon

measures and a subsequence (not relabeled) such that $\rho_n \xrightarrow{*} \rho_*$. In fact, ρ_* is absolutely continuous with respect to $\tilde{\rho}$ as otherwise $G_b(\rho_*) = -\infty$, which contradicts that $G_b(\cdot) > -\infty$ somewhere. We conclude that $\rho_* \in L^1(\mathbb{R}^d)$ since $\tilde{\rho} \in L^1(\mathbb{R}^d)$. To ensure $\rho_* \in \mathcal{P}(\mathbb{R}^d)$, we require (3) tightness of the minimizing sequence (ρ_n) . By Markov's inequality [117] it is sufficient to establish a uniform bound on the second moments:

$$\int \|z\|^2 d\rho_n(z) < C \quad \forall n \in \mathbb{N}. \quad (2.35)$$

To see this we proceed in a similar way as in the proof of Proposition 2.25. Defining

$$K(\rho) := - \int [f_1(z, b(\rho)) + \alpha \log \tilde{\rho}(z)] d\rho(z) + \alpha \int \rho \log \rho dz + \frac{1}{2} \int \rho W * \rho dz,$$

we have $K(\rho) = -G_b(\rho) + \int f_2(z, b(\rho)) d\pi(z) + \frac{\kappa}{2} \|b(\rho) - x_0\|^2$. Then using again the bound on $b(\rho)$ from Lemma 2.44,

$$\begin{aligned} K(\rho) &\leq -G_b(\rho) + \sup_{x \in B_R(0)} \int f_2(z, x) d\pi(z) + \kappa (R^2 + \|x_0\|^2) \\ &\leq -G_b(\rho) + c_2 + \kappa (R^2 + \|x_0\|^2), \end{aligned}$$

where the last inequality is thanks to (2.10). Hence, using the estimates (2.33) and (2.34) from above,

$$\frac{(\alpha\tilde{\lambda} - \Lambda_1)}{4} \int \|z\|^2 d\rho_n(z) \leq c_1 - \int [f_1(z, b(\rho_n)) + \alpha \log \tilde{\rho}(z)] d\rho_n(z).$$

Applying Lemma 2.63 with $\varepsilon = \frac{\alpha\tilde{\lambda} - \Lambda_1}{8\alpha}$ in the same fashion as in the proof of Proposition 2.25 and noting that the sequence (ρ_n) is minimizing $(-G_b)$,

$$\begin{aligned} \frac{(\alpha\tilde{\lambda} - \Lambda_1)}{4} \int \|z\|^2 d\rho_n(z) &\leq c_1 + K(\rho_n) + \alpha\varepsilon \int \|z\|^2 d\rho_n(z) + \alpha c_\varepsilon \\ \Rightarrow \frac{(\alpha\tilde{\lambda} - \Lambda_1)}{8} \int \|z\|^2 d\rho_n(z) &\leq c_1 - G_b(\rho_n) + c_2 + \kappa (R^2 + \|x_0\|^2) + \alpha c_\varepsilon \\ &\leq c_1 - G_b(\rho_1) + c_2 + \kappa (R^2 + \|x_0\|^2) + \alpha c_\varepsilon < \infty, \end{aligned}$$

which uniformly bounds the second moments of (ρ_n) . This concludes the proof for the estimate (2.35) and also ensures that $\rho_* \in \mathcal{P}_2(\mathbb{R}^d)$. $\square$

Corollary 2.47. *Any maximizer ρ_* of G_b is a steady state for equation (2.8), and satisfies $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho}) = \mathbb{R}^d$ and $\rho_* \in C^2(\mathbb{R}^d)$.*

Proof. To show that ρ_* is a steady state according to Lemma 2.38 we can follow exactly the same argument as in the proof of Proposition 2.36, just replacing $\int f_1(z, x) \mu_*(x)$ with $f_1(z, b(\rho_*))$.

It remains to show that $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$. As ρ^* is a maximizer, it is in particular a critical point, and therefore satisfies that $\delta_\rho G_b[\rho_*](z)$ is constant on all connected components of $\text{supp}(\rho_*)$. Thanks to Proposition 2.41, this means there exists a constant $c[\rho_*]$ (which may be different on different components of $\text{supp}(\rho_*)$) such that

$$f_1(z, b(\rho_*)) - \alpha \log \left(\frac{\rho_*(z)}{\tilde{\rho}(z)} \right) - W * \rho_*(z) = c[\rho_*] \quad \text{on } \text{supp}(\rho_*) .$$

Rearranging, we obtain (for a possibly different constant $c[\rho_*] \neq 0$)

$$\rho_*(z) = c[\rho_*] \tilde{\rho}(z) \exp \left[\frac{1}{\alpha} (f_1(z, b(\rho_*)) - W * \rho_*(z)) \right] \quad \text{on } \text{supp}(\rho_*) . \quad (2.36)$$

Firstly, $\text{supp}(\rho_*) \subset \text{supp}(\tilde{\rho})$ since ρ_* is absolutely continuous with respect to $\tilde{\rho}$. Secondly, note that $\exp \frac{1}{\alpha} f_1(z, b(\rho_*)) > 0$ for all $z \in \mathbb{R}^d$ since $f_1 \in \mathbb{R}$. Finally, we claim that $\exp \left(-\frac{1}{\alpha} W * \rho_*(z) \right) > 0$ for ρ_* -a.e. $z \in \mathbb{R}^d$. In other words, we claim that $W * \rho_*(z) < \infty$ for ρ_* -a.e. $z \in \mathbb{R}^d$. Else, if $W * \rho_*(z) = \infty$ on a set of non-zero measure, then $\int \rho_* W * \rho_* = +\infty$. Since all other terms are upper bounded due to the x -player minimization, this implies that $G_b(\rho_*) = -\infty$. Since ρ^* is a maximizer of G_b , this would mean $G_b \equiv -\infty$, which is a contradiction. We conclude that $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$. Finally, $\rho_* \in C^2(\mathbb{R}^d)$ thanks to $\tilde{\rho}, f_1, W \in C^2(\mathbb{R}^d)$. $\square$

Remark 2.48. Note that $\tilde{\rho} \in L^\infty(\mathbb{R}^d)$ by Assumption 2.10. If we have in addition that $f_1(\cdot, x) \in L^\infty(\mathbb{R}^d)$ for all $x \in \mathbb{R}^d$, then the maximizer ρ_* of G_b is in $L^\infty(\mathbb{R}^d)$ as well. This follows directly by bounding the right-hand side of (2.36).

With the notion of steady state given in Lemma 2.38, we can obtain improved regularity for any steady state ρ_∞ under mild additional assumptions.

Lemma 2.49. Any steady state ρ_∞ for equation (2.8) has continuous Lebesgue density, $\rho_\infty \in C(\mathbb{R}^d)$.

Proof. Thanks to our assumptions, we have $f_1(\cdot, b(\rho_\infty)) + \alpha \log \tilde{\rho}(\cdot) \in C^1$, which implies that $\nabla_z(f_1(\cdot, b(\rho_\infty)) + \alpha \log \tilde{\rho}(\cdot)) \in L_{loc}^\infty$. By the definition of a steady state, $\rho_\infty \in L^1 \cap L_{loc}^\infty$ and $\nabla W * \rho_\infty \in L_{loc}^\infty$. Let

$$h(z) := \rho_\infty(z) \nabla_z [f_1(z, b(\rho_\infty)) + \alpha \log \tilde{\rho}(z) - (W * \rho_\infty)(z)] .$$

Then by the aforementioned regularity, we obtain $h \in L_{loc}^1 \cap L_{loc}^\infty$. By interpolation, it follows that $h \in L_{loc}^p$ for all $1 < p < \infty$. This implies that $\text{div}(h) \in W_{loc}^{-1,p}$. Since ρ_∞ is a weak $W_{loc}^{1,2}$ -solution of (2.25), we have

$$\Delta \rho_\infty = \text{div}(h) ,$$

and so by classic elliptic regularity theory we conclude $\rho_\infty \in W_{loc}^{1,p}$. Finally, applying Morrey's inequality, we have $\rho_\infty \in C^{0,k}$ where $k = \frac{p-d}{p}$ for any $d < p < \infty$. Therefore $\rho_\infty \in C(\mathbb{R}^d)$ (after possibly being redefined on a set of measure zero). $\square$

With the above preliminary results, we can now show the HWI inequality, which implies again a Talagrand-type inequality and a generalized logarithmic Sobolev inequality.

Proposition 2.50 (HWI inequalities). *Define the dissipation functional*

$$D_b(\gamma) := \iint \|\nabla_z \delta_\rho G_b[\rho](z)\|^2 d\rho(z).$$

Let Assumption 2.18 hold. Denote by ρ_* the unique maximizer of G_b .

(HWI) *Let $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho_0), G_b(\rho_1), D_b(\rho_0) < \infty$. Then*

$$G_b(\rho_0) - G_b(\rho_1) \leq \mathcal{W}_2(\rho_0, \rho_1) \sqrt{D_b(\rho_0)} - \frac{\lambda_b}{2} \mathcal{W}_2(\rho_0, \rho_1)^2 \quad (2.37)$$

(LogSobolev) *Any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho), D_b(\rho) < \infty$ satisfies*

$$D_b(\rho) \geq 2\lambda_b G_b(\rho | \rho_*). \quad (2.38)$$

(Talagrand) *For any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho) < \infty$, we have*

$$\mathcal{W}_2(\rho, \rho_*)^2 \leq \frac{2}{\lambda_b} G_b(\rho | \rho_*). \quad (2.39)$$

Proof. The proof for this result follows analogously to the arguments presented in the proofs of Proposition 2.27, Corollary 2.29 and Corollary 2.30, using the preliminary results established in Proposition 2.42 and Proposition 2.46. $\square$

Proof of Theorem 2.19. Following the same approach as in the proof of Theorem 2.13, the results in Theorem 2.19 immediately follow by combining Proposition 2.46, Corollary 2.47 and Proposition 2.50 applied to solutions of the PDE (2.8). $\square$

2.B Competitive Objective, Fast Population (Proof of Theorem 2.20)

In this section, let Assumptions 2.9(ii), 2.10 and 2.11 hold throughout with $\alpha\tilde{\lambda} > \Lambda_1$. The proof for Theorem 2.20 uses similar strategies as that of Theorem 2.19, but considers the evolution of an ODE rather than a PDE. Recall that for any $x \in \mathbb{R}^d$ the best response $r(x)(\cdot) \in \mathcal{P}(\mathbb{R}^d)$ in (2.9) is defined as

$$r(x) := \operatorname{argmax}_{\hat{\rho} \in \mathcal{P}} G(\hat{\rho}, x),$$

where the energy $G(\rho, x) : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, \infty]$ is given as in Appendix 2.A by

$$G(\rho, x) = \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2 - \alpha KL(\rho \mid \tilde{\rho}) - \frac{1}{2} \int (W * \rho)(z) d\rho(z).$$

We start by showing that the best response $r(x)$ exists and is uniquely defined for each $x \in \mathbb{R}^d$ (Proposition 2.51). In order to show convergence to equilibrium for the algorithm dynamics $x(t)$ (Theorem 2.20), we need to differentiate the energy $G_d(x) := G(\rho, x)|_{\rho=r(x)}$ in x , which is a non-trivial task given that it is not obvious that the best response $r(x)$ is even differentiable in x . Our goal is therefore to achieve a Danskin-like result of the form

$$\nabla_x G_d(x) = (\nabla_x G(\rho, x))|_{\rho=r(x)}$$

(Proposition 2.60). Danskin's theorem, also known as the envelope theorem, is a classical result in (Euclidean) game theory. In zero-sum games, the derivative of a cost function through the best response of the other player is equal to the derivative with the best response plugged in after differentiating, and is used to prove convexity through the implicitly-defined best response function; see [18] for details. To achieve this result, we will need to use that the best response $r(x)$ is continuous in x in $\mathcal{W}_2$ (Corollary 2.58). For the rest of this section, we use the following notation: denote $m(z, x) := -f_1(z, x) - \alpha \log \tilde{\rho}(z)$ and

$$F(\rho, x) := \alpha \int \rho(z) \log \rho(z) dz + \int m(z, x) d\rho(z) + \frac{1}{2} \int (W * \rho)(z) d\rho(z).$$

Then $G(\rho, x) = \int f_2(z, x) d\pi(z) + \frac{\kappa}{2} \|x - x_0\|^2 - F(\rho, x)$, and so maximizing $G(\rho, x)$ in ρ corresponds to minimizing $F(\rho, x)$ in ρ . Showing that the best response $r(x)$ is continuous means that for any sequence $(x_n)_{n \in \mathbb{N}}$ in $\mathbb{R}^d$ converging to $\bar{x}$ as $n \rightarrow \infty$, we have $\mathcal{W}_2(r(x_n), r(\bar{x})) \rightarrow 0$. For a sequence $x_n \rightarrow \bar{x}$, define

$$F_n(\rho) := F(\rho, x_n), \quad \bar{F}(\rho) := F(\rho, \bar{x}), \quad \rho_n := \operatorname{argmin}_{\rho \in \mathcal{P}_2(\mathbb{R}^d)} F_n(\rho), \quad m_n(z) := m(z, x_n). \quad (2.40)$$

To obtain the continuity of $r(x)$, we first show that the sequence $(\rho_n)_{n \in \mathbb{N}}$ has uniformly bounded second moments (obtained from convergence of the second moments Proposition 2.57) and that $F_n \xrightarrow{\Gamma} \bar{F}$ in $\mathcal{W}_2$ (Proposition 2.55). For these results, we will need an extension of Fatou's Lemma (Theorem 2.53) which gives conditions under which limits and integrals can be exchanged when not only the integrand but also the measure of integration depend on $n \rightarrow \infty$. The notion of *asymptotically uniformly integrable (a.u.i)* (Definition 2.52) relaxes the notion of lower-boundedness of the integrand, appears as a condition in the extension of Fatou's Lemma (Theorem 2.53), and is a key ingredient to show the Γ -convergence result in Proposition 2.55.

Proposition 2.51. *For each $x \in \mathbb{R}^d$ there exists a unique maximizer $\rho_* := r(x)$ solving $\operatorname{argmax}_{\hat{\rho} \in \mathcal{P}} G(\hat{\rho}, x)$. Further, $r(x) \in L^1(\mathbb{R}^d) \cap \mathcal{P}_2(\mathbb{R}^d)$, $\operatorname{supp}(r(x)) = \operatorname{supp}(\tilde{\rho})$, and there exists a function $c : \mathbb{R}^d \rightarrow \mathbb{R}$ such that the best response $\rho_*(z) = r(x)(z)$ solves the Euler-Lagrange equation*

$$\delta_\rho G[\rho_*, x](z) := \alpha \log \rho_*(z) + m(z, x) + (W * \rho_*)(z) = c(x) \quad (2.41)$$

for all $(z, x) \in \operatorname{supp}(\tilde{\rho}) \times \mathbb{R}^d$.

Proof. Equivalently, consider the minimization problem for $F(\rho, x)$ for some fixed x . Note that $m(z, x)$ is strictly convex in z for fixed x by Assumptions 2.9(ii) and 2.10. Together with Assumption 2.11, we can directly apply the uniqueness and existence result from [50, Theorem 2.1 (i)].

The result on the support of $r(x)$ and the expression for the Euler-Lagrange equation follows by the same arguments as in Corollary 2.26 and Corollary 2.47, using Assumption 2.9(ii) and that $\alpha\tilde{\lambda} > \Lambda_1$. The proof that $r(x) \in L_1(\mathbb{R}^d) \cap \mathcal{P}_2(\mathbb{R}^d)$ uses the same arguments as those in Proposition 2.46, where instead of using $x = b(\rho)$, x is some fixed value. $\square$

Definition 2.52 (asymptotically uniformly integrable, [107]). *A sequence of measurable $(\mathbb{R} \cup \{\pm\infty\})$ -valued functions $(f_n)_{n \in \mathbb{N}}$ is called asymptotically uniformly integrable (a.u.i) with respect to a sequence of measures $(\mu_n)_{n \in \mathbb{N}} \subset \mathcal{P}(\mathbb{R}^d)$ if*

$$\lim_{R \rightarrow +\infty} \limsup_{n \in \mathbb{N}} \int_{\mathbb{R}^d} |f_n(z)| \mathbb{I}\{z \in \mathbb{R}^d : |f_n(z)| \geq R\} d\mu_n(z) = 0.$$

Theorem 2.53 (Theorem 2.4 in [107], Fatou's Lemma for Weakly Converging Measures). *Let $\mathbf{S}$ be a metric space, $(\mu_n)_{n \in \mathbb{N}}$ be a sequence of measures on $\mathbf{S}$ that converges narrowly to $\mu \in \mathcal{M}(\mathbf{S})$, and let $(f_n)_{n \in \mathbb{N}}$ be a sequence of measurable $(\mathbb{R} \cup \{\pm\infty\})$ -valued functions on $\mathbf{S}$ such that $\lim_{n \rightarrow \infty, s' \rightarrow s} f_n(s')$ exists for μ -a.e. $s \in \mathbf{S}$. Denote $f_n^-(z) = -\min\{f_n(z), 0\}$. Let $(f_n^-)_{n \in \mathbb{N}}$ be a.u.i with respect to $(\mu_n)_{n \in \mathbb{N}}$. Then*

$$\liminf_{n \rightarrow \infty} \int_{\mathbf{S}} f_n(s) \mu_n(s) (ds) \leq \int_{\mathbf{S}} \liminf_{n \rightarrow \infty, s' \rightarrow s} f_n(s') \mu(ds).$$

Remark 2.54. *If instead the condition on f_n is strengthened to $(f_n)_{n \in \mathbb{N}}$ being a.u.i, then the limit bound holds with equality [107, Corollary 2.8]; this result will also be utilized.*

Proposition 2.55 (Γ -Convergence of Energy). *Let Assumption 2.15(b) hold. For any sequence (x_n) converging to some limit $\bar{x}$ in $\mathbb{R}^d$, we have $F_n \xrightarrow{\Gamma} \bar{F}$ as $n \rightarrow \infty$ in $\mathcal{W}_2$ for sequences $(\mu_n) \in \mathcal{P}_2^{ac}$.*

Proof. We write

$$F_n(\rho) = H_0(\rho) + H_n(\rho)$$

where we define

$$H_0(\rho) = \alpha \int \rho \log \rho + \frac{1}{2} \int \rho W * \rho, \quad H_n(\rho) = \int m_n(z) d\rho(z).$$

where $m_n(z) = m(z, x_n)$. Convergence in $\mathcal{W}_2$ is equivalent to narrow convergence $\mu_n \xrightarrow{*} \bar{\mu}$ in $\mathcal{P}(\mathbb{R}^d)$ together with convergence of second moments $\int \|z\|^2 d\mu_n(z) \rightarrow \int \|z\|^2 d\bar{\mu}(z)$. This is equivalent to $\int \varphi(z) d\mu_n(z) \rightarrow \int \varphi(z) d\bar{\mu}(z)$ for all φ with at most quadratic growth, and implies that μ_n also converges weakly. First, we will show that for every $(\mu_n) \in \mathcal{P}_2^{ac}$ converging to $\bar{\mu}$ in $\mathcal{W}_2$, we have

$$\bar{F}(\bar{\mu}) \leq \liminf_{n \rightarrow \infty} F_n(\mu_n).$$

We accomplish this by considering individual terms, noting that

$$\liminf_{n \rightarrow \infty} F_n(\mu_n) \geq \liminf_{n \rightarrow \infty} H_0(\mu_n) + \liminf_{n \rightarrow \infty} H_n(\mu_n),$$

see [30, Chapter 1, (1.2)], and thus lower bounds for each term sufficiently lower-bound $\liminf_{n \rightarrow \infty} F_n(\mu_n)$. By Lemma 2.64,

$$\liminf_{n \rightarrow \infty} \int \mu_n \log \mu_n \geq \int \bar{\mu} \log \bar{\mu}.$$

The interaction kernel term can be written as $\iint W(z - z') d(\mu_n \times \mu_n)(z, z')$. Since $W \geq 0$ and is lower-semicontinuous due to Assumption 2.11, applying [202, Portmanteau Theorem 1.3.4 (iv)] gives that

$$\liminf_{n \rightarrow \infty} \int W(z - z') d(\mu_n \times \mu_n)(z, z') \geq \int W(z - z') d(\bar{\mu} \times \bar{\mu})(z, z').$$

This concludes the proof that for any sequence $(\mu_n) \in \mathcal{P}_2^{ac}$ such that $\mathcal{W}_2(\mu_n, \bar{\mu}) \rightarrow 0$, we have

$$\liminf_{n \rightarrow \infty} H_0(\mu_n) \geq H_0(\bar{\mu}).$$

For the second term $H_n(\mu_n)$, define

$$\bar{H}(\bar{\mu}) := - \int m(z, \bar{x}) d\bar{\mu}(z).$$

We now show that $\liminf_{n \rightarrow \infty} H_n(\mu_n) \geq \bar{H}(\bar{\mu})$ via Fatou's Lemma for narrowly converging measures, Theorem 2.53. Note that $\nabla_z^2 m(z, x) \succeq \lambda_b I_d$ for all $x, z \in \mathbb{R}^d$, which implies $\|\nabla_z^2 m(z, x)\|_2 \geq \lambda_b$ for all $x, z \in \mathbb{R}^d$. We also have from Assumption 2.15(b) that $\|\nabla_{xz}^2 m(z, x)\|_2 \leq L$. Hence it follows from Lemma 2.66 that $\{m_n^-\}_{n \in \mathbb{N}}$ is asymptotically uniformly integrable with respect to (μ_n) . This allows us to apply Theorem 2.53, which gives us that

$$\liminf_{n \rightarrow \infty} H_n(\mu_n) = \liminf_{n \rightarrow \infty} \int m_n(z) d\mu_n(z) \geq \int \liminf_{n \rightarrow \infty, z' \rightarrow z} m_n(z') \bar{\mu}(z) dz = \bar{H}(\bar{\mu}).$$

Secondly, we must show that for any measure $\bar{\mu}$ there exists a recovery sequence $(\mu_n) \in \mathcal{P}_2^{ac}$ with μ_n converging in $\mathcal{W}_2$ to $\bar{\mu}$, such that

$$\bar{F}(\bar{\mu}) \geq \limsup_{n \rightarrow \infty} F_n(\mu_n).$$

We select the constant sequence $\mu_n = \bar{\mu}$, and evaluate

$$\limsup_{n \rightarrow \infty} F_n(\bar{\mu}) = \limsup_{n \rightarrow \infty} (H_0(\bar{\mu}) + H_n(\bar{\mu})) \leq H_0(\bar{\mu}) + \int (\limsup_{n \rightarrow \infty} m_n(z)) d\bar{\mu}(z)$$

using Fatou's Lemma with the uniform upper bound on $m_n(z)$ from Lemma 2.65. We conclude that

$$\limsup_{n \rightarrow \infty} F_n(\bar{\mu}) \leq H_0(\bar{\mu}) + \bar{H}(\bar{\mu}).$$

□

Proposition 2.56 (Uniform Second Moment Bound). *Given a sequence $x_n \rightarrow \bar{x}$ in $\mathbb{R}^d$, the sequence (ρ_n) as defined in (2.40) has uniformly bounded second moments.*

Proof. Applying Lemma 2.65 (1) with $m(z, x) := -f_1(z, x) - \alpha \log \tilde{\rho}(z)$, we have $m_n(z) := m(z, x_n) \geq -\hat{c}_1 + \hat{c}_2 \|z\|^2$ for all $n \in \mathbb{N}$ with $\hat{c}_2 > 0$ since $\lambda_b > 0$. Integrating the inequality with respect to ρ_n results in

$$\begin{aligned} \hat{c}_2 \int \|z\|^2 d\rho_n(z) &\leq \hat{c}_1 + \int m_n(z) d\rho_n(z) \\ &\leq \hat{c}_1 + \int m_n(z) d\rho_n(z) + \alpha \int \rho_n \log \rho_n + \frac{\hat{c}_2}{2} \int \|z\|^2 d\rho_n(z) + \hat{c}_3 + \frac{1}{2} \int \rho_n W * \rho_n \end{aligned}$$

using Lemma 2.63 with $\varepsilon = \hat{c}_2/(2\alpha)$, denoting $\hat{c}_3 = \alpha c_\varepsilon$. Using (2.10), there exists $N > 0$ and $c_2 > 0$ such that $\int f_2(z, x_n) d\pi(z) \leq c_2$ for all $n \geq N$. Then for all $n \geq N$,

$$\begin{aligned} \frac{\hat{c}_2}{2} \int \|z\|^2 d\rho_n(z) &\leq C + \int m_n(z) d\rho_n(z) + \alpha \int \rho_n \log \rho_n + \frac{1}{2} \int \rho_n W * \rho_n \\ &\quad - \frac{\kappa}{2} \|x_n - x_0\|^2 - \int f_2(z, x_n) d\pi(z) \\ &= C - G(\rho_n, x_n) \end{aligned}$$

where $C := \hat{c}_1 + \hat{c}_3 + c_2 + \max_{n \geq 1} \frac{\kappa}{2} \|x_n - x_0\|^2$. Since ρ_n maximizes G , we have $G(\rho_n, x_n) \geq G(\tilde{\rho}, x_n)$ for any choice of $n \geq 0$. Then for all $n \geq N$,

$$\frac{\hat{c}_2}{2} \int \|z\|^2 d\rho_n(z) \leq C - G(\tilde{\rho}, x_n) \leq C + \max_x (-G(\tilde{\rho}, x)).$$

The right-hand side is finite and bounded by some constant independent of n since $G(\tilde{\rho}, \cdot)$ is twice continuously differentiable and strongly λ_d -convex over $\mathbb{R}^d$ (see proof of Lemma 2.39). □

Proposition 2.57 (Convergence of Second Moments). *Let Assumptions 2.15-2.17 hold. If $x_n \rightarrow \bar{x}$ in $\mathbb{R}^d$, then the sequence (ρ_n) as defined in (2.40) satisfies $\int \|z\|^2 d\rho_n(z) \rightarrow \int \|z\|^2 d\bar{\rho}(z)$.*

Proof. Note that $(\rho_n) \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$ by Proposition 2.51. We will show convergence of the second moments by showing that $(\|z\|^2 \rho_n(z))$ is equi-integrable, that is,

$$\lim_{K \rightarrow \infty} \sup_n \int \|z\|^2 \rho_n(z) \mathbb{1}_{\{z : \|z\|^2 \rho_n(z) \geq K\}} dz = 0.$$

Since ρ_n is a minimizer of $F(\cdot, x_n)$, ρ_n satisfies the re-arranged EL equation (2.41)

$$\rho_n(z) = \tilde{c}(x_n) \exp\left(-\frac{1}{\alpha} m(z, x_n) - \frac{1}{\alpha} W * \rho_n\right) \leq \tilde{c}(x_n) \exp\left(-\frac{1}{\alpha} m(z, x_n)\right),$$

where we used that $\exp(-W * \rho) \leq 1$ since $W * \rho \geq 0$. Using (1) from Lemma 2.65 which provides a lower bound on m , we obtain

$$\rho_n(z) \leq \tilde{c}(x_n) \exp\left(\frac{\hat{c}_1 - \hat{c}_2 \|z\|^2}{\alpha}\right)$$

with $\hat{c}_2 > 0$ since $\lambda_b > 0$. In order to obtain an upper-bound for ρ_n independent of x_n , we show that $\sup_n \tilde{c}(x_n) < \infty$. Starting again from the Euler-Lagrange equation, using that $\int \rho_n = 1$

$$\begin{aligned} \tilde{c}(x_n) \int \exp\left(-\frac{1}{\alpha} m(z, x_n) - \frac{1}{\alpha} W * \rho_n\right) dz &= 1 \\ \Rightarrow \tilde{c}(x_n) &= \left(\int \exp\left(-\frac{1}{\alpha} m(z, x_n) - \frac{1}{\alpha} W * \rho_n\right) dz \right)^{-1}. \end{aligned}$$

Next, we upper-bound $\tilde{c}(x_n)$. Using the upper estimate (2) for $m(z, x_n)$ from Lemma 2.65, which requires Assumptions 2.15 and 2.16, we obtain $m(z, x_n) \leq c_1 + c_2 \|z\|^2$ with $c_2 > 0$. To control the term with $W * \rho_n$, we use Assumption 2.17 to obtain a growth condition for W ,

$$\begin{aligned} W(z) &= W(0) + \nabla W(0) \cdot z + \int_0^1 z^\top \nabla^2 W(sz) z ds \leq W(0) + \|\nabla W(0)\| \|z\| + \Lambda_W \|z\|^2 \\ &\leq \tilde{c}_1 + \tilde{c}_2 \|z\|^2, \end{aligned}$$

and so

$$\begin{aligned} W * \rho_n(z) &= \int W(z - z') \rho_n(z') dz' \leq \tilde{c}_1 + \tilde{c}_2 \int \|z - z'\|^2 d\rho_n(z') \\ &\leq \tilde{c}_1 + 2\tilde{c}_2 \int \|z'\|^2 d\rho_n(z') + 2\tilde{c}_2 \|z\|^2 \leq \tilde{C}_1 + \tilde{C}_2 \|z\|^2 \end{aligned}$$

with $\tilde{C}_1, \tilde{C}_2 > 0$, using that the second moments of ρ_n are uniformly bounded by Proposition 2.56. Therefore,

$$\tilde{c}(x_n) \leq \left(\int \exp\left(-\frac{1}{\alpha} [c_1 + \tilde{C}_1 + (c_2 + \tilde{C}_2) \|z\|^2]\right) dz \right)^{-1} =: \tilde{c}_0$$

The upper-bound on ρ_n is therefore

$$\rho_n(z) \leq \tilde{c}_0 \exp\left(\hat{c}_1 - \hat{c}_2 \|z\|^2\right) \quad \text{for a.e. } z \in \mathbb{R}^d, \quad (2.42)$$

which shows that $\|z\|^2 \rho_n$ is equi-integrable since the Gaussian-like shape of the upper bound in (2.42) has a finite second moment. Applying [26, Theorem 4.5.4] results in the convergence of the second moment. $\square$

Corollary 2.58. *Let Assumptions 2.15-2.17 hold. The best response $r(x) \in \mathcal{P}_2(\mathbb{R}^d)$ is $\mathcal{W}_2$ -continuous in $x \in \mathbb{R}^d$, and $G_d : \mathbb{R}^d \rightarrow \mathbb{R}$ is continuous in $\mathbb{R}^d$.*

Proof. Given any sequence $(x_n)_{n \in \mathbb{N}}$ converging to $\bar{x}$ in $\mathbb{R}^d$ as $n \rightarrow \infty$, we have that $F_n \xrightarrow{\Gamma} \bar{F}$ in $\mathcal{W}_2$ for sequences such that $(\mu_n) \in \mathcal{P}_2^{ac}(\mathbb{R}^d)$ by Proposition 2.55. From Lemma 2.51, $r(x_n) = \arg \min_{\rho \in \mathcal{P}_2} F_n(\rho)$ exists and is unique for every x_n ; similarly, $r(\bar{x})$ is the unique minimizer of $\bar{F}$. Uniform boundedness of second moments for $(r(x_n))_{n \in \mathbb{N}}$ follows from Proposition 2.56, and tightness of $(r(x_n))_{n \in \mathbb{N}}$ follows using [9, Proposition 7.1.5]. Then precompactness of $(r(x_n))_{n \in \mathbb{N}}$ in the narrow topology follows from [9, Prokhorov's Theorem 5.1.3]. Together with convergence of second moments (Proposition 2.57), we conclude that $r(x_n)$ converges to $r(\bar{x})$ in $\mathcal{W}_2$ [29, Theorem 2.10]. In other words, the best response $r(x)$ is continuous in x in $\mathcal{W}_2$.

Recall that $G_d(x) := \max_{\rho \in \mathcal{P}} G(\rho, x)$. For any fixed ρ , $G(x, \rho)$ is continuous in x due to our assumptions. Because the maximum over continuous functions is lower-semicontinuous, $G_d(x)$ is lower-semicontinuous in x with respect to $\mathbb{R}^d$. From Proposition 2.55, it holds that F_n is lower-semicontinuous:

$$\liminf_{n \rightarrow \infty} F_n(\mu_n) \geq \bar{F}(\bar{\mu}).$$

Since $G_d(x_n) = G(\rho_n, x_n) = \frac{\kappa}{2} \|x - x_0\|^2 + \int f_2(z, x) d\pi(z) - F_n(\rho_n)$, F_n is lower-semicontinuous, and all other terms in G_d are continuous in x_n , we have

$$\limsup_{n \rightarrow \infty} G_d(x_n) = -\liminf_{n \rightarrow \infty} F_n(\rho_n) + \frac{\kappa}{2} \|\bar{x} - x_0\|^2 + \int f_2(z, \bar{x}) d\pi(z) \leq G_d(\bar{x}),$$

which means that G_d is upper-semicontinuous in x with respect to $\mathbb{R}^d$. Since both lower- and upper-semicontinuity have been shown, G_d is continuous in x . $\square$

Lemma 2.59. *Let Assumptions 2.15-2.17 hold. Fix $i \in \{1, \dots, d\}$ and let $x^{(h)} = \bar{x} + h e^{(i)}$ and $\rho_h = \arg \max_{\rho \in \mathcal{P}} G(\rho, x^{(h)})$, where $e^{(i)}$ denotes the i th standard unit vector in $\mathbb{R}^d$. Then for any sequence $y^{(h)} \rightarrow \bar{x}$ as $h \rightarrow 0$, we have*

$$\lim_{h \rightarrow 0} \partial_{x_i} G(\rho_h, y^{(h)}) = \partial_{x_i} G(\bar{\rho}, \bar{x}), \quad (2.43)$$

with x_i denoting the i th coordinate of $x \in \mathbb{R}^d$ and $\bar{\rho} = \operatorname{argmax}_{\rho \in \mathcal{P}} G(\rho, \bar{x})$. Moreover for any sequence $x_n \rightarrow \bar{x}$,

$$\lim_{n \rightarrow \infty} \int f_1(z, x_n) d\mathbf{r}(x_n)(z) = \int f_1(z, \bar{x}) d\mathbf{r}(\bar{x})(z). \quad (2.44)$$

Proof. To show that (2.43) holds, we want to compute

$$\begin{aligned} & \lim_{h \rightarrow 0} \left(\int \partial_{x_i} f_1(z, y^{(h)}) d\rho_h + \int \partial_{x_i} f_2(z, y^{(h)}) d\pi(z) + \kappa(y_i^{(h)} - (x_0)_i) \right) \\ &= \lim_{h \rightarrow 0} \left(\int \partial_{x_i} f_1(z, y^{(h)}) d\rho_h(z) \right) + \int \partial_{x_i} f_2(z, \bar{x}) d\pi(z) + \kappa(\bar{x}_i - (x_0)_i), \end{aligned}$$

where the limiting value of the last two terms follows immediately from the assumption that $f_2 \in C^2$. From a corollary of Fatou's lemma for weakly converging measures [107, Corollary 2.8], it is sufficient to show that $(\partial_{x_i} f_1(z, y^{(h)}))_{h \geq 0}$ is a.u.i with respect to $(\rho_h)_{h \geq 0}$. Using the Taylor expansion around $z = 0$ with remainder gives

$$\partial_{x_i} f_1(z, y^{(h)}) = \partial_{x_i} f_1(0, y^{(h)}) + \int_0^1 \nabla_z \partial_{x_i} f_1(tz, y^{(h)}) \cdot z dt.$$

From Assumption 2.15(b), we obtain

$$\left| \int_0^1 \nabla_z \partial_{x_i} f_1(tz, y^{(h)}) \cdot z dt \right| \leq L \|z\|,$$

and since $\partial_{x_i} f_1$ is continuous in x , we have

$$|g_h(z)| \leq L \|z\| + C_0$$

for $g_h(z) := \partial_{x_i} f_1(z, y^{(h)})$ and for some constant $C_0 > 0$ depending only on $\bar{x}$. For $K > C_0$, the key term for the a.u.i condition can be bounded by

$$\int |g_h(z)| \mathbb{1}\{z : |g_h(z)| \geq K\} d\rho_h(z) \leq \int_{B_{\frac{K-C_0}{L}}^c} (L \|z\| + C_0) d\rho_h(z),$$

where $B_{\frac{K-C_0}{L}}$ denotes the open ball of radius $\frac{K-C_0}{L}$ centered at zero. Taking the limit as $h \rightarrow 0$ results in

$$\begin{aligned} \lim_{h \rightarrow 0} \int |g_h(z)| \mathbb{1}\{z : |g_h(z)| \geq K\} d\rho_h(z) &\leq \lim_{h \rightarrow 0} \int_{B_{\frac{K-C_0}{L}}^c} (L \|z\| + C_0) d\rho_h(z) \\ &= \int_{B_{\frac{K-C_0}{L}}^c} (L \|z\| + C_0) d\bar{\rho}(z), \end{aligned}$$

where convergence to $\bar{\rho}$ is due to $\rho_h \rightarrow \bar{\rho}$ in $\mathcal{W}_2$ by Corollary 2.58. Now taking K to infinity,

$$\begin{aligned} & \lim_{K \rightarrow \infty} \lim_{h \rightarrow 0} \int |g_h(z)| \mathbb{1}\{z : |g_h(z)| \geq K\} d\rho_h(z) \\ & \leq \lim_{K \rightarrow \infty} L \int_{B^c_{\frac{K-C_0}{L}}} \|z\| \bar{\rho}(z) dz + \lim_{K \rightarrow \infty} C_0 \int_{B^c_{\frac{K-C_0}{L}}} \bar{\rho}(z) = 0, \end{aligned}$$

since $\int \|z\| \bar{\rho}(z) dz < \infty$ and $\int \bar{\rho}(z) dz = 1$. Therefore $\partial_{x_i} f_1(z, y^{(h)})$ is a.u.i. with respect to (ρ_h) . We conclude that

$$\lim_{h \rightarrow 0} \int g_h(z) d\rho_h(z) = \int \left(\lim_{h \rightarrow 0} g_h(z) \right) d\bar{\rho}(z) = \int \partial_{x_i} f_1(z, \bar{x}) d\bar{\rho}(z)$$

by applying Corollary [107, Corollary 2.8] and the limit $\lim_{h \rightarrow 0} \partial_{x_i} f_1(z, y^{(h)})$ exists because f_1 is C^2 in x .

To show that (2.44) holds, it is sufficient to show that $(f_1(z, x_n))_{n \in \mathbb{N}}$ is a.u.i with respect to (ρ_n) . Note that ρ_h denotes the best response to $x^{(h)} = \bar{x} + h e_i$, whereas ρ_n is the best response to *any* sequence $x_n \rightarrow \bar{x}$. Again using the Taylor expansion around $z = 0$,

$$f_1(z, x_n) = f_1(0, x_n) + \nabla_z f_1(0, x_n) \cdot z + \int_0^1 z^\top \nabla_z^2 f_1(tz, x_n) z dt.$$

Using that $f_1(0, \cdot) \in C^2$ and applying Assumptions 2.9(ii) and 2.15(a) results in

$$|f_1(z, x_n)| \leq c_1 + c_2 \|z\|^2, \quad \forall z \in \mathbb{R}^d, n \in \mathbb{N},$$

for some $c_1, c_2 \in \mathbb{R}_+$. The a.u.i condition, with $\tilde{B}^c := B^c_{\sqrt{\frac{K-c_1}{c_2}}}$ for $K > c_1$, is given by

$$\begin{aligned} \int |f_1(z, x_n)| \mathbb{1}\{z : |f_1(z, x_n)| \geq K\} d\rho_n(z) & \leq \int |f_1(z, x_n)| \mathbb{1}\{z : c_1 + c_2 \|z\|^2 \geq K\} d\rho_n(z) \\ & \leq \int_{\tilde{B}^c} (c_1 + c_2 \|z\|^2) d\rho_n(z). \end{aligned}$$

Taking the limit as $n \rightarrow \infty$ using the convergence of second moments given by Proposition 2.57 results in

$$\lim_{n \rightarrow \infty} \left(c_1 \int_{\tilde{B}^c} d\rho_n(z) + c_2 \int_{\tilde{B}^c} \|z\|^2 d\rho_n(z) \right) = c_1 \int_{\tilde{B}^c} d\bar{\rho}(z) + c_2 \int_{\tilde{B}^c} \|z\|^2 d\bar{\rho}(z).$$

Now taking the limit as K goes to ∞ ,

$$\lim_{K \rightarrow \infty} \left(c_1 \int_{\tilde{B}^c} d\bar{\rho}(z) + c_2 \int_{\tilde{B}^c} \|z\|^2 d\bar{\rho}(z) \right) = 0,$$

since $\int \bar{\rho}(z) = 1$ and $\int \|z\|^2 d\bar{\rho}(z) < \infty$. Therefore f_1 is a.u.i with respect to ρ_n , and applying [107, Corollary 2.8] results in

$$\lim_{n \rightarrow \infty} \int f_1(z, x_n) dr(x_n)(z) = \int f_1(z, \bar{x}) dr(\bar{x})(z) .$$

□

Proposition 2.60 (Version of Danskin's Theorem). *Let Assumptions 2.15-2.17 hold. Let $r(x)$ be as defined in (2.9). Then $G_d \in C^1(\mathbb{R}^d)$ and*

$$\nabla_x G_d(x) = (\nabla_x G(\rho, x))|_{\rho=r(x)} .$$

Proof. We will show that $G_d(x)$ is differentiable by showing that component-wise left and right derivatives coincide. More precisely, for any function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ (not necessarily differentiable), we denote

$$\begin{aligned} \partial_{x_i}^{\sup+} g(x) &= \limsup_{h \downarrow 0} \frac{g(x + he^{(i)}) - g(x)}{h}, & \partial_{x_i}^{\inf+} g(x) &= \liminf_{h \downarrow 0} \frac{g(x + he^{(i)}) - g(x)}{h}, \\ \partial_{x_i}^{\sup-} g(x) &= \limsup_{h \uparrow 0} \frac{g(x + he^{(i)}) - g(x)}{h}, & \partial_{x_i}^{\inf-} g(x) &= \liminf_{h \uparrow 0} \frac{g(x + he^{(i)}) - g(x)}{h}. \end{aligned}$$

Fix $\tilde{x} \in \mathbb{R}^d$ and recall that $G_d(\tilde{x}) = G(r(\tilde{x}), \tilde{x})$. We begin with the $\liminf$ for G . Because $r(\tilde{x})$ cannot make $G(\cdot, x)$ larger than $G(r(x), x)$ by definition of the best response, we have that

$$G_d(x) \geq G(r(\tilde{x}), x) \quad \forall x \in \mathbb{R}^d .$$

Then for any $h > 0$ and standard normal vector $e^{(i)} \in \mathbb{R}^d$,

$$\frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \geq \frac{G(r(\tilde{x}), \tilde{x} + he^{(i)}) - G(r(\tilde{x}), \tilde{x})}{h} .$$

Taking the $\liminf$ in $h > 0$ on both sides, we have

$$\partial_{x_i}^{\inf+} G_d(\tilde{x}) = \liminf_{h \downarrow 0} \frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \geq \partial_{x_i}^{\inf+} G(\rho, \tilde{x})|_{\rho=r(\tilde{x})} = \partial_{x_i} G(\rho, \tilde{x})|_{\rho=r(\tilde{x})} ,$$

since $G(\rho, \cdot) \in C^2(\mathbb{R}^d)$ for any $\rho \in \mathcal{P}(\mathbb{R}^d)$. Likewise, for $h < 0$,

$$\frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \leq \frac{G(r(\tilde{x}), \tilde{x} + he^{(i)}) - G(r(\tilde{x}), \tilde{x})}{h} ,$$

which, after taking $\limsup$ on both sides, results in

$$\partial_{x_i}^{\sup-} G_d(\tilde{x}) \leq \partial_{x_i} G(\rho, \tilde{x})|_{\rho=r(\tilde{x})} .$$

Next, consider the inequality

$$G_d(\tilde{x}) \geq G(r(x), \tilde{x}) \quad \forall x \in \mathbb{R}^d$$

and again setting $x = \tilde{x} + he^{(i)}$, we have for any $h > 0$,

$$\frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \leq \frac{G(r(\tilde{x} + he^{(i)}), \tilde{x} + he^{(i)}) - G(r(\tilde{x} + he^{(i)}), \tilde{x})}{h}. \quad (2.45)$$

Our goal is to apply Lemma 2.59 so that we can take the lim sup on the right-hand side. By the mean value theorem,

$$G(\rho, \tilde{x} + he^{(i)}) = G(\rho, \tilde{x}) + h\partial_{x_i}G(\rho, \xi^{(h)}) \quad (2.46)$$

for some $\xi^{(h)} \in [\tilde{x}, \tilde{x} + he^{(i)}]$ for any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$. Taking the lim sup on both sides of (2.45) and using (2.46), we have

$$\begin{aligned} \partial_{x_i}^{\sup+} G_d(\tilde{x}) &= \limsup_{h \downarrow 0} \frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \\ &\leq \limsup_{h \downarrow 0} \frac{G(r(\tilde{x} + he^{(i)}), \tilde{x} + he^{(i)}) - G(r(\tilde{x} + he^{(i)}), \tilde{x})}{h} \\ &= \limsup_{h \downarrow 0} \partial_{x_i}G(r(\tilde{x} + he^{(i)}), \xi^{(h)}) = \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, \end{aligned}$$

where the last line follows from (2.43) in Lemma 2.59. Now taking $h < 0$, we have

$$\frac{G_d(\tilde{x} + he^{(i)}) - G_d(\tilde{x})}{h} \geq \frac{G(r(\tilde{x} + he^{(i)}), \tilde{x} + he^{(i)}) - G(r(\tilde{x} + he^{(i)}), \tilde{x})}{h}.$$

which results in, after taking the lim inf and using the continuity of $r(\cdot)$ from Corollary 2.58,

$$\partial_{x_i}^{\inf-} G_d(\tilde{x}) \geq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}.$$

Collecting inequalities, we have shown

$$\begin{aligned} \partial_{x_i}^{\inf+} G_d(\tilde{x}) &\geq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, & \partial_{x_i}^{\sup-} G_d(\tilde{x}) &\leq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, \\ \partial_{x_i}^{\sup+} G_d(\tilde{x}) &\leq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, & \partial_{x_i}^{\inf-} G_d(\tilde{x}) &\geq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}. \end{aligned}$$

Chaining the inequalities together, we obtain

$$\begin{aligned} \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})} &\leq \partial_{x_i}^{\inf+} G_d(\tilde{x}) \leq \partial_{x_i}^{\sup+} G_d(\tilde{x}) \leq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, \\ \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})} &\leq \partial_{x_i}^{\inf-} G_d(\tilde{x}) \leq \partial_{x_i}^{\sup-} G_d(\tilde{x}) \leq \partial_{x_i}G(\rho, \tilde{x})|_{\rho=r(\tilde{x})}, \end{aligned}$$

and therefore $\partial_{x_i}^{\inf\pm} G_d(\tilde{x}) = \partial_{x_i}^{\sup\pm} G_d(\tilde{x}) = \partial_{x_i}G_d(\tilde{x})$ for any $\tilde{x}$ and so all partial derivatives of G_d exist at any $x \in \mathbb{R}^d$ with partial derivative given by $\partial_{x_i}G(\rho, x)|_{\rho=r(x)}$. If $x_n \rightarrow \bar{x}$, then $r(x_n) \rightarrow r(\bar{x})$ in $\mathcal{M}_2$ and $G_d(x_n) \rightarrow G_d(\bar{x})$ by Corollary 2.58. Further, from the expression for $\partial_{x_i}G(\rho, x)$ it is clear that $\partial_{x_i}G_d(x)$ is continuous for all i , and so we conclude that $G_d \in C^1(\mathbb{R}^d)$ and $\nabla_x G(\rho, x)|_{\rho=r(x)} = \nabla_x G_d(x)$. $\square$

Proof of Theorem 2.20. As the energy $G(\rho, x)$ is strongly λ_d -convex in x for each ρ , the energy $G_d(x) = \max_{\rho \in \mathcal{P}(\mathbb{R}^d)} G(\rho, x)$ is also strongly convex as a supremum of strongly convex functions. It follows that G_d is coercive, and has a unique minimizer $x_\infty \in \mathbb{R}^d$. By Proposition 2.60, $G_d \in C^1(\mathbb{R}^d)$. Convergence in norm then immediately follows from strong convexity of G_d : for solutions $x(t)$ to (2.9), we have

$$\frac{1}{2} \frac{d}{dt} \|x(t) - x_\infty\|^2 = -\nabla_x (G_d(x(t)) - G_d(x_\infty)) \cdot (x(t) - x_\infty) \leq -\lambda_d \|x(t) - x_\infty\|^2.$$

A similar result holds for convergence in entropy using the Polyák-Łojasiewicz convexity inequality

$$\frac{1}{2} \|\nabla G_d(x)\|_2^2 \geq \lambda_d (G_d(x) - G_d(x_\infty)),$$

which is itself a direct consequence of strong convexity of G_d . Then

$$\begin{aligned} \frac{d}{dt} (G_d(x(t)) - G_d(x_\infty)) &= \nabla_x G_d(x(t)) \cdot \dot{x}(t) \\ &= -\|\nabla_x G_d(x(t))\|^2 \leq -2\lambda_d (G_d(x(t)) - G_d(x_\infty)), \end{aligned}$$

and so the result in Theorem 2.20 follows. $\square$

2.C Auxiliary Lemmas

Lemma 2.61. *Let Assumptions 2.9(ii), 2.10, and 2.11 hold with $\lambda_c := \min\{\lambda_{f,1} + \lambda_{V,1}, \lambda_{f,2} + \lambda_{V,2}\} > 0$. For any $\gamma, \tilde{\gamma} \in \mathcal{P}_2 \times \mathcal{P}_2$ for which the left-hand side below is well-defined, the functional F_c satisfies*

$$\iint \left[\frac{z - \nabla \varphi(z)}{x - \nabla \psi(x)} \right] \cdot \left[\frac{\nabla_z \delta_\rho F_c[\tilde{\rho}, \tilde{\mu}](\nabla \varphi(z)) - \nabla_z \delta_\rho F_c[\rho, \mu](z)}{\nabla_x \delta_\mu F_c[\rho, \mu](x) - \nabla_x \delta_\mu F_c[\tilde{\rho}, \tilde{\mu}](\nabla \psi(x))} \right] d\rho(z) d\mu(x) \geq \lambda_c \overline{W}(\gamma, \tilde{\gamma})^2,$$

where (φ, ψ) are optimal transport maps such that $\tilde{\rho} = \nabla \varphi_\# \rho$ and $\tilde{\mu} = \nabla \psi_\# \mu$.

Proof. We break the left-hand side of the inequality into four parts; one with the diffusion terms, one with the coupling potential term, one with the convolution terms, and one with the external potentials:

$$\begin{aligned} C_1(\gamma, \tilde{\gamma}) &= -\alpha \int (z - \nabla \varphi(z)) \cdot (\nabla \log \tilde{\rho}(\nabla \varphi(z)) - \nabla \log \rho(z)) d\rho(z) \\ &\quad - \beta \int (x - \nabla \psi(x)) \cdot (\nabla \log \tilde{\mu}(\nabla \psi(x)) - \nabla \log \mu(x)) d\mu(x), \\ C_2(\gamma, \tilde{\gamma}) &= \iint (z - \nabla \varphi(z)) \cdot (\nabla_1 f(\nabla \varphi(z), \nabla \psi(x)) - \nabla_1 f(z, x)) d\gamma(z, x) \\ &\quad - \iint (x - \nabla \psi(x)) \cdot (\nabla_2 f(\nabla \varphi(z), \nabla \psi(x)) - \nabla_2 f(z, x)) d\gamma(z, x) \\ C_3(\gamma, \tilde{\gamma}) &= - \int (z - \nabla \varphi(z)) \cdot [(\nabla W_1 * \tilde{\rho})(\nabla \varphi(z)) - (\nabla W_1 * \rho)(z)] d\rho(z) \\ &\quad - \int (x - \nabla \psi(x)) \cdot [(\nabla W_2 * \tilde{\mu})(\nabla \psi(x)) - (\nabla W_2 * \mu)(x)] d\mu(x), \end{aligned}$$

and

$$C_4(\gamma, \tilde{\gamma}) = - \int (z - \nabla\varphi(z)) \cdot (\nabla V_1(\nabla\varphi(z)) - \nabla V_1(z)) d\rho(z), \\ - \int (x - \nabla\psi(x)) \cdot (\nabla V_2(\nabla\psi(x)) - \nabla V_2(x)) d\mu(x).$$

where ∇_i is the gradient operator with respect to the i^{th} argument. With the above definitions, it holds that

$$\sum_{i=1}^4 C_i(\gamma, \tilde{\gamma}) = \iint \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix} \cdot \begin{bmatrix} \nabla_z \delta_\rho F_c[\tilde{\rho}, \tilde{\mu}](\nabla\varphi(z)) - \nabla_z \delta_\rho F_c[\rho, \mu](z) \\ \nabla_x \delta_\mu F_c[\rho, \mu](x) - \nabla_x \delta_\mu F_c[\tilde{\rho}, \tilde{\mu}](\nabla\psi(x)) \end{bmatrix} d\gamma(z, x).$$

We claim that $C_1 = 0$ for all $\gamma, \tilde{\gamma} \in P_2^{ac} \times \mathcal{P}_2^{ac2}$. We prove this claim for one of the terms in C_1 , and the rest follow by relabeling variables. Let

$$c = \int \nabla\varphi(z) \cdot (\nabla \log \tilde{\rho}(\nabla\varphi(z)) - \nabla \log \rho(z)) d\rho(z).$$

We use the pushforward maps to write

$$c = \int \nabla\varphi(z) \cdot \frac{\nabla \tilde{\rho}(\nabla\varphi(z))}{\tilde{\rho}(\nabla\varphi(z))} d\rho(z) - \int \nabla\varphi(z) \cdot \frac{\nabla \rho(z)}{\rho(z)} d\rho(z) \\ = \int z \cdot \frac{\nabla \tilde{\rho}(z)}{\tilde{\rho}(z)} d\tilde{\rho}(z) - \int \nabla\varphi(z) \cdot \nabla \rho(z) dz \\ = \int z \cdot \nabla \tilde{\rho}(z) dz + \int \Delta\varphi(z) \rho dz = - \int \nabla \cdot (z) d\tilde{\rho}(z) + \int \nabla \cdot (z) d\tilde{\rho}(z) = 0,$$

Next, we expand the expression for $C_2(\gamma, \tilde{\gamma})$. We show the exact form of the Taylor expansion of $\nabla_1 f$ and the expansion for $\nabla_2 f$ follows similarly. Define the stacked variable $y = [z, x]$ and let $y_s := (1-s)[z, x] + s[\nabla\varphi(z), \nabla\psi(x)]$. Computing the Taylor expansion with respect to y results in

$$\nabla_1 f(\nabla\varphi(z), \nabla\psi(x)) = \nabla_1 f(z, x) + \int_0^1 \nabla_y \nabla_1 f(y_s) \cdot \begin{bmatrix} \nabla\varphi(z) - z \\ \nabla\psi(x) - x \end{bmatrix} ds \\ = \nabla_1 f(z, x) - \int_0^1 \begin{bmatrix} \nabla_1^2 f(x_s, z_s) \\ \nabla_{12}^2 f(x_s, z_s) \end{bmatrix} \cdot \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix} ds.$$

Plugging this expansion into C_2 results in

$$C_2(\gamma, \tilde{\gamma}) = \int_0^1 \iint \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix}^\top \cdot \begin{bmatrix} -\nabla_1^2 f(y_s) & -\nabla_{12}^2 f(y_s) \\ \nabla_{12}^2 f(y_s) & \nabla_2^2 f(y_s) \end{bmatrix} \cdot \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix} d\gamma(z, x) ds.$$

²Note that absolute continuity is required for the left-hand side in Lemma 2.61 to be well defined if $\alpha, \beta > 0$.

Since C_2 is a scalar, we use that $C_2(\gamma, \tilde{\gamma}) = \frac{1}{2}C_2(\gamma, \tilde{\gamma}) + \frac{1}{2}C_2(\gamma, \tilde{\gamma})^\top$, giving

$$C_2(\gamma, \tilde{\gamma}) = \int_0^1 \iint \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix}^\top \cdot \begin{bmatrix} -\nabla_1^2 f(y_s) & 0 \\ 0 & \nabla_2^2 f(y_s) \end{bmatrix} \cdot \begin{bmatrix} z - \nabla\varphi(z) \\ x - \nabla\psi(x) \end{bmatrix} d\gamma(z, x) ds.$$

By Assumption 2.9(ii), we use that $-\nabla_z^2 f(z, x) \succeq \lambda_{f,1} \mathbf{I}_{d_1}$ and $\nabla_x^2 f(z, x) \succeq \lambda_{f,2} \mathbf{I}_{d_2}$ to obtain

$$C_2(\gamma, \tilde{\gamma}) \geq \lambda_{f,1} \mathcal{W}_2(\rho, \tilde{\rho})^2 + \lambda_{f,2} \mathcal{W}_2(\mu, \tilde{\mu})^2.$$

Next, we show a lower-bound for C_3 . We show the exact Taylor expansion calculation explicitly for the first term:

$$\begin{aligned} (\nabla W_1 * \tilde{\rho})(\nabla\varphi(z)) - (\nabla W_1 * \rho)(z) &= \int \nabla W_1(\nabla\varphi(z) - z') d\tilde{\rho}(z') - \int \nabla W_1(z - z') d\rho(z') \\ &= \int [\nabla W_1(\nabla\varphi(z) - \nabla\varphi(z')) - \nabla W_1(z - z')] d\rho(z'). \end{aligned}$$

Define $z_s := (1-s)(z - z') + s(\nabla\varphi(z) - \nabla\varphi(z'))$. The expansion of ∇W_1 around $z - z'$ is

$$\begin{aligned} \nabla W_1(\nabla\varphi(z) - \nabla\varphi(z')) &= \nabla W_1(z - z') + \int \nabla^2 W_1(z_s)(\nabla\varphi(z) - z) ds \\ &\quad - \int \nabla^2 W_1(z_s)(\nabla\varphi(z') - z') ds, \end{aligned}$$

which, after integrating against $(z - \nabla\varphi(z))\rho(z')\rho(z)$ and relabeling, results in a quadratic and can be bound via Assumption 2.11 using that $\nabla^2 W_1$ is even:

$$\frac{1}{2} \iint \int_0^1 \begin{bmatrix} z - \nabla\varphi(z) \\ z' - \nabla\varphi(z') \end{bmatrix}^\top \begin{bmatrix} \nabla^2 W_1(z_s) & -\nabla^2 W_1(z_s) \\ -\nabla^2 W_1(z_s) & \nabla^2 W_1(z_s) \end{bmatrix} \begin{bmatrix} z - \nabla\varphi(z) \\ z' - \nabla\varphi(z') \end{bmatrix} ds d\rho(z') d\rho(z) \geq 0,$$

since the eigenvalues of $\begin{bmatrix} \nabla^2 W_1(z) & -\nabla^2 W_1(z) \\ -\nabla^2 W_1(z) & \nabla^2 W_1(z) \end{bmatrix}$ have a tight lower-bound of zero ($\lambda_{W,1} \geq 0$).

Computing a similar bound for the term dependent on μ , we have

$$C_3(\gamma, \tilde{\gamma}) \geq 0.$$

Lastly, we show the bound for C_4 . Computing for just the ρ -dependent term, we have

$$\begin{aligned} &\int (z - \nabla\varphi(z)) \cdot (\nabla V_1(\nabla\varphi(z)) - \nabla V_1(z)) d\rho(z) \\ &= - \int \int_0^1 (z - \nabla\varphi(z)) \nabla^2 V_1((1-s)z + s\nabla\varphi(z)) (z - \nabla\varphi(z)) d\rho(z) \\ &\leq -\lambda_{V,1} \int \|z - \nabla\varphi(z)\|^2 d\rho(z) = -\lambda_{V,1} \mathcal{W}_2(\rho, \tilde{\rho})^2, \end{aligned}$$

using Assumption 2.10. Computing similarly for the term dependent on μ , we have $C_4(\gamma, \tilde{\gamma}) \geq \lambda_{V,1} \mathcal{W}_2(\rho, \tilde{\rho})^2 + \lambda_{V,2} \mathcal{W}_2(\mu, \tilde{\mu})^2$. We now have shown that $C_1(\gamma, \tilde{\gamma}) + C_2(\gamma, \tilde{\gamma}) + C_3(\gamma, \tilde{\gamma}) + C_4(\gamma, \tilde{\gamma}) \geq \lambda_c \overline{W}(\gamma, \tilde{\gamma})^2$, concluding the proof. $\square$

Lemma 2.62. *Let Assumptions 2.9(ii), 2.10, and 2.11 hold with $\lambda_c > 0$. Fix $T > 0$. Let γ_t and γ'_t be any two solutions of the dynamics (2.3), with initial conditions $\rho_0, \rho'_0 \in \mathcal{P}_2^{ac}(\mathbb{R}^{d_1})$ and $\mu_0 = \delta_{x_0}, \mu'_0 = \delta_{x'_0}$, with $\alpha > 0$ and $\beta = 0$ such that $\rho_t, \rho'_t \in \mathcal{P}_2^{ac}(\mathbb{R}^{d_1})$ and $\mu_t = \delta_{x_t}, \mu'_t = \delta_{x'_t}$ for all $t \in (0, T)$. Assume $D_c(\gamma_0) < \infty$, $D_c(\gamma'_0) < \infty$ and $\nabla_z \delta_\rho F_c[\gamma_t](z), \nabla_z \delta_\rho F_c[\gamma'_t](z)$ are locally Lipschitz in z for all $t \in [0, T)$. Then γ_t and γ'_t satisfy*

$$\overline{W}(\gamma_t, \gamma'_t) \leq e^{-\lambda_c t} \overline{W}(\gamma_0, \gamma'_0) \quad \text{for all } t \in [0, T).$$

Proof. Define $\nabla \varphi_t(z)$ and $T_t(x)$ such that $\rho_t = \nabla \varphi_{t\#} \rho'_t$ and $x_t = T_t(x'_t)$. Note that $\nabla \varphi_t(z)$ is invertible because ρ_t, ρ'_t are absolutely continuous. The time derivative of $\overline{W}(\gamma_t, \gamma'_t)^2$ is

$$\frac{d}{dt} \overline{W}(\gamma_t, \gamma'_t)^2 = \frac{d}{dt} \mathcal{W}_2(\rho_t, \rho'_t)^2 + \frac{d}{dt} \|x_t - x'_t\|^2.$$

To compute the time derivative of the Wasserstein-2 squared distance via [203, Theorem 23.9], we use that

$$\int \|\nabla_z \delta_\rho F_c[\gamma_t](z)\|^2 d\rho_t(z) + \|\nabla_x \delta_\mu F_c[\gamma_t](x_t)\|^2 < \infty,$$

thanks to Lemma 2.32, with the same holding true for γ'_t . Then by [203, Theorem 23.9], the time derivative of $\mathcal{W}_2(\rho_t, \rho'_t)^2$ is

$$\begin{aligned} \frac{d}{dt} \mathcal{W}_2(\rho_t, \rho'_t)^2 &= 2 \int \langle \nabla_z \delta_\rho F_c[\gamma'_t](z), z - \nabla \varphi_t(z) \rangle d\rho'_t(z) \\ &\quad + 2 \int \langle \nabla_z \delta_\rho F_c[\gamma_t](z), z - (\nabla \varphi_t)^{-1}(z) \rangle d\rho_t(z) \\ &= 2 \int \langle \nabla_z \delta_\rho F_c[\gamma'_t](z) - \nabla_z \delta_\rho F_c[\gamma_t](\nabla \varphi_t(z)), z - \nabla \varphi_t(z) \rangle d\rho'_t(z) \end{aligned}$$

Since μ evolves as a Dirac for all time, the dynamics for x_t are

$$\dot{x}_t = v(\rho_t, x_t) := -\nabla_x \left(\int f(z, x_t) d\rho_t(z) + V_2(x_t) \right),$$

Note that the interaction term does not appear in the dynamics because

$$\nabla_x \int \delta_{x_t} W_2 * \delta_{x_t} = \nabla_x W_2(x_t - x_t) = \nabla_x W_2(0) = 0.$$

The time derivative of $\|x_t - x'_t\|^2$ is therefore

$$\begin{aligned} \frac{d}{dt} \|x_t - x'_t\|^2 &= 2 \langle x_t - x'_t, v(\rho_t, x_t) - v(\rho'_t, x'_t) \rangle \\ &= 2 \langle x'_t - T_t(x'_t), v(\rho'_t, x'_t) - v(\rho_t, T_t(x'_t)) \rangle \\ &= -2 \left\langle x'_t - T_t(x'_t), \nabla_x \left(\int (f(z, x'_t) - f(\nabla \varphi_t(z), T_t(x'_t))) d\rho'_t(z) + V_2(x'_t) - V_2(T_t(x'_t)) \right) \right\rangle \\ &= -2 \int \langle x - T_t(x), \nabla_x \delta_\mu F_c[\gamma'_t](x) - \nabla_x \delta_\mu F_c[\gamma_t](T_t(x)) \rangle d\mu'_t(x), \end{aligned}$$

where $\mu'_t = \delta_{x'_t}$. Summing $\frac{d}{dt} \|x_t - x'_t\|^2$ and $\frac{d}{dt} \mathcal{W}_2(\rho_t, \rho'_t)^2$ results in

$$\frac{d}{dt} \overline{W}(\gamma_t, \gamma'_t)^2 = -2 \iint \begin{bmatrix} z - \nabla \varphi_t(z) \\ x_t - T_t(x) \end{bmatrix} \cdot \begin{bmatrix} \nabla_z \delta_\rho F_c[\gamma_t](\nabla \varphi_t(z)) - \nabla_z \delta_\rho F_c[\gamma'_t](z) \\ \nabla_x \delta_\mu F_c[\gamma'_t](x) - \nabla_x \delta_\mu F_c[\gamma_t](T_t(x)) \end{bmatrix} d\gamma'_t(z, x),$$

which is the same expression given in the proof of contraction in the setting where both $\mu_t, \rho_t \in \mathcal{P}_2^{ac}$ (Proposition 2.33) by choosing $\psi_t(x) := \frac{1}{2} \|x - x'_t + x_t\|^2$ resulting in $T_t = \nabla \psi_t$ on $\text{supp } \mu'_t$. Contraction, boundedness of the second moments, and convergence to the steady state follow similarly. $\square$

Lemma 2.63. *For any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and any $\varepsilon > 0$, it holds that $\int \rho \log \rho \geq -\varepsilon \int \|z\|^2 d\rho(z) - c_\varepsilon$, for some $c_\varepsilon \in [0, \infty)$.*

Proof. Consider the function $a(x) : [0, \infty) \rightarrow \mathbb{R}$ defined as $a(x) = x \log x$. The Legendre dual, given by $a^*(y) = \sup_{x \geq 0} [x \cdot y - a(x)]$, is $a^*(y) = e^{y-1}$. By definition, for all $x, y \in [0, \infty)$, it holds that $a(x) + a^*(y) \geq x \cdot y$. Selecting $x = \rho(z)$ and $y = -\varepsilon \|z\|^2$ for any $\varepsilon > 0$, for any value of $z \in \mathbb{R}^d$, we have

$$\rho(z) \log \rho(z) + e^{-\varepsilon \|z\|^2 - 1} \geq -\|z\|^2 \varepsilon \rho(z),$$

and integrating this function over all z results in

$$\int \rho \log \rho dz \geq -\varepsilon \int \|z\|^2 d\rho(z) - c_\varepsilon,$$

where $c_\varepsilon := \int e^{-\varepsilon \|z\|^2 - 1} dz \in [0, \infty)$. $\square$

Lemma 2.64 (Lower Semicontinuity of Entropy). *For every sequence $\rho_n \rightharpoonup \rho$ converging weakly according to Definition 2.1, we have $\liminf_{n \rightarrow \infty} \int \rho_n \log \rho_n \geq \int \rho \log \rho$.*

Proof. A uniform bound on the second moment implies a uniform bound on the first moment:

$$\int \|z\| d\rho_n(z) \leq \int \left(\|z\|^2 + \frac{1}{4} \right) d\rho_n(z) \leq C + \frac{1}{4},$$

then the result follows from [193, Proposition 2.1]. $\square$

Lemma 2.65 (Upper and Lower Bounds on m). *Let $m(z, x) := -f_1(z, x) - \alpha \log \tilde{\rho}(z) \in C^2(\mathbb{R}^d \times \mathbb{R}^d)$ and $x_n \rightarrow \bar{x}$ in $\mathbb{R}^d$.*

(1) *Under Assumptions 2.9(ii) and 2.10, there exist constants $\hat{c}_1, \hat{c}_2 \in \mathbb{R}$ such that*

$$m(z, x_n) \geq -\hat{c}_1 + \hat{c}_2 \|z\|^2 \quad \forall z \in \mathbb{R}^d, n \in \mathbb{N}.$$

Further, if $\lambda_b > 0$, then one can choose $\hat{c}_2 > 0$.

(2) Under Assumptions 2.15 and 2.16, there exists constants $c_1 \in \mathbb{R}$ and $c_2 > 0$ such that

$$m(z, x_n) \leq c_1 + c_2 \|z\|^2 \quad \forall z \in \mathbb{R}^d, n \in \mathbb{N}.$$

Proof. Taylor expanding $m(z, x)$ first in z and then the second term in x , we have

$$\begin{aligned} m(z, x) &= m(0, x) + z^\top \nabla_z m(0, x) + \frac{1}{2} z^\top \nabla_z^2 m(\xi_2, x) z \\ &= m(0, x) + z^\top \nabla_z m(0, 0) + z^\top \nabla_{xz}^2 m(0, \xi_1) x + \frac{1}{2} z^\top \nabla_z^2 m(\xi_2, x) z, \end{aligned}$$

with $\xi_1 = sx$ for some $s \in [0, 1]$, $\xi_2 = \tau z$ for some $\tau \in [0, 1]$.

Proof of (1): By Assumptions 2.9(ii) and 2.10, we have for some continuous function $\hat{c}(x) > 0$

$$m(z, x) \geq m(0, x) + z^\top \nabla_z m(0, 0) - L \|z\| \|x\| + \frac{\lambda_b}{2} \|z\|^2 \geq \hat{c}(x) + \hat{c}_2 \|z\|^2,$$

with $\hat{c}(x) \in \mathbb{R}$ and $\hat{c}_2 > 0$ if $\lambda_b > 0$. Substituting $x = x_n$ in the above estimate and using that x_n converges to $\bar{x}$, there exists a constant $\hat{c}_1 \in \mathbb{R}$ such that

$$m(z, x_n) \geq -\hat{c}_1 + \hat{c}_2 \|z\|^2. \quad (2.47)$$

Proof of (2): Similarly for the upper bound,

$$m(z, x) \leq m(0, x) + z^\top \nabla_z m(0, 0) + L \|z\| \|x\| + \frac{-\ell_1 + \alpha \tilde{\Lambda}}{2} \|z\|^2 \leq c(x) + c_2 \|z\|^2,$$

where $c_2 > 0$ and $c(x)$ is a continuous function. The upper bounds $-\ell_1$ and L come from Assumption 2.15 and $\tilde{\Lambda}$ from Assumption 2.16. Again using that x_n converges to $\bar{x}$, there exists a constant c_1 such that

$$m(z, x_n) \leq c_1 + c_2 \|z\|^2.$$

□

Lemma 2.66. Let $m \in C^2(\mathbb{R}^d \times \mathbb{R}^d)$ such that $\|\nabla_z^2 m(z, x)\|_2 \geq \lambda$ and $\|\nabla_{xz}^2 m(z, x)\|_2 \leq L$ for all $z, x \in \mathbb{R}^d$ for some $\lambda, L > 0$. Consider a sequence of vectors $x_n \in \mathbb{R}^d$ converging to some limit $\bar{x} \in \mathbb{R}^d$ and any sequence of measures $(\mu_n)_{n \in \mathbb{N}} \in \mathcal{P}(\mathbb{R}^d)$ narrowly converging to some limit $\bar{\mu} \in \mathcal{P}(\mathbb{R}^d)$. Then there exists a constant $c \in \mathbb{R}$ such that $m_n(z) := m(z, x_n) \geq c$ for all $z \in \mathbb{R}^d$, for all $n \geq 0$. In particular, $(m_n^-)_{n \in \mathbb{N}}$ is asymptotically uniformly integrable with respect to $(\mu_n)_{n \in \mathbb{N}}$, where $m_n^-(z) = -\min\{m_n(z), 0\}$.

Proof. Let $z^*(x) : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $m^*(x) : \mathbb{R}^d \rightarrow \mathbb{R}$, and $\bar{z} \in \mathbb{R}^d$ be given by

$$m^*(x) := \min_{z \in \mathbb{R}^d} m(z, x), \quad z^*(x) := \operatorname{argmin}_{z \in \mathbb{R}^d} m(z, x), \quad \bar{z} = \operatorname{argmin}_{z \in \mathbb{R}^d} m(z, \bar{x}).$$

Showing that $(m^*(x_n))_{n \in \mathbb{N}}$ has a uniform lower bound is sufficient to achieve the desired a.u.i result for $(m_n^-)_{n \in \mathbb{N}}$. We will show this uniform lower bound using the implicit function theorem [101, Theorem C.7]. By definition of the best response,

$$\nabla_z m(z, x) \Big|_{z=z^*(x)} = 0 \quad \Rightarrow \quad \nabla_z^2 m(z^*(x), x) \nabla_x z^*(x) + \nabla_{zx}^2 m(z^*(x), x) = 0,$$

and since $\nabla_z^2 m(z^*(x), x)$ is invertible due to strong convexity,

$$\nabla_x z^*(x) = - \left(\nabla_z^2 m(z, x)^{-1} \nabla_{zx}^2 m(z, x) \right) \Big|_{z=z^*(x)}.$$

Using the matrix bounds results in the following bound for $\nabla_x z^*(x)$

$$\|\nabla_x z^*(x)\|_2 \leq \frac{L}{\lambda} \quad \forall x \in \mathbb{R}^d.$$

This gradient bound provides a bound on the distance of $z^*(x_n)$ from $z^*(\bar{x})$ as

$$\|z^*(x_n) - z^*(\bar{x})\|_2 \leq \frac{L}{\lambda} \|x_n - \bar{x}\|.$$

In particular, $z^*(x)$ is continuous and therefore $m^*(x) := m(z^*(x), x)$ is continuous in x . Hence $m(z, x_n) \geq m^*(x_n) \geq c$ for all $n > 0$ for some $c \in \mathbb{R}$. If $c \geq 0$, then the a.u.i condition is immediately satisfied because $m_n^-(z) = 0$ for all n . If $c < 0$, then for all $K > |c|$,

$$\limsup_{n \rightarrow \infty} \int |m_n^-(z)| \mathbb{1}_{\{z : |m_n^-(z)| \geq K\}} d\mu_n(z) = 0$$

and m_n^- is therefore asymptotically uniformly integrable. $\square$

Acknowledgments

The authors are grateful for helpful discussions with José A. Carrillo and Filippo Santambrogio, as well as Eitan Levin, Matthieu Darcy and Pau Batlle. In particular, Filippo Santambrogio suggested the strategy for the proof of Proposition 2.60 using the Γ -convergence result Proposition 2.55, and pointed us to the version of Lemma 2.63 in [193]. LC is funded by an NDSEG fellowship from the AFOSR and a PIMCO fellowship, FH is supported by start-up funds at the California Institute of Technology and by NSF CAREER Award 2340762, EM is supported in part from NSF award 2240110, LJR is supported in part by NSF 1844729, NSF 2312775, and ONR YIP N000142012571.

Chapter 3

MONOTONE MULTISPECIES FLOWS

L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Monotone Multispecies Flows, June 2025.
DOI: 10.48550/arXiv.2506.22947. arXiv:2506.22947

ABSTRACT

We present a novel notion of λ -monotonicity for an n -species system of partial differential equations governed by mass-preserving flow dynamics, extending monotonicity in Banach spaces to the Wasserstein-2 metric space. We show that monotonicity implies the existence of and convergence to a unique steady state, convergence of the velocity fields and second moments, and contraction in the Wasserstein-2 metric, at rates dependent on λ . In the special setting of Wasserstein-2 gradient descent of different energies for each species, we prove convergence to the unique Nash equilibrium of the associated energies and delineate the relationship between monotonicity and displacement convexity. This extends known zero-sum results in infinite-dimensional game theory to the general-sum setting. We provide a number of examples of monotone coupled gradient flow systems, including cross-diffusion, gradient flows with potentials, nonlocal interaction, linear and nonlinear diffusion, and min-max systems, and draw connections to a class of mean-field games. Numerically, we demonstrate convergence of a four-player economic model for service providers and strategic users competing in a market, and a degenerately monotone game.

3.1 Introduction

In the analysis of nonlinear dynamical systems, understanding the long-time behavior is important for modeling applications. When the dynamical system evolves via an operator on Banach spaces, *monotonicity* has proven to be a useful criteria under which convergence can be shown. In this work, we investigate mass-preserving flows, which are not in a normed space, and develop a notion of monotonicity in Wasserstein space which gives similar results to those for Banach spaces. Mass-preserving flow equations are commonly used to describe physical systems, such as fluid flow, biological systems [183, Ch 8], traffic [160], and increasingly appear in emerging areas like generative modeling [140], opinion dynamics [95], and strategic learning environments [62, 63]. Importantly, the monotonicity condition we propose has a structure that allows for systems where there is more than one entity evolving, such as in coupled flow equations. Understanding the long-time behavior of the coupled dynamics is a key question in these applications; in particular, monotonicity allows us to determine the existence of and convergence to a unique steady state.

Our notion of monotonicity is a natural extension of the well-studied condition for Euclidean dynamical systems in finite dimensions and Banach spaces in infinite dimensions, providing a condition for contraction of nonlinear dynamics. We show how this idea can be applied to velocity fields driving the evolution of probability measures. In the monotone flow setting, we consider an n -species system where each species $\rho_i \in \mathcal{P}(\mathbb{R}^{d_i})$, with $\mathcal{P}(\mathbb{R}^{d_i})$ the space of probability measures on $\mathbb{R}^{d_i}$, evolves according to a velocity field $v_i[\rho] : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^{d_i}$, which can depend on all species, denoted $\rho = [\rho_1, \dots, \rho_n]$. The mass of each ρ_i is conserved due to the divergence structure. We prove that monotonicity is a sufficient condition for dynamics of the form

$$\partial_t \rho_i + \operatorname{div}(\rho_i v_i[\rho]) = 0 \quad \forall i \in [1, \dots, n], \quad (3.1)$$

to contract in the joint Wasserstein-2 metric.

We additionally study a particular choice of velocity fields which results in a system of coupled gradient flows. In the Euclidean setting, all norms are equivalent, whereas in infinite dimensions, the choice of metric determines the topology and the gradient flow trajectory. We focus on the Wasserstein-2 metric, which results in a system of partial differential equations (PDEs) where the evolution of each species is written as a gradient flow with respect to its own energy. Let $\mathcal{P}_2(\mathbb{R}^d)$ be the space of probability measures with bounded second moments on $\mathbb{R}^d$, denote the product space $\overline{\mathcal{P}_2} = \mathcal{P}_2(\mathbb{R}^{d_1}) \times \dots \times \mathcal{P}_2(\mathbb{R}^{d_n})$, and denote by $\mathcal{W}_2$ the Wasserstein-2 metric (see section 3.A for details). Consider energy functionals $F_i : \mathcal{A} \subseteq \overline{\mathcal{P}_2} \rightarrow \mathbb{R}$, for a suitable choice of domain $\mathcal{A} \subseteq \overline{\mathcal{P}_2}$. Selecting $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ results in a system in which each species evolves in the direction of steepest descent of F_i in the Wasserstein-2 metric, given by

$$\partial_t \rho_i = -\nabla_{\mathcal{W}_2, \rho_i} F_i[\rho] = \operatorname{div}_{x_i}(\rho_i \nabla_{x_i} \delta_{\rho_i} F_i[\rho]) \quad \forall i \in [1, \dots, n]. \quad (3.2)$$

This system also has a game theoretic interpretation; in a game where each player i has a cost functional F_i which it seeks to minimize, the player i evolves a distribution $\rho_i \in \mathcal{P}_2(\mathbb{R}^{d_i})$ to minimize a cost functional $F_i[\rho_i, \rho_{-i}] = F_i[\rho] : \mathcal{A} \rightarrow \mathbb{R}$, where ρ_{-i} denotes the non- i^{th} player measures. We show that monotonicity implies displacement convexity of each F_i ; however, displacement convexity of each F_i is not enough to ensure monotonicity because monotonicity also depends on the coupling structure among all the species. Additionally, under mild conditions, the steady state of (3.2) corresponds to the unique Nash equilibrium of the game characterized by (F_i) .

A primary contribution of this work is the definition of monotonicity for velocity fields and energy functionals in this setting. Denote by $\Gamma_i(\rho_i^{(0)}, \rho_i^{(1)})$ the set of couplings whose marginals are $\rho_i^{(0)}, \rho_i^{(1)} \in \mathcal{P}_2(\mathbb{R}^{d_i})$, and let $\Gamma(\rho^{(0)}, \rho^{(1)}) := \Gamma_1(\rho_1^{(0)}, \rho_1^{(1)}) \times \cdots \times \Gamma_n(\rho_n^{(0)}, \rho_n^{(1)})$. Let Γ^* be the set of optimal transport plans, where $\Gamma_i^* = \{\gamma_i : \gamma_i \in \arg \min_{\gamma_i \in \Gamma_i(\rho_i^{(0)}, \rho_i^{(1)})} \int \|x_i - y_i\|^2 d\gamma_i(x_i, y_i)\}$. Let $x = [x_1, \dots, x_n] \in \mathbb{R}^d$, with $d = \sum_{i=1}^n d_i$. Let $\bar{W}$ be the joint metric over all species, given by $\bar{W}(\rho, \mu)^2 = \sum_{i=1}^n \mathcal{W}_2(\rho_i, \mu_i)^2$. The velocity field $v[\rho] : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is given by

$$v[\rho] := \begin{bmatrix} v_1[\rho] \\ \vdots \\ v_n[\rho] \end{bmatrix}.$$

Definition 3.1 (Monotonicity). *Let $\mathcal{A} \subseteq \bar{\mathcal{P}}_2$ be a geodesically convex set. If for any $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$,*

$$-\int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \int \|x - y\|^2 d\gamma(x, y) = \lambda \bar{W}(\rho^{(0)}, \rho^{(1)})^2, \quad (3.3)$$

for all $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$, then the velocity fields $\{v_i : \mathcal{A} \rightarrow \mathbb{R}^{d_i}\}_{i=1}^n$ are λ -monotone in $\mathcal{A}$. Similarly, we say the energy functionals $\{F_i : \mathcal{A} \rightarrow \mathbb{R}\}_{i=1}^n$ are λ -monotone in $\mathcal{A}$ if $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ are λ -monotone.

This definition of monotonicity is a multispecies version of the dissipativity condition given in [56], for single-valued v instead of set-valued v . Some texts use the word *accretive* instead of monotone to describe operators which satisfy positive-definiteness over a Hilbert space [31]; however, the space of measures endowed with the Wasserstein-2 metric is not a Hilbert space, and so one can think of our notion of monotonicity as an extension of the notion of accretivity. It was recently shown in [55, Theorem 7.6] that if a velocity field is dissipative, the dissipation inequality holds for all admissible plans, not just optimal couplings. In particular, when the dissipation inequality (3.3) holds for all pairs of measures on $\mathcal{P}_2$, it holds for all optimal plans between a pair of measures if and only if it holds for all admissible plans between those two measures. We prove a similar type of result for the multispecies setting.

Theorem 3.2 (Admissible Couplings). *Let $v : \overline{\mathcal{P}_2} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ be demicontinuous according to Definition 3.43. If for any $\rho^{(0)}, \rho^{(1)} \in \overline{\mathcal{P}_2}$ and $\gamma^* \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$ it holds*

$$- \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma^*(x, y) \geq \lambda \int \|x - y\|^2 d\gamma^*(x, y),$$

then for any $\gamma \in \Gamma(\rho^{(0)}, \rho^{(1)})$,

$$- \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \int \|x - y\|^2 d\gamma(x, y).$$

For a proof of this theorem, and a counterexample of why this result does not hold when v is λ -monotone on an arbitrary geodesically convex subset $\mathcal{A} \subseteq \overline{\mathcal{P}_2}$, see section 3.C. When the velocity field takes the form of a coupled gradient flow, we use the following notation for the row-wise gradient of the first variation

$$\overline{\nabla \delta_\rho} F[\rho](x) := \begin{pmatrix} \nabla_{x_1} \delta_{\rho_1} F_1[\rho](x_1) \\ \vdots \\ \nabla_{x_n} \delta_{\rho_n} F_n[\rho](x_n) \end{pmatrix}.$$

Then λ -monotonicity for $F = (F_i)$ writes as

$$\int \langle x - y, \overline{\nabla \delta_\rho} F[\rho^{(0)}](x) - \overline{\nabla \delta_\rho} F[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2. \quad (3.4)$$

Theorem 3.3 (Convergence, informal). *Let v be λ -monotone in $\mathcal{A}$ and $\rho(t) \in \mathcal{A}$ evolve according to (3.1) for all $t \geq 0$. Then*

- (a) *there exists a unique steady state ρ^∞ of (3.1), and $\rho(t)$ converges exponentially fast to ρ^∞ in $\overline{W}$ with rate λ , and*
- (b) *if $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$ for all $i \in [1, \dots, n]$, and the energies F_i are lower semicontinuous, then ρ^∞ is the unique Nash equilibrium in $\mathcal{A}$ for the game specified by F .*

For the full statement of this result, see Theorem 3.17. Key supporting results include uniform second moment bounds and convergence of the L^2 norm of the velocity fields, which is related to the energy dissipation in the single-species gradient flow setting. For a summary of assumptions and results, see fig. 3.1. The monotone structure mirrors existing finite-dimensional results (see [114] and references therein), which are outlined in what follows.

Definition 3.4 (Finite-Dimensional Monotonicity). *Let $u = [u_1, \dots, u_n] : \mathbb{R}^d \rightarrow \mathbb{R}^d$ where $u_i \in C^1(\mathbb{R}^d; \mathbb{R}^{d_i})$, and define $J(x) := \nabla u(x) \in \mathbb{R}^{d \times d}$. Then u is λ -monotone if it satisfies*

$$\begin{aligned} \text{(first-order)} \quad & \langle u(x) - u(x'), x - x' \rangle \geq \lambda \|x - x'\|^2 \quad \forall x, x' \in \mathbb{R}^d, \text{ or} \\ \text{(second-order)} \quad & J(x) + J(x)^\top \succeq 2\lambda I_d \quad \forall x \in \mathbb{R}^d. \end{aligned}$$

for some $\lambda \in \mathbb{R}$; u is strongly monotone if $\lambda > 0$.

In the finite-dimensional monotone setting, for dynamics $\dot{x}(t) = -u(t)$, commonly-used Lyapunov functions include

$$V_1(x) = \frac{1}{2} \|x - x^*\|^2, \quad V_2(x) = \frac{1}{2} \sum_{i=1}^n \|u_i(x)\|^2, \quad (3.5)$$

where x^* satisfies $u(x^*) = 0$. When u is strongly monotone, it can be shown that $\dot{V}_1(x) \leq -2\lambda V_1(x)$ using the first-order monotonicity condition, and that $\dot{V}_2(x) \leq -2\lambda V_2(x)$ using the second-order monotonicity condition. This results in exponential convergence of V_1 and V_2 to zero, which is sufficient to conclude exponential convergence of x to x^* . We will analyze the infinite-dimensional equivalents:

$$\frac{1}{2} \|x - x^*\|^2 \Leftrightarrow \frac{1}{2} \overline{W}(\rho, \rho^\infty)^2, \quad \frac{1}{2} \sum_{i=1}^n \|u_i(x)\|^2 \Leftrightarrow \frac{1}{2} \sum_{i=1}^n \int_{\mathbb{R}^{d_i}} \|v_i[\rho](x_i)\|^2 d\rho_i(x_i),$$

where $V_1(x)$ is analogous to the squared joint metric distance $\overline{W}$ between a measure ρ and the steady state ρ^∞ , and V_2 is analogous to the weighted L^2 norm of the velocity fields. For a finite dimensional monotone system where the time derivative for each species is given by $-u_i(x)$, we can define a lifted infinite-dimensional dynamical system with a velocity field $v[\rho](x) = -[\int u_1(x) d\rho_{-1}(x_{-1}), \dots, \int u_n(x) d\rho_{-n}(x_{-n})]$. Importantly, for monotone velocity fields in finite dimensions, the corresponding formulation in infinite dimensions is also monotone.

Proposition 3.5 (Finite- to Infinite-Dimensional Monotonicity). *Consider the finite dimensional velocity fields $u = [u_1, \dots, u_n] : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and define $v[\rho] = -[\int u_1(x) d\rho_{-1}(x_{-1}), \dots, \int u_n(x) d\rho_{-n}(x_{-n})]$. Then u is λ -monotone according to Definition 3.4 if and only if the velocity field v is λ -monotone in $\overline{\mathcal{P}}_2$ according to Definition 3.1.*

Proof. Assume u is λ -monotone according to Definition 3.4. Then the monotonicity condition for $v[\rho]$ to check is

$$\sum_{i=1}^n \int \left\langle x_i - y_i, \int \nabla_{x_i} u_i(x) d\rho_{-i}^{(0)}(x_{-i}) - \int \nabla_{y_i} u_i(y) d\rho_{-i}^{(1)}(y_{-i}) \right\rangle d\gamma_i(x_i, y_i) \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)}).$$

Since the marginals of γ_i are $\rho_i^{(0)}$ and $\rho_i^{(1)}$, we can rewrite the integral and apply the monotonicity condition for u

$$\begin{aligned} & \sum_{i=1}^n \int \left\langle x_i - y_i, \int \nabla_{x_i} u_i(x) d\rho_{-i}^{(0)}(x_{-i}) - \int \nabla_{y_i} u_i(y) d\rho_{-i}^{(1)}(y_{-i}) \right\rangle d\gamma_i(x_i, y_i) \\ &= \int \sum_{i=1}^n \langle x_i - y_i, \nabla_{x_i} u_i(x) - \nabla_{y_i} u_i(y) \rangle d\gamma(x, y) \geq \lambda \int \|x - y\| d\gamma(x, y), \end{aligned}$$

which completes the proof of finite-dimensional monotonicity implying infinite dimensional monotonicity. To show the other direction, selecting Dirac deltas for each species gives the desired inequality. $\square$

In contrast with the finite-dimensional setting, it is possible for monotonicity to fail for measures that are a single Dirac, while it holds for non-atomic measures; this corresponds to searching over mixed-Nash equilibria rather than just pure Nash equilibria. This occurs, for example, when $v_i[\rho](x_i)$ includes terms like $\log \rho(x_i)$ corresponding to linear diffusion. Our setting includes the finite dimensional setting as a special case (Proposition 3.5), and the notion of monotonicity introduced in Definition 3.1 is more general.

3.1.1 Contributions

In this work, we define a notion of monotonicity for a system of coupled flows in Wasserstein space which results in convergence to a steady state. We then consider the setting of coupled gradient flows, and provide examples of systems that satisfy this monotonicity condition, both analytically for general functional classes and numerically for specific applications.

Monotone and Coupled Gradient Flows In the monotone flow setting, we prove convergence in a weighted L^2 -norm of the velocity fields, convergence of the second moments to a ball where they remain for all time, and convergence of solutions to the monotone flow in a joint Wasserstein-2 metric to a unique steady state. In the coupled gradient flow setting, we show that the steady state is a Nash equilibrium of the game specified by the cost functions encoded by the dynamics. We show that our infinite-dimensional generalization of monotonicity encapsulates the Euclidean definition, and prove that monotonicity of the full system implies displacement convexity of each energy functional with respect to its own species.

Applications We include examples of monotone coupled gradient flow systems, with terms such as cross-diffusion, nonlocal interaction, and linear and nonlinear diffusion. We simulate a four-team game in which two service providers compete for users in two strategic populations, which manipulate their data over time. Then we draw a connection to multi-team mean-field games parameterized by a particular class of cost functions whose solution is a monotone coupled gradient flow. Finally, we empirically demonstrate convergence of a degenerately monotone game (a game which is not monotone, but almost monotone), using diffusion to ensure sufficient probability mass over the monotone region to still observe convergence. This suggests the possibility for our results to be extended to perturbations of monotone games.

3.1.2 Related Literature

This work draws from the optimization, game theory and PDE literature, applying convergence criteria and techniques from game theory to systems of flow equations. Contraction techniques for single-species PDEs fit elegantly into the multispecies framework, and the analysis of infinite-dimensional games utilizes classical PDE techniques.

Monotone Variational Inequalities in Optimization The notion of (nonlinear) monotone operators in Banach spaces is well-studied; it was developed at least as early as 1965 in [35]. While the results include infinite-dimensional Banach spaces, it does not apply to the space of probability measures, as the Wasserstein-2 metric is not induced by a norm. Therefore, one can view our notion of monotonicity as an extension of classical monotonicity results to a metric space. For an overview of related results, see [133, Ch 3], [66]. Other results detail algorithms for finding critical points, which are guaranteed to exist under monotonicity. The extragradient method [138] iterates by computing the gradient at a nearby point in the direction of the current gradient, and taking a step using the neighboring gradient. The results in [36] show that monotonicity implies weak convergence of a weighted history of a gradient descent trajectory, and the formulation of monotonicity in [57] hints at our formulation, by relating a monotone mapping on $\mathbb{R}^n$ to a set-valued mapping defined via the tangent space, which the authors use to derive a numerical method for finding solutions to a variational inequality. All of these works assume a Banach space structure; our work extends this to a non-Banach-space setting.

Game Theory In the context of game theory, monotonicity emerges in *diagonally strictly concave games* [191], as a condition under which existence of a unique Nash equilibrium can be proven for an n -player game. More recent results focus on learning algorithms which minimize functions under the assumption of monotone cost structures [201], or in the presence of strategic data [172, 173]. From the dynamical systems perspective, monotonicity is closely tied to the notion of contractive dynamics [38]; when the dynamics evolve a vector in Euclidean space, a λ -monotone system is said to be λ -contractive.

The extension of finite-dimensional games to infinite-dimensional games has recently been explored in related work. One such work provides a gradient ascent-descent scheme for solving a minimax problem over measures [154], while [207] presents a particle method for computing mixed Nash equilibria for two-player zero-sum games. A similar problem was formulated as an adversarially-robust optimization problem in [116], in which the authors also analyze a Wasserstein gradient ascent-descent particle method. Convergence to the mean-field limit of a particular zero-sum particle system was shown in [155]. Convergence in time in a zero-sum, infinite-dimensional setting was proven in [63], and it was shown in [62] that in a timescale-separated setting, the minimax problem

over measures is a gradient flow, providing convergence to the steady state under suitable conditions. Our framework generalizes these settings to n players and classes of cost functions beyond zero-sum, i.e. beyond the setting in which one player minimizes a cost that the other player maximizes. Many real-world settings have a non-zero-sum structure, where results such as Danskin’s theorem are no longer available and timescale separation can change the steady state.

Some recent work [122] gives a few monotonicity conditions for a single-species mean field game, including the Lasry–Lions monotonicity condition; while our condition is similar, we consider a dynamical system rather than the optimization problem of a mean field game. In one of our examples, we investigate the application of our monotonicity condition to a particular multi-population game with functionals of the form in [213]. Many results consider multi-population games, such as [78, 60], examining the existence of equilibria or as a method for clustering [13]. In our example, we analyze the equilibrium of the mean field game, which is a set of velocities, and show how the resulting dynamics have monotone structure.

Coupled PDEs The literature on multispecies interactions via convolution kernels can be viewed through a game-theoretic lens. These models are common in biology for instance for modeling dynamics of chemotaxis among other phenomena; see references in [62]. The majority of the PDE literature in this area concerns proofs for existence of solutions; for example, [112] provides a proof for existence, with uniqueness under additional assumptions, for solutions of two species with self interaction and cross-interaction; the kernel can be mildly singular. Under stronger regularity assumptions and for a large class of kernels, we give conditions for convergence to a steady state. Another work [49] characterizes the stationary states of a number of interaction systems. The results in [54] as well as references therein are very relevant to our setting due to their characterization of the equilibria of a two-species gradient flow system with nonlocal interaction terms, in addition to proving existence and uniqueness of solutions. While our condition for convergence to the steady state is generally more restrictive than settings considered in existing literature on existence of solutions, we offer a starting point for a new set of results on convergence conditions for these dynamics. For example, quadratic and attractive-repulsive kernels satisfy our notion of monotonicity. The recent work [92] poses a mean field variational inference problem as an optimization over n measures. This can be equivalently stated as a cooperative n -player game, in which all players aim to minimize the same energy functional. Our setting generalizes this setting to non-cooperative energies. In fact, the n -player dynamics in this work can be considered within the framework of vector-valued optimal transport that has been proposed very recently in [68]. There, the authors formulate a multispecies optimal transport problem, where mass moves both within a species and between species. Analyzing dynamics which are a gradient flow with respect to the metric proposed in [68] is an interesting extension of this work; see [68, Section 1.4] for details. The notion of

monotonicity (or accretivity, or dissipativity) for operators on Hilbert spaces used in [31], which we extend here to Wasserstein spaces, was used as a criterion for existence of solutions for measure evolution equations [56], in which the authors consider multi-valued vector fields; see also references therein.

3.1.3 Structure

We summarize relevant notation and definitions in section 3.2. Next, we present the existence of and convergence to the steady state results and proofs in section 3.3. section 3.4 details properties of monotonicity in the coupled gradient flow setting. section 3.5 presents applications in machine learning and game theory, illustrating how our theoretical results can be applied. In section 3.A we provide additional monotonicity results, such the equivalence between first order and second order monotonicity conditions (section 3.A.1), and a comparison of finite dimensional and infinite dimensional monotonicity as well as comparison between convexity (single species) and monotonicity (multi-species) (section 3.A.2). Next, section 3.B includes detailed calculations for interaction kernels supporting the results in section 3.4.4, and in section 3.C, we prove Theorem 3.2, which shows that optimal plans can be replaced by any admissible plan for the monotonicity condition in $\overline{\mathcal{P}_2}$.

3.2 Preliminaries

We consider a dynamical system with n species. The set of measures with finite second moment is denoted by $\mathcal{P}_2$, and each species ρ_i lives in $\mathcal{P}_2(\mathbb{R}^{d_i})$. Let $x = [x_1, \dots, x_n] \in \mathbb{R}^d$ where $d = \sum_{i=1}^n d_i$, $\rho := [\rho_1, \dots, \rho_n]$ be the vector of all measures, and ρ_{-i} denotes the vector of all non- i^{th} species, $\rho_{-i} := [\rho_1, \dots, \rho_{i-1}, \rho_{i+1}, \dots, \rho_n]$. Similarly, for a vector-valued function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^n$, $\overline{\nabla \varphi}_{-i} = [\nabla_{x_1} \varphi_1, \dots, \nabla_{x_{i-1}} \varphi_{i-1}, \nabla_{x_{i+1}} \varphi_{i+1}, \dots, \nabla_{x_n} \varphi_n]$ and $x_{-i} = [x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n]$. The notation $\int \|x\|^2 d\rho(x)$ means

$$\int \|x\|^2 d\rho(x) = \sum_{i=1}^n \int \|x_i\|^2 d\rho_i(x_i) .$$

Let C_c^∞ denote the set of infinitely-differentiable functions with compact support, C^k the set of k -times continuously differentiable functions, and $\mathcal{P}^{ac}(\mathbb{R}^{d_i})$ the set of probability measures on $\mathbb{R}^{d_i}$ that are absolutely continuous with respect to the Lebesgue measure. The Lebesgue measure on $\mathbb{R}^d$ is denoted by $\mathcal{L}^d$. The set of functions which are square integrable with respect to ρ_i are denoted by $L^2(\rho_i)$. The narrow topology is defined as convergence in duality with continuous bounded functions and the weak-* topology is defined in duality with continuous functions vanishing at infinity. We will use the following related notion of convergence, which we refer to as *weak topology*.

Definition 3.6 (Weak Convergence). *A sequence of measures (ρ_n) converges in the weak topology, denoted by $\rho_n \rightharpoonup \rho$, when (ρ_n) converges narrowly, that is, in duality with all continuous bounded functions, and there exists a uniform second moment bound for (ρ_n) .*

Let δ_{ρ_i} denote the first variation with respect to ρ_i and ∇_{x_i} the gradient with respect to x_i . For matrices A, B of the same dimension, the operator $A \succ B$ denotes that $A - B$ is positive-definite. Let I_d denote the identity matrix of dimension $d \times d$ and id the identity operator. Throughout this paper, we use the following notion for interpolation along transport plans and geodesics.

Definition 3.7 (Displacement Interpolants). *Define the displacement interpolant (or geodesic) between $\rho^{(0)}$ and $\rho^{(1)}$ as $\rho_i^{(s)} = I_{\#}^{(s)} \gamma_i$, where $I^{(s)}(x, y) = (1 - s)x + sy$, for all $s \in [0, 1]$ and $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$. If a transport map exists from $\rho_i^{(0)} \in \mathcal{P}_2$ to $\rho_i^{(1)} \in \mathcal{P}_2$ for all i , then there exists convex functions $\varphi_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}$ such that $\rho_i^{(1)} = [\nabla \varphi_i]_{\#} \rho_i^{(0)}$. Then $\rho^{(s)}$ can also be expressed as $\rho_i^{(s)} = [(1 - s)\text{id} + s\nabla \varphi_i]_{\#} \rho_i^{(0)}$. We define the operator $\overline{\nabla \varphi}(x) := [\nabla \varphi_1(x_1), \dots, \nabla \varphi_n(x_n)]$ as the vector of the convex functions which define the parameterization.*

Any displacement interpolant $\rho^{(s)}$ satisfies the following geodesic conservation of mass and momentum equations:

$$\partial_s \rho_i^{(s)} + \text{div}_{x_i} \left(\rho_i^{(s)} w_i^{(s)} \right) = 0, \quad \partial_s (\rho_i^{(s)} w_i^{(s)}) + \text{div}_{x_i} \left(\rho_i^{(s)} w_i^{(s)} \otimes w_i^{(s)} \right) = 0, \quad (3.6)$$

and if an invertible pushforward map exists between $\rho^{(0)}$ and $\rho^{(1)}$, then $w^{(0)}(x) = \overline{\nabla \varphi}(x) - x$ and $w^{(1)}(x) = x - (\overline{\nabla \varphi})^{-1}(x)$, where $(\overline{\nabla \varphi})^{-1}(x)$ is defined for a.e. $x \in \text{supp } \rho^{(1)}$.

A number of convexity results follow from the monotonicity definition (Definition 3.1). The notion of convexity we use is (uniform) displacement convexity, which can be viewed as convexity of a functional along $\mathcal{W}_2$ geodesics between any two measures in the domain of the functional.

Definition 3.8 (Displacement Convexity). *A functional $G : \mathcal{P}_2(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ is displacement convex if for all $\rho_i^{(0)}, \rho_i^{(1)}$ that are atomless we have*

$$G(\rho_i^{(s)}) \leq (1 - s)G(\rho_i^{(0)}) + sG(\rho_i^{(1)}),$$

where $\rho_i^{(s)}$ is given in Definition 3.7. Further, $G : \mathcal{P}_2(\mathbb{R}^{d_i}) \rightarrow \mathbb{R}$ is uniformly displacement convex with constant $\lambda > 0$ if

$$G(\rho_i^{(s)}) \leq (1 - s)G(\rho_i^{(0)}) + sG(\rho_i^{(1)}) - s(1 - s)\frac{\lambda}{2}\mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2.$$

Uniform convexity is analogous to strong convexity for Euclidean functions, with the inequality using the Wasserstein-2 distance instead of the two-norm. Equivalent first-order conditions for

uniform displacement convexity with constant λ are, where $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$,

$$\int \langle \nabla \delta_\rho G[\rho^{(0)}](x), y - x \rangle d\gamma(x, y) \leq G(\rho^{(1)}) - G(\rho^{(0)}) - \frac{\lambda}{2} \mathcal{W}_2(\rho^{(0)}, \rho^{(1)})^2, \quad (3.7a)$$

$$\text{or} \quad \int \langle x - y, \nabla \delta_\rho G[\rho^{(0)}](x) - \nabla \delta_\rho G[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \mathcal{W}_2(\rho^{(0)}, \rho^{(1)})^2. \quad (3.7b)$$

For details regarding these convexity inequalities, see [203, Proposition 16.2]. A sufficient second-order condition for uniform displacement convexity with constant λ is

$$\frac{d^2}{ds^2} G(\rho_i^{(s)}) \geq \lambda \mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2;$$

see [162] for displacement convexity details and tables 3.2 and 3.3 for comparison with monotonicity conditions.

For this work, we are not concerned with existence results for solutions to (3.1) and assume enough regularity of the velocity fields so that solutions $\rho(t) \in \overline{\mathcal{P}_2}$ exist. To analyze the long-time behavior of the system, we define the steady-state of the dynamics (3.1).

Definition 3.9 (Steady states for (3.1)). *Given $\rho^\infty \in \overline{\mathcal{P}_2}$, then ρ^∞ is a steady state for the system (3.1) if it satisfies*

$$\int \nabla \phi \cdot v_i[\rho^\infty](x_i) d\rho_i^\infty(x_i) = 0 \quad \forall \phi \in C_c^\infty(\mathbb{R}^{d_i}), \quad \forall i \in [1, \dots, n]. \quad (3.8)$$

In the coupled gradient flow setting, each player i evolves in the direction of steepest descent of F_i with respect to its own measure, in the Wasserstein-2 metric. Importantly, the dynamics (3.2) are *not* necessarily a joint gradient flow; rather, each player is doing gradient descent along its own energy functional that depends on the other agents. We will prove conditions under which dynamics with corresponding functionals that satisfy monotonicity will converge to a unique steady state, which is also the unique Nash equilibrium of the game $F = (F_i)$. Using the standard definition of Nash equilibrium, we present the definition of a Nash equilibrium in terms of the action ρ_i and objective F_i of each player.

Definition 3.10 (Nash Equilibrium). *A set of measures $\rho^* \in \mathcal{A}$ is a Nash equilibrium over $\mathcal{A}$ if,*

$$F_i[\rho_i^*, \rho_{-i}^*] \leq F_i[\rho_i, \rho_{-i}^*] \quad \forall \rho_i \in \mathcal{A}_i, \quad \forall i \in [1, \dots, n]. \quad (3.9)$$

A set of actions ρ^* is a Nash equilibrium when for each player i , all other players' actions are fixed and the i^{th} player has no incentive to deviate from its current strategy. Our notion of monotonicity can be seen as a generalization of monotonicity for mixed Nash equilibria.

3.3 Main Results

Our key result is convergence of the coupled flow system (3.1) to a unique steady state. We prove that monotonicity implies that the L^2 norm of the velocity fields converge to zero, and the second moments converge to a ball where they remain for all time. We prove a contraction result in $\overline{W}$, which results in the existence of a unique steady state. We can then show convergence from an initial condition to that steady state. In the setting where the system is a coupled gradient flow, we also show that the steady state corresponds to a Nash equilibrium.

For the remainder of the manuscript, any λ -monotone vector field v will be considered with a domain that is geodesically convex in $\overline{\mathcal{P}_2}$. Below, we summarize additional assumptions used for our main results to hold.

Assumption 3.11 (Absolutely Continuous Curves). *The solution $\rho(t)$ to (3.1) is absolutely continuous, that is, there exists some $g \in L^1([0, \infty))$ such that*

$$\overline{W}(\rho(t), \rho(s)) \leq \int_t^s g(\tau) d\tau, \quad \forall 0 \leq t \leq s < \infty.$$

Assumption 3.12 (Lower Semicontinuity). *For all $i \in [1, \dots, n]$ and for any fixed $\rho \in \mathcal{A}$, the functional $F_i[\cdot, \rho_{-i}]$ is lower-semicontinuous with respect to the weak topology over the domain $\mathcal{A}_i$.*

Assumption 3.11 is needed to compute the time derivative of the $\overline{W}$ metric. Assumption 3.12 is used to prove the existence of minimizers for each $F_i[\cdot, \rho_{-i}]$ in the coupled gradient flow setting, which is then applied together with monotonicity to show that the steady state is the unique Nash equilibrium.

The next three assumptions are not used in the main theorem; they are relevant for proving equivalence of first- and second-order notions of monotonicity, which is then used to show convergence of velocity fields. More precisely, Assumption 3.13 allows us to compute the time derivative of the L^2 norm of the velocities, and Assumptions 3.14 and 3.15 are used to construct a local, second-order notion of monotonicity (Lemma 3.37). These assumptions are not chosen to be optimal, and we expect relaxations to be possible.

Assumption 3.13 (Differentiability of v). *The terms $J_1[\rho] : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ and $J_2[\rho] : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ given by*

$$(J_1[\rho](x))_{ij} := \begin{cases} \nabla_{x_i} v_i[\rho](x_i) & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \in \mathbb{R}^{d_i \times d_i},$$

$$(J_2[\rho](x, \hat{x}))_{ij} := \nabla_{\hat{x}_j} \delta_{\rho_j} v_i[\rho](x_i, \hat{x}_j) \in \mathbb{R}^{d_i \times d_j},$$

are well-defined for ρ -a.e. x for all $\rho \in \mathcal{A}$.

Assumption 3.14 (Continuity of v). *The functions J_1, J_2 defined in Assumption 3.13 satisfy the following continuity property: If $\rho_\varepsilon \rightarrow \bar{\rho}$ in $\bar{W}$ as $\varepsilon \rightarrow 0$ and $\rho_\varepsilon, \bar{\rho} \in \mathcal{A}$, then for any $w \in L^2(\bar{\rho})$,*

$$\begin{aligned} & \iint \langle w(x), J_2[\rho_\varepsilon](x + \varepsilon w(x), \hat{x} + \varepsilon w(\hat{x})) \cdot w(\hat{x}) \rangle d\bar{\rho}(x) d\bar{\rho}(\hat{x}) \\ & + \int \langle w(x), J_1[\rho_\varepsilon](x + \varepsilon w(x)) \cdot w(x) \rangle d\bar{\rho}(x) \\ & \rightarrow \iint \langle w(x), J_2[\bar{\rho}](x, \hat{x}) \cdot w(\hat{x}) \rangle d\bar{\rho}(x) d\bar{\rho}(\hat{x}) + \int \langle w(x), J_1[\bar{\rho}](x) \cdot w(x) \rangle d\bar{\rho}(x) \text{ as } \varepsilon \rightarrow 0. \end{aligned}$$

Assumption 3.15 (Admissible Velocity Fields). *The velocity fields satisfy $v_i[\rho] \in Y_i(\rho_i, \mathcal{A}_i)$ for all $\rho \in \mathcal{A}$ and all $i \in [1, \dots, n]$, where*

$$Y_i(\rho_i, \mathcal{A}_i) = \overline{\{c_i(r_i - \text{id}) \in L^2(\rho_i) : (\text{id} \times r_i)_\# \rho_i \in \Gamma^*(\rho_i, (r_i)_\# \rho_i), (r_i)_\# \rho_i \in \mathcal{A}_i\}}^{L^2(\rho_i)}$$

where $r_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^{d_i}$ and $c_i > 0$ is a constant.

We can view the set $Y_i(\rho_i, \mathcal{A}_i)$ as a subset of the tangent space at ρ_i ; for the setting where $\mathcal{A}_i = \mathcal{P}_2(\mathbb{R}^{d_i})$, it was shown in [9, Theorem 8.5.1] that Y_i is exactly the Wasserstein-2 tangent space at ρ_i . For example, if $A_i = L^1(\mathbb{R}^{d_i})$, then $r_i = 0$ is not in Y_i since it pushes $\rho_i \in L^1(\mathbb{R}^{d_i})$ to a Dirac, which is not in $L^1(\mathbb{R}^{d_i})$; using the entire Wasserstein-2 tangent space may push densities outside of A_i .

Remark 3.16. *Absolute continuity of solution curves (Assumption 3.11) can be obtained as a consequence of related properties.*

(i) *For instance, a sufficient condition for Assumption 3.11 to hold is $v[\rho(t)](\cdot) \in L^2(\rho(t))$ for all $t \geq 0$ and $\int_0^\infty \int \|v[\rho(t)](x)\|^2 d\rho(t) dt < \infty$ [194, Theorem 5.14].*

(ii) *If $\rho(t) \in \mathcal{A}$ for all $t \geq 0$, and Assumptions 3.13 to 3.15 together with*

$$\frac{1}{2} \sum_{i=1}^n \int \|v_i[\rho(0)](x_i)\|^2 d\rho_i(0, x_i) < \infty$$

hold, it follows from λ -monotonicity for $\lambda > 0$ (Proposition 3.18) that the sufficient condition in (i) holds, and so Assumption 3.11 is satisfied.

For a given velocity field $v : \mathcal{A} \rightarrow \mathbb{R}^d$, we consider the set

$$S = \left\{ \rho \in \overline{\mathcal{P}_2} : \int_\Omega \|v[\rho](x)\|^2 d\rho(x) < \infty \quad \forall \Omega \subset \mathbb{R}^d, \Omega \text{ compact} \right\}.$$

Note that for any steady state ρ^∞ of the flow (3.1) satisfying Definition 3.9, we have $v_i[\rho^\infty](x_i) = 0$ for ρ_i^∞ -a.e. x_i and therefore $\rho^\infty \in S$.

Theorem 3.17 (Main Result). *Let Assumption 3.11 hold. Suppose (3.1) is λ -monotone in $\mathcal{A} \subseteq \overline{\mathcal{P}_2}$ with $\lambda > 0$, and $\rho(t)$ evolves according to (3.1). If $\rho(t) \in \mathcal{A}$ for all $t \geq 0$, then*

- (a) *there exists a unique steady state $\rho^\infty \in S$ of (3.1) and if $\mathcal{A}$ is compact in $\overline{W}$, then $\rho^\infty \in S \cap \mathcal{A}$;*
- (b) *if $\rho^\infty \in S \cap \mathcal{A}$, the solution to (3.1) $\rho(t)$ converges exponentially fast to ρ^∞ in $\overline{W}$ with rate λ ,*

$$\overline{W}(\rho(t), \rho^\infty) \leq e^{-\lambda t} \overline{W}(\rho(0), \rho^\infty).$$

- (c) *If, additionally, $v_i = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$, Assumption 3.12 holds, and $\rho^\infty \in \mathcal{A} \cap S$, then ρ^∞ is the unique Nash equilibrium of F in $\mathcal{A}$.*

We postpone the proof of this result to section 3.3.2. In the case where $\mathcal{A}$ is not compact, for specific functionals it can still be shown that $\rho^\infty \in \mathcal{A}$. See [63] for an example where diffusion is included in the evolution, and the steady state is shown to be in $\mathcal{A} = L^1(\mathbb{R}^d)$, which is not compact in the weak topology on $\mathcal{P}_2(\mathbb{R}^d)$.

In finite-dimensional game theory, convergence to Nash equilibria can be proven using a Lyapunov function dependent on the velocity, $V_2(x) = \frac{1}{2} \sum_{i=1}^n \|u_i(x)\|^2$. The analogous infinite-dimensional Lyapunov function offers insight into the L^2 -norm of the velocity fields. In the single-player gradient flow setting where $v_1[\rho_1] = -\nabla \delta_{\rho_1} F_1[\rho_1]$, the Lyapunov function is closely related to the *energy dissipation*. It is the energy dissipation that would result from a single, decoupled gradient flow, but in the multispecies setting, each player's velocity field in its dynamics is dependent on the other players' actions. We present a convergence result for this functional and comment on the additional requirements needed to show convergence. We define the Lyapunov function $D : \mathcal{A} \rightarrow \mathbb{R}$

$$D(\rho) = \sum_{i=1}^n D_i(\rho), \quad D_i(\rho) = \frac{1}{2} \int \|v_i[\rho](x_i)\|^2 d\rho_i(x_i). \quad (3.10)$$

We show that $D(\rho)$ converges to zero exponentially with a rate dependent on the monotonicity parameter.

Proposition 3.18 (Convergence of D). *Let v be λ -monotone in $\mathcal{A}$, let Assumptions 3.13 to 3.15 hold, and suppose $\rho(t) \in \mathcal{A}$ for $t \geq 0$, with $D(\rho(0)) < \infty$. Then D as defined in (3.10) along solutions to (3.1) satisfies*

$$D(\rho(t)) \leq e^{-2\lambda t} D(\rho(0)) \quad \forall t \geq 0. \quad (3.11)$$

Proof. First, computing the derivative of D_i along solutions of (3.1),

$$\begin{aligned}\dot{D}_i(\rho) &= \sum_{j=1}^n \iint \langle v_i[\rho](x_i), \delta_{\rho_j} v_i[\rho](x_i, \bar{x}_j) \rangle \partial_t \rho_j(\bar{x}_j) d\rho_i(x_i) d\bar{x}_j \\ &\quad + \frac{1}{2} \int \|v_i[\rho](x_i)\|^2 \partial_t \rho_i(x_i) dx_i.\end{aligned}$$

Substituting the dynamics and using integration by parts we have

$$\begin{aligned}\dot{D}_i(\rho) &= \sum_{j=1}^n \left(\iint \langle v_i[\rho](x_i), \nabla_{\bar{x}_j} \delta_{\rho_j} v_i[\rho](x_i, \bar{x}_j) \cdot v_j[\rho](\bar{x}_j) \rangle d\rho_j(\bar{x}_j) d\rho_i(x_i) \right. \\ &\quad \left. + \int \langle v_i[\rho](x_i), \nabla_{x_i} v_i[\rho](x_i) \cdot v_i[\rho](x_i) \rangle d\rho_i(x_i) \right).\end{aligned}$$

Combining the derivatives for all species into vectors and matrices, we have

$$\dot{D}(\rho) = \int v[\rho](x)^\top J_1[\rho](x) v[\rho](x) d\rho(x) + \iint v[\rho](x)^\top J_2[\rho](x, \bar{x}) v[\rho](\bar{x}) d\rho(x) d\rho(\bar{x}).$$

Since $D(\rho(0)) < \infty$ and the second-order derivative terms are well-defined by Assumption 3.13, $\dot{D}(\rho(0))$ is also well-defined. Since $v_i[\rho^{(0)}] \in Y_i(\rho_i^{(0)}, \mathcal{A}_i)$ thanks to Assumption 3.15, applying the monotonicity result Lemma 3.37 at $t = 0$ with $w(x) = v[\rho^{(0)}](x)$, we have

$$\dot{D}(\rho^{(0)}) \leq -\lambda \int \|v[\rho^{(0)}](x)\|^2 d\rho^{(0)}(x) = -2\lambda D(\rho^{(0)}),$$

and Grönwall's lemma gives an exponential estimate of D with rate 2λ , allowing the above argument to be applied for any $t \geq 0$. $\square$

Remark 3.19. If $\mathcal{A} = \overline{\mathcal{P}_2}$, then we can use any admissible map in the proof of Lemma 3.37, due to Theorem 3.2, and so Proposition 3.18 holds without Assumptions 3.14 and 3.15.

3.3.1 Second Moments

Monotonicity is closely tied to properties of the second moment of each species. We detail conditions under which the second moment is uniformly bounded. We use $M(\rho) : \overline{\mathcal{P}_2} \rightarrow \mathbb{R}$ to denote the sum of the second moments of all species,

$$M(\rho) = \frac{1}{2} \int \|x\|^2 d\rho(x) = \frac{1}{2} \sum_{i=1}^n \int \|x_i\|^2 d\rho_i(x_i). \quad (3.12)$$

The time derivative of the second moment along solutions to (3.1) is

$$\frac{d}{dt} M(\rho(t)) = \int \langle x, v[\rho](x) \rangle d\rho(t, x) \quad \forall t \geq 0. \quad (3.13)$$

Using monotonicity, we upper-bound the right hand side of (3.13). The following properties specify settings where the second moment is not blowing up in finite time (Property 1) and where the second moment is uniformly bounded for all time (Property 2).

Property 1 (Finite Second Moments). *Given an initial condition $\rho(0) \in \overline{\mathcal{P}_2}$, the solution $\rho(t)$ to (3.1) satisfies*

$$M(\rho(t)) < \infty \quad \forall t \geq 0.$$

Property 2 (Uniformly Bounded Second Moments). *Given an initial condition $\rho(0) \in \overline{\mathcal{P}_2}$, the solution $\rho(t)$ to (3.1) satisfies, for some $c \in \mathbb{R}_+$,*

$$M(\rho(t)) \leq c \quad \forall t \geq 0,$$

While Property 1 is a sufficient condition for applying a contraction theorem in the analysis of $\overline{W}$ convergence, it is possible for it to hold while monotonicity fails. In the single-species setting, the solution to $\partial_t \rho_1(t) = \operatorname{div}(\rho_1(t) \nabla \log \rho_1(t)) = \Delta \rho_1(t)$ satisfies Property 1 but does not satisfy strict monotonicity or Property 2. Property 2 can be shown using monotonicity, with additional assumptions needed depending on the specific choice of $\mathcal{A}$. First, we show that monotonicity in $\overline{\mathcal{P}_2}$ results in Property 2 holding under mild conditions. Let $\bar{\delta}_0 := [\delta_0, \dots, \delta_0]$ be the Dirac at the origin; in the proof of the following lemma, we will use that the second moment of ρ can be characterized via $M(\rho) = \frac{1}{2} \overline{W}(\rho, \bar{\delta}_0)^2$.

Lemma 3.20 (Uniformly Bounded Second Moments in Domains with Diracs). *Assume $\bar{\delta}_0 \in \mathcal{A}$. Let Assumptions 3.13 to 3.15 hold, $v : \mathcal{A} \rightarrow \mathbb{R}^d$ be λ -monotone in $\mathcal{A}$ with $\lambda > 0$, and $\rho(t) \in \mathcal{A} \forall t \geq 0$ evolves according to (3.1). If $\|v[\bar{\delta}_0](0)\| < \infty$, $D(\rho(0)) < \infty$ and for all $\rho \in \mathcal{A}$, the velocity fields $v_i[\rho]$ satisfy $\nabla_{x_i} v_i[\rho] \in L_{loc}^\infty(\mathbb{R}^{d_i}; \mathbb{R}^{d_i} \times \mathbb{R}^{d_i})$, then Property 2 holds.*

Proof of Lemma 3.20. The sum of the second moments can be equated to $\overline{W}(\bar{\delta}_0, \rho(t))^2$ since $\bar{\delta}_0 = \overline{\nabla \varphi}(t) \# \rho(t)$ with $\overline{\nabla \varphi}(t, x) = 0$ for a.e. $x \in \operatorname{supp} \rho$ an optimal map, and therefore

$$\overline{W}(\bar{\delta}_0, \rho(t))^2 = \sum_{i=1}^n \int \|x_i\|^2 d\rho_i(t, x_i).$$

Because (3.1) is λ -monotone in $\mathcal{A}$ and $\mathcal{A}$ contains $\bar{\delta}_0$, the monotonicity condition ensures that

$$\int \langle x, v[\rho(t)](x) - v[\bar{\delta}_0](0) \rangle d\rho(t, x) \leq -\lambda \overline{W}(\bar{\delta}_0, \rho(t))^2.$$

We use Assumptions 3.13 to 3.15 to conclude from Proposition 3.18 that the L^2 norm of the velocity field is bounded. This allows us to apply [203, Theorem 23.9] to compute the time derivative of the Wasserstein metric,

$$\frac{d}{dt} \overline{W}(\rho(t), \bar{\delta}_0)^2 = 2 \int \langle x, v[\rho(t)](x) \rangle d\rho(t, x).$$

Substituting the monotonicity estimate for $\int \langle x, v[\rho(t)](x) - v[\bar{\delta}_0](0) \rangle d\rho(t, x)$ results in

$$\begin{aligned} \frac{d}{dt} \bar{W}(\rho(t), \bar{\delta}_0)^2 &\leq -2\lambda \bar{W}(\rho(t), \bar{\delta}_0)^2 + 2 \int \langle x, v[\bar{\delta}_0](0) \rangle d\rho(t, x) \\ &\leq -2\lambda \bar{W}(\rho(t), \bar{\delta}_0)^2 + 2 \left(\int \|x\|^2 d\rho(t, x) \right)^{1/2} \|v[\bar{\delta}_0](0)\| \\ &\leq -2\lambda \bar{W}(\rho(t), \bar{\delta}_0)^2 + c \bar{W}(\rho(t), \bar{\delta}_0), \end{aligned}$$

where $c = 2 \|v[\bar{\delta}_0](0)\| < \infty$ by assumption, and the Hölder inequality was used on the second line. Then we have that $\bar{W}(\rho(t), \bar{\delta}_0)^2 = 2M(\rho(t))$ exponentially converges to a ball with radius dependent on c and λ . Therefore $M(\rho(t))$ is uniformly bounded for all time inside a ball of size R , where

$$R = \max \left\{ \frac{c}{4\lambda}, M(\rho(0)) \right\}.$$

□

While Lemma 3.20 applies to the setting when $\mathcal{A}$ contains a Dirac, it is possible for the results to hold in greater generality. When Diracs are not included in $\mathcal{A}$, such as when v contains linear or nonlinear diffusion, applying monotonicity of (3.1) in $\bar{\mathcal{P}}_2$ to control the second moment is not feasible; we provide an extension of the previous lemma below. To show that the second moments are bounded under more general conditions, we will use a mollifier $\rho^{(\tau)} \in \mathcal{A}$ as a proxy for $\bar{\delta}_0$, which satisfies $\bar{W}(\rho^{(\tau)}, \bar{\delta}_0) < \infty$. For example, when $\mathcal{A} = L^1(\mathbb{R}^{d_i})^{\otimes n}$, selecting

$$(\rho^{(\tau)}(x))_i := \frac{1}{(2\pi\tau)^{d_i/2}} \exp\left(-\frac{\|x_i\|^2}{2\tau}\right),$$

results in $\bar{W}(\rho^{(\tau)}, \bar{\delta}_0) = \sqrt{d\tau}$, and allows us to apply the monotonicity inequality on $\rho^{(\tau)}$ when it cannot be applied to $\bar{\delta}_0$. This mollifier can be selected to have tails appropriate for the given velocity fields. The following lemma gives conditions under which Property 2 is satisfied.

Lemma 3.21 (Bounded Second Moments). *Let v be λ -monotone in $\mathcal{A}$, $\rho(t) \in \mathcal{A}$ evolve according to (3.1) for all $t \geq 0$, and Assumption 3.11 hold. Fix $\tau > 0$. If there exists $\rho^{(\tau)} \in \mathcal{A}$ such that*

$$c_\tau := \left(\int \|v[\rho^{(\tau)}](x)\|^2 d\rho^{(\tau)}(x) \right)^{1/2} < \infty \quad \text{and} \quad \bar{W}(\rho^{(\tau)}, \bar{\delta}_0) < \infty,$$

then

$$\sqrt{2M(\rho(t))} \leq \begin{cases} c_1 + c_0 e^{-\lambda t} & \text{if } \lambda \neq 0 \\ c_2 + c_\tau t & \text{if } \lambda = 0, \end{cases}$$

where

$$c_0 := \max \left\{ 0, \overline{W}(\rho(0), \rho^{(\tau)}) - \frac{c_\tau}{\lambda} \right\}, \quad c_1 = \frac{c_\tau}{\lambda} + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0), \quad c_2 = \overline{W}(\rho(0), \rho^{(\tau)}) + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0).$$

In particular, if $\lambda > 0$, then $\rho(t)$ satisfies Property 2.

Proof. The second moment (3.12) can be upper-bounded using the triangle inequality:

$$\sqrt{2M(\rho(t))} = \overline{W}(\rho(t), \bar{\delta}_0) \leq \overline{W}(\rho(t), \rho^{(\tau)}) + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0).$$

Let $\gamma_i(t) \in \Gamma^*(\rho_i(t), \rho_i^{(\tau)})$. Using Assumption 3.11 to apply [56, Theorem 3.14] for the time derivative of $\overline{W}(\rho(t), \rho^{(\tau)})^2$,

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \overline{W}(\rho(t), \rho^{(\tau)})^2 &= \int \langle x - y, v[\rho(t)](x) \rangle d\gamma(t, x, y) \\ &\leq -\lambda \overline{W}(\rho(t), \rho^{(\tau)})^2 + \int \langle x - y, v[\rho^{(\tau)}](y) \rangle d\gamma(t, x, y), \end{aligned}$$

where the last line follows from adding and subtracting $\int \langle x - y, v[\rho^{(\tau)}](y) \rangle d\gamma(t, x, y)$ and applying the monotonicity inequality. We use Hölder's inequality to obtain

$$\begin{aligned} \int \langle x - y, v[\rho^{(\tau)}](y) \rangle d\gamma(t, x, y) &\leq \left(\int \|x - y\|^2 d\gamma(t, x, y) \int \|v[\rho^{(\tau)}](x)\|^2 d\rho^{(\tau)}(x) \right)^{1/2} \\ &= c_\tau \overline{W}(\rho(t), \rho^{(\tau)}). \end{aligned}$$

Combining the last two estimates,

$$\frac{1}{2} \frac{d}{dt} \overline{W}(\rho(t), \rho^{(\tau)})^2 \leq -\lambda \overline{W}(\rho(t), \rho^{(\tau)})^2 + c_\tau \overline{W}(\rho(t), \rho^{(\tau)}).$$

When $\lambda \neq 0$, we have

$$\overline{W}(\rho(t), \rho^{(\tau)}) \leq \frac{c_\tau}{\lambda} + c_0 e^{-\lambda t},$$

where $c_0 = \max\{0, \overline{W}(\rho(0), \rho^{(\tau)}) - \frac{c_\tau}{\lambda}\}$. When $\lambda = 0$,

$$\overline{W}(\rho(t), \rho^{(\tau)}) \leq \overline{W}(\rho(0), \rho^{(\tau)}) + c_\tau t.$$

Using the triangle inequality, the second moments of $\rho(t)$ satisfy

$$\sqrt{2M(\rho(t))} = \overline{W}(\rho(t), \bar{\delta}_0) \leq \begin{cases} \frac{c_\tau}{\lambda} + c_0 e^{-\lambda t} + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0) & \text{if } \lambda \neq 0 \\ \overline{W}(\rho(0), \rho^{(\tau)}) + c_\tau t + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0) & \text{if } \lambda = 0 \end{cases}.$$

Setting $c_1 = \frac{c_\tau}{\lambda} + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0)$ and $c_2 = \overline{W}(\rho(0), \rho^{(\tau)}) + \overline{W}(\rho^{(\tau)}, \bar{\delta}_0)$ gives the desired inequality. From this inequality, we conclude that if $\lambda > 0$, then Property 2 is satisfied and if $\lambda \leq 0$, then Property 1 is satisfied. $\square$

3.3.2 Proof of main theorem

In this section, we will prove our main result, Theorem 3.17. The components of the proof and related assumptions are illustrated in fig. 3.1. Monotonicity is used to prove convergence of the velocity fields and a uniform bound on the second moments; in the setting where invertible pushforwards exist between any two points along the trajectory, these two key results could then be used to prove the main contraction and convergence results. However, to generalize the results to transport plans, Assumption 3.11 takes the place of these results. In the coupled gradient flow setting, we use the fact that monotonicity of F implies displacement convexity of each F_i , which we discuss in detail later in section 3.4. This, along with the lower semicontinuity assumption, allows us to conclude that the steady state ρ^∞ is the Nash equilibrium of F . To prove properties about the steady state and




Figure 3.1: The relationships between main results (in blue), assumptions (in violet), and the monotonicity condition (in orange) are illustrated. Assuming that solutions to (3.1) are absolutely continuous curves (Assumption 3.11) allows us to compute the time derivative of the Wasserstein-2 metric via [56, Thm 3.14] to show contraction. Then via [52, Lemma 7.3], a unique steady state $\rho^\infty \in \overline{\mathcal{P}_2}$ exists (Theorem 3.17a). Convergence of $\rho(t)$ to ρ^∞ follows by similar arguments (Theorem 3.17b). If the dynamics are a coupled gradient flow, that is, take the form (3.2), then using Assumption 3.11, Assumption 3.12 and displacement convexity of each F_i (Lemma 3.28a), we show that ρ^∞ is critical point and the unique Nash equilibrium of F (Theorem 3.17c). The convergence of the velocity fields (Proposition 3.18) and second moment bounds (Lemma 3.21) for solutions to (3.1) follow from monotonicity.

the Nash equilibrium in the coupled gradient flow setting, we utilize some additional lemmas.

Lemma 3.22 (Uniqueness of Critical Points). *There is no more than one critical point $\rho^* \in \mathcal{A}$ for the game specified by F under the monotonicity condition (Definition 3.1) when $\lambda > 0$.*

Proof. Consider $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ which are both critical points. The Euler-Lagrange condition

for each F_i requires that $\delta_{\rho_i} F_i[\rho^{(0)}](x_i) = b_i$ for $\rho_i^{(0)}$ -a.e. x_i , and $\delta_{\rho_i} F_i[\rho^{(1)}](x_i) = c_i$ for $\rho_i^{(1)}$ -a.e. x_i for some constants $b_i, c_i \in \mathbb{R}$. Therefore $\nabla_{x_i} \delta_{\rho_i} F_i[\rho^{(0)}](x_i) = 0$ for $\rho_i^{(0)}$ -a.e. x_i and $\nabla_{x_i} \delta_{\rho_i} F_i[\rho^{(1)}](x_i) = 0$ for $\rho_i^{(1)}$ -a.e. x_i . Using these values in the monotonicity definition, we have

$$\begin{aligned} \int \langle x - y, \overline{\nabla \delta_\rho F}[\rho^{(0)}](x) - \overline{\nabla \delta_\rho F}[\rho^{(1)}](y) \rangle d\gamma(x, y) &= \int \langle x - y, 0 \rangle d\gamma(x, y) \\ &= 0 \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2. \end{aligned}$$

Hence, $\overline{W}(\rho^{(0)}, \rho^{(1)}) = 0$, and there is at most one critical point in $\mathcal{A}$. $\square$

Lemma 3.23 (Steady State is Unique Critical Point). *Under λ -monotonicity of F in $\mathcal{A}$, a steady state of (3.2) $\rho^\infty \in \mathcal{A} \cap S$ is the unique critical point of F in $\mathcal{A}$.*

Proof. The steady state $\rho^\infty \in S$ satisfies

$$\int \nabla \phi \cdot \nabla_{x_i} \delta_{\rho_i} F_i[\rho^\infty](x_i) d\rho_i^\infty(x_i) = 0 \quad \forall \phi \in C_c^\infty(\mathbb{R}^{d_i}),$$

implying $\nabla_{x_i} \delta_{\rho_i} F_i[\rho^\infty](x_i) = 0$ for ρ_i^∞ a.e. x_i , and so $\delta_{\rho_i} F_i[\rho^\infty](x_i) = c_i$ for some constant $c_i \in \mathbb{R}$, ρ_i^∞ a.e. x_i ; therefore ρ^∞ is a critical point of each F_i . From Lemma 3.22, there is at most one critical point in $\mathcal{A}$, so $\rho^\infty \in \mathcal{A}$ is the unique critical point in $\mathcal{A}$. $\square$

Lemma 3.24 (Steady State is Nash Equilibrium). *Under λ -monotonicity in $\mathcal{A}$ and Assumption 3.12, a steady state $\rho^\infty \in \mathcal{A} \cap S$ of (3.2) is the unique Nash equilibrium of F in $\mathcal{A}$, according to Definition 3.10.*

Proof. From Lemma 3.23, the steady state ρ^∞ is the unique critical point of F . It remains to show that the critical point is also the Nash equilibrium, that is, that ρ^∞ satisfies

$$F_i[\rho_i^\infty, \rho_{-i}^\infty] \leq F_i[\rho_i, \rho_{-i}^\infty] \quad \forall \rho_i \in \mathcal{A}, \quad \forall i \in [1, \dots, n].$$

By Lemma 3.28(a), $F_i[\cdot, \rho_{-i}]$ is displacement convex for all $\rho_{-i} \in \mathcal{A}$. By Assumption 3.12, $F_i[\cdot, \rho_{-i}]$ is lower semi-continuous in the weak topology of $(\mathcal{P}_2, \mathcal{W}_2)$. From [9, Lemma 2.4.8], a unique minimizer of $F_i[\cdot, \rho_{-i}]$ exists in $\mathcal{P}_2$, and because the critical point of the game F is unique in $\mathcal{A} \subseteq \overline{\mathcal{P}_2}$, it must be the minimizer. Applying this for all $i \in [1, \dots, n]$, the Nash condition holds and therefore the steady state is a Nash equilibrium. $\square$

Proof of Theorem 3.17. Recall the definition of the set S

$$S = \left\{ \rho \in \overline{\mathcal{P}_2} : \int_{\Omega} \|v[\rho](x)\|^2 d\rho(x) < \infty \quad \forall \Omega \subset \mathbb{R}^d, \Omega \text{ compact} \right\}.$$

Proof of (a): The proof follows from showing that the dynamics are a contraction in the joint Wasserstein-2 metric. Let $\rho(t)$ and $\eta(t)$ be two solutions to (3.1). We have from Assumption 3.11 that $\rho(t), \eta(t) \in \mathcal{A}$ are absolutely continuous curves. Hence, the time derivative of the joint Wasserstein-2 metric between the ρ and η is

$$\frac{d}{dt} \overline{W}(\rho(t), \eta(t))^2 = 2 \int \langle x - y, v[\eta(t)](x) - v[\rho(t)](y) \rangle d\gamma(t, x, y),$$

by [56, Theorem 3.14], where $\gamma_i(t) \in \Gamma^*(\rho_i(t), \eta_i(t))$. The right-hand side of the inequality is exactly the negative of the left-hand side of the monotonicity definition, so by λ -monotonicity in $\mathcal{A}$,

$$\frac{d}{dt} \overline{W}(\rho(t), \eta(t))^2 \leq -2\lambda \overline{W}(\rho(t), \eta(t))^2.$$

By Grönwall's inequality, $\overline{W}(\rho(t), \eta(t)) \leq e^{-\lambda t} \overline{W}(\rho(0), \eta(0))$, and therefore the dynamics are contractive in the joint Wasserstein-2 metric. From [9, Proposition 7.1.5], $\mathcal{P}_2(\mathbb{R}^d)$ endowed with the Wasserstein-2 metric is complete because $\mathbb{R}^d$ is complete, and since $\rho(t), \eta(t)$ are absolutely continuous, we can apply [52, Lemma 7.3] to obtain the existence of a unique stationary point $\rho^\infty \in \overline{\mathcal{P}_2}$. If $\mathcal{A}$ is compact and the solution stays in $\mathcal{A}$ for all time, the long-time limit ρ^∞ is also in $\mathcal{A}$.

Proof of (b): This step mirrors the convergence shown for $V_1(x) = \frac{1}{2} \|x - x^*\|^2$ in the Euclidean setting of monotonicity. We use that $\frac{d}{dt} \overline{W}(\rho(t), \rho^\infty)^2 = \sum_{i=1}^n \frac{d}{dt} \mathcal{W}_2(\rho_i(t), \rho_i^\infty)^2$ and let $\gamma_i(t) \in \Gamma^*(\rho_i(t), \rho_i^\infty)$. Using [56, Theorem 3.11] for the time derivative of the Wasserstein-2 metric, we have

$$\sum_{i=1}^n \frac{d}{dt} \mathcal{W}_2(\rho_i(t), \rho_i^\infty)^2 = 2 \sum_{i=1}^n \int \langle x_i - y_i, v_i[\rho(t)](x_i) \rangle d\gamma_i(t, x_i, y_i)$$

Since $v_i[\rho^\infty](y_i) = 0$ for ρ_i^∞ -a.e. y_i ,

$$\left| \int \langle x_i, v_i[\rho^\infty](y_i) \rangle d\gamma_i(t, x_i, y_i) \right| \leq \sqrt{\int \|x_i\|^2 d\rho_i(t, x_i) \int \|v_i[\rho^\infty](y_i)\|^2 d\rho_i^\infty(y_i)} = 0,$$

and so we can add this term to the time derivative:

$$\sum_{i=1}^n \frac{d}{dt} \mathcal{W}_2(\rho_i(t), \rho_i^\infty)^2 = 2 \sum_{i=1}^n \int \langle x_i - y_i, v_i[\rho(t)](x_i) - v_i[\rho^\infty](y_i) \rangle d\gamma(t, x, y).$$

Applying the monotonicity condition with the optimal plan, using that $\rho^\infty \in \mathcal{A}$, gives

$$\sum_{i=1}^n \frac{d}{dt} \mathcal{W}_2(\rho_i(t), \rho_i^\infty)^2 \leq -2\lambda \sum_{i=1}^n \mathcal{W}_2(\rho_i(t), \rho_i^\infty)^2,$$

which results in $\frac{d}{dt} \overline{W}(\rho(t), \rho^\infty)^2 \leq -2\lambda \overline{W}(\rho(t), \rho^\infty)^2$. Exponential convergence follows from Grönwall's inequality.

Part (c) follows from Lemma 3.24. □

3.4 Coupled Gradient Flow Setting

In this section, we present and discuss results for the setting where the dynamics are a coupled gradient flow. Monotonicity can be viewed as a *extension of displacement convexity to the multispecies setting*, where instead of having a condition on a single energy functional, the condition holds for the set of energies $F = [F_1, \dots, F_n]$. The notion of monotonicity also generalizes the results in [63], which proves exponential convergence for a two-species min-max flow with $F_1 = -F_2$. For a displacement convex-concave game $F = [F_1, F_2]$ with $F_1 = -F_2$, F is an example of a λ -monotone set of energies. Recall the dynamics (3.2), where each species i evolves in the direction of steepest descent for $F_i[\cdot, \rho_{-i}]$ in the $\mathcal{W}_2$ metric:

$$\partial_t \rho_i(t) = -\nabla_{\mathcal{W}_2, \rho_i} F_i[\rho(t)] = \operatorname{div}_{x_i} (\rho_i \nabla_{x_i} \delta_{\rho_i} F_i[\rho(t)]) , \quad \forall i \in [1, \dots, n] .$$

Recall also the definition of monotonicity for coupled gradient flows, Definition 3.1. Let $\mathcal{A}$ be a geodesically convex set. If for any $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$, it holds that

$$\lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2 \leq \int \left\langle x - y, \overline{\nabla \delta_\rho F}[\rho^{(0)}](x) - \overline{\nabla \delta_\rho F}[\rho^{(1)}](y) \right\rangle d\gamma(x, y) ,$$

for $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$, then the set of *energy functionals* $\{F_i : \mathcal{A} \rightarrow \mathbb{R}\}_{i=1}^n$ is λ -monotone in $\mathcal{A}$. In this section, we show how monotonicity implies convexity but convexity does not necessarily imply monotonicity, because monotonicity accounts for dependencies among species. Like convexity, monotonicity is additive, and we show how to relate lower-bounds on Hessians of cross-interaction kernels to the monotonicity coefficient λ . The results are summarized in the following properties:

- (C1) Monotonicity implies convexity: if F is λ -monotone in $\mathcal{A}$, then each $F_i[\cdot, \rho_{-i}]$ is λ -displacement convex in $\mathcal{A}_i$ for all $\rho_{-i} \in \mathcal{A}_{-i}$. See Lemma 3.28, Example 3.29 (section 3.4.2).
- (C2) Additivity of monotonicity: if F is λ -monotone in $\mathcal{A}$ and E is λ -monotone in $\mathcal{B}$, then $F + E$ is λ -monotone in $\mathcal{A} \cap \mathcal{B}$. See Lemma 3.30, Example 3.31 (section 3.4.3).
- (C3) Monotonicity of cross-interaction kernels: for $\mathcal{W}_{ij}$ defined as $\mathcal{W}_{ij} = \int W_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j)$, $F_i = \sum_{j \neq i} \mathcal{W}_{ij}$, let α_{ij} be the Lipschitz coefficient of $\nabla_{x_i} W_{ij}(x_i, \cdot)$ and $\sum_{i \neq j} \nabla_{x_i}^2 W_{ij}(x_i, x_j) \succeq c_i \mathbf{I}_{d_i}$. Define

$$\Lambda := \begin{bmatrix} c_1 & -\alpha_{12} & \dots & -\alpha_{1n} \\ -\alpha_{21} & c_2 & & -\alpha_{2n} \\ \vdots & & \ddots & \vdots \\ -\alpha_{n1} & -\alpha_{n2} & \dots & c_n \end{bmatrix} .$$

Then F is λ -monotone in $\mathcal{A}$, where $\frac{1}{2}(\Lambda + \Lambda^\top) \succeq \lambda \mathbf{I}_d$. See Lemma 3.32, Example 3.34 (section 3.4.4).

To prove these properties, we use an equivalent, second-order notion of monotonicity. The second-order definition parallels the finite-dimension notion; we first present the second-order finite-dimension definition with the infinite-dimensional version following. In a finite-dimensional game in $\mathbb{R}^d$, each player selects a deterministic vector $x_i \in \mathbb{R}^{d_i}$ minimizing a cost function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, which depends on the selections of all other players. When each player does gradient descent on its cost function, the dynamics are $\dot{x}_i = -\nabla_{x_i} f_i(x)$. The definition of monotonicity for a finite-dimensional game is given in terms of the vector of gradients.

Definition 3.25 (Finite-Dimensional Monotonicity). *Let $f = [f_1, \dots, f_n] : \mathbb{R}^d \rightarrow \mathbb{R}^n$ where $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, and define $\bar{\nabla} f(x) := [\nabla_{x_1} f_1(x), \dots, \nabla_{x_n} f_n(x)] \in \mathbb{R}^d$ and $\bar{\nabla}^2 f(x) \in \mathbb{R}^{d \times d}$ such that $(\bar{\nabla}^2 f(x))_{ij} = \nabla_{x_j x_i}^2 f_i(x) \in \mathbb{R}^{d_i \times d_j}$. A game f is monotone if it satisfies*

$$\text{(first order)} \quad \left\langle \bar{\nabla} f(x) - \bar{\nabla} f(x'), x - x' \right\rangle \geq \lambda \|x - x'\|^2 \quad \forall x, x' \in \mathbb{R}^d, \text{ or}$$

$$\text{(second order)} \quad \bar{\nabla}^2 f(x) + \bar{\nabla}^2 f(x)^\top \succeq 2\lambda I_d \quad \forall x \in \mathbb{R}^d,$$

for some $\lambda \in \mathbb{R}$. The game is strongly monotone if $\lambda > 0$.

Importantly, $\bar{\nabla}$ is *not* a gradient; it is an entry-wise operator. In order to present the corresponding infinite-dimensional result, we define Hessian-like operators $\overline{\nabla^2 \delta_\rho}$ and $\overline{\nabla^2 \delta_\rho^2}$.

Definition 3.26 (Matrix Operators). *For sufficiently regular F , the operators $\overline{\nabla^2 \delta_\rho}$ and $\overline{\nabla^2 \delta_\rho^2}$ are defined as follows, where $(\cdot)_{ij}$ denotes the i^{th} row block and j^{th} column block:*

$$\begin{aligned} \left(\overline{\nabla^2 \delta_\rho} F[\rho](x) \right)_{ij} &:= \begin{cases} \nabla_{x_i}^2 \delta_{\rho_i} F_i[\rho](x_i) & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \in \mathbb{R}^{d_i \times d_i}, \\ \left(\overline{\nabla^2 \delta_\rho^2} F[\rho](x, \hat{x}) \right)_{ij} &:= \nabla_{x_i \hat{x}_j}^2 \delta_{\rho_i, \rho_j}^2 F_i[\rho](x_i, \hat{x}_j) \in \mathbb{R}^{d_i \times d_j}. \end{aligned}$$

Proposition 3.27 (Reformulation of Monotonicity). *Let Assumption 3.13 hold. Then λ -monotonicity in $\mathcal{A}$ holds if and only if for any two $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$,*

$$\begin{aligned} &\int_0^1 \left(\iint \left\langle x - y, \overline{\nabla^2 \delta_\rho^2} F[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) \cdot (\hat{x} - \hat{y}) \right\rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) \right. \\ &\quad \left. + \int \left\langle x - y, \overline{\nabla^2 \delta_\rho} F[\rho^{(s)}](I^{(s)}(x, y)) \cdot (x - y) \right\rangle d\gamma(x, y) \right) ds \geq \lambda \bar{W}(\rho^{(0)}, \rho^{(1)})^2, \end{aligned}$$

where $\rho^{(s)}$ is the displacement interpolant between $\rho^{(0)}$ and $\rho^{(1)}$, and $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$.

If additionally Assumptions 3.14 and 3.15 hold, then $F : \mathcal{A} \rightarrow \mathbb{R}^n$ is λ -monotone in $\mathcal{A}$ if and only if

$$\begin{aligned} &\iint \left\langle w(x), \overline{\nabla^2 \delta_\rho^2} F[\rho](x, \hat{x}) \cdot w(\hat{x}) \right\rangle d\rho(x) d\rho(\hat{x}) \\ &\quad + \int \left\langle w(x), \overline{\nabla^2 \delta_\rho} F[\rho](x) \cdot w(x) \right\rangle d\rho(x) \geq \lambda \int \|w(x)\|^2 d\rho(x). \end{aligned}$$

for all $\rho \in \mathcal{A}$, $w_i \in Y_i(\rho_i, \mathcal{A}_i)$ for all $i \in [1, \dots, n]$.

For the proof, see section 3.A.1. The second reformulation of monotonicity above provides a characterization that depends on only a single element $\rho \in \mathcal{A}$ and a velocity field set $\{w_i \in Y_i(\rho_i, \mathcal{A}_i)\}_i$, instead of a pair $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$. This can therefore be viewed as a "local" notion of monotonicity; however, it does require sufficient regularity for the second order derivatives of the energies.

The key difference between convexity and monotonicity for both the Euclidean ($\mathbb{R}^d$) and measure $\mathcal{P}(\mathbb{R}^d)$ settings is that for monotonicity, the terms in the inner product are no longer gradients of a single object; they are player-specific gradients. The distance is the Euclidean distance $\|\cdot\|$ or the $\overline{W}$ distance, and the inner product is either the Euclidean inner product or the integral of the inner product weighted by a measure. For the second-order conditions in Proposition 3.27, similar to the first-order case, the gradients for monotonicity are dependent on all of the cost functions and in general cannot be written as the Hessian of a single cost function. Notably, there are two matrices for the $\mathcal{P}_2(\mathbb{R}^d)$ monotonicity setting whereas only one for the Euclidean monotone setting; this is due to the fact that the gradient of the Euclidean metric tensor is zero while the gradient of the Wasserstein-2 metric tensor is not. For more details and a summary of the relationship between monotonicity and convexity in finite and infinite dimensions, see section 3.A.2.

In section 3.4.1, we present an example illustrating how monotonicity properties assist in satisfying the requirements for Theorem 3.17, resulting in convergence of a coupled gradient flow system to a unique Nash equilibrium. In the next subsections, we present lemmas and examples for each of the properties (C1)-(C3).

3.4.1 Example: Two Species Zero-Sum Game with Nonlinear Diffusion

In this example, we consider the two-species setting with $F = [F_1, -F_1]$, where

$$F_1[\rho_1, \rho_2] = \iint f(x_1, x_2) d\rho_1(x_1) d\rho_2(x_2) + \alpha_1 \int U_{m_1}(\rho_1) dx_1 - \alpha_2 \int U_{m_2}(\rho_2) dx_2$$

for some $\alpha_1, \alpha_2 > 0$, with $f \in C^2(\mathbb{R}^{d_1} \times \mathbb{R}^{d_2}; \mathbb{R})$ such that $[f, -f]$ satisfies finite-dimensional monotonicity (see Definition 3.25) for some $\lambda > 0$. For a measure $\eta \in \mathcal{P}(\mathbb{R}^k)$, the nonlinear diffusion term $U_m(\eta) : \mathcal{P}(\mathbb{R}^k) \rightarrow \{\mathbb{R} \cup +\infty\}$ is defined as

$$U_m(\eta) = \begin{cases} \eta \log \eta & \text{for } m = 1, \eta \ll \mathcal{L}^k, \\ \frac{\eta^m}{m-1} & \text{for } m > 1, \eta \ll \mathcal{L}^k, \\ +\infty & \text{else.} \end{cases} \quad (3.14)$$

For $\rho \in \overline{L}^1 := L^1(\mathbb{R}^{d_1}) \times \cdots \times L^1(\mathbb{R}^{d_n})$, the corresponding flow (3.1) has velocity field $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho] = w_i[\rho] - \alpha_i \nabla_{x_i} U'_{m_i}(\rho_i)$, where

$$w[\rho] := \left[-\int \nabla_{x_1} f(x_1, x_2) d\rho_2(x_2), \int \nabla_{x_2} f(x_1, x_2) d\rho_1(x_1) \right].$$

We will use (C1)-(C2) to show that the game F is λ -monotone in $\bar{L}^1 \cap \bar{\mathcal{P}}_2$ in the sense of Definition 3.1. Let $\mathcal{A}_i = L^1(\mathbb{R}^{d_i}) \cap \mathcal{P}_2(\mathbb{R}^{d_i})$ and $\mathcal{B}_i = \mathcal{P}_2(\mathbb{R}^{d_i})$. Since $[f, -f]$ satisfies λ -monotonicity in finite dimensions (Definition 3.25), we have that w is λ -monotone in $\mathcal{B} = \bar{\mathcal{P}}_2$ by Proposition 3.5, so

$$E := \left[\iint f(x_1, x_2) d\rho_1(x_2) d\rho_2(x_2), - \iint f(x_1, x_2) d\rho_1(x_1) d\rho_2(x_2) \right]$$

is λ -monotone in $\mathcal{B}$. Using that $\int U_{m_i}(\rho_i) dx_i$ is 0-convex in L^1 [50, Section 3.1], it follows from Lemma 3.28(b) (C1, monotonicity vs convexity) that $\tilde{E}$ with $\tilde{E}_i := \alpha_i \int U_{m_i}(\rho_i) dx_i$ is 0-monotone in $\mathcal{A}$. Then applying Lemma 3.30 (C2, additivity), we conclude that $F = E + \tilde{E}$ is λ -monotone in $\mathcal{A} = \bar{\mathcal{P}}_2 \cap \bar{L}^1$.

To illustrate that the theoretical results on second moment control as well as existence of and convergence to a unique Nash equilibrium apply to this example, we assume from now on that Assumption 3.11 holds. First, we show that the second moments of solutions to (3.2) defined by the game F are uniformly bounded. Assuming $\rho(t) \in \mathcal{A} = \bar{L}^1 \cap \bar{\mathcal{P}}_2$ for all $t \geq 0$, Lemma 3.21 results in a uniform bound on the second moments.

Next, we establish that Assumption 3.12 holds. The finite-dimensional monotonicity of $[f, -f]$ with $\lambda > 0$ implies $f(\cdot, x_2)$ is strongly convex for any $x_2 \in \mathbb{R}^{d_2}$ and $f(x_1, \cdot)$ is strongly concave for any $x_1 \in \mathbb{R}^{d_1}$; in combination with $f \in C^2$, we obtain $\int f(\cdot, x_2) d\rho_2(x_2) \geq c_2$, $\int f(x_1, \cdot) d\rho_1(x_1) \leq c_1$ where c_1, c_2 depend on ρ_1, ρ_2 respectively. Together with continuity of these terms in x_1 and x_2 respectively, we have lower semicontinuity of $\iint f(x_1, x_2) d\rho_1(x_1) d\rho_2(x_2)$ with respect to ρ_1 for any fixed ρ_2 and upper semicontinuity of $\iint f(x_1, x_2) d\rho_1(x_1) d\rho_2(x_2)$ with respect to ρ_2 for any fixed ρ_1 , due to [194, Proposition 7.1]. For the diffusion term, if $m_i > 1$, the L^{m_i} norm is lower semicontinuous [194, Proposition 7.7, Remark 7.8]. For the case $m_i = 1$, lower semicontinuity of the diffusion term follows from the uniform second moment bound and [63, Lemma C.4]. Therefore Assumption 3.12 holds.

Using Assumption 3.11 and applying Theorem 3.17(a), there exists a unique steady state in $\rho^\infty \in S$. In [63] it was shown that $\rho^\infty \in \bar{L}^1 \cap \bar{\mathcal{P}}_2$ as long as $\alpha_1, \alpha_2 > 0$ in the case of linear diffusion, and the result for $m_i > 1$ follows in the same way by considering the Euler-Lagrange condition

$$\frac{\alpha_i m_i}{m_i - 1} \rho_i^{m_i-1}(x_i) + (-1)^{i+1} \int f(x_1, x_2) d\rho_{-i}(x_{-i}) = c_i \quad \text{on } \text{supp } \rho_i.$$

Then, since $\rho^\infty \in \mathcal{A} \cap S$, by Theorem 3.17(b), we have exponential convergence in $\bar{W}$ with rate λ . And moreover, thanks to Assumption 3.12, we can apply Theorem 3.17(c) to conclude that ρ^∞ is the unique Nash equilibrium.

3.4.2 Monotonicity Implies Convexity

Here we state the lemma which describes property (C1). While monotonicity implies convexity, the converse is not true, as illustrated in the example in this section.

Lemma 3.28 (Monotonicity and Convexity). *The following statements detail the relationship between monotonicity and convexity.*

- (a) Monotonicity implies convexity. *Let $F : \mathcal{A} \rightarrow \mathbb{R}^n$ be λ -monotone in $\mathcal{A}$. Then $F_i[\cdot, \rho_{-i}]$ is λ -displacement convex for all fixed $\rho_{-i} \in \mathcal{A}_{-i}$, for all $i \in [1, \dots, n]$.*
- (b) Uncoupled convexity implies monotonicity. *Let $E_i : \mathcal{A}_i \rightarrow \mathbb{R}$ be a λ_i -displacement convex functional for all $\rho_i \in \mathcal{A}_i$, for all $i \in [1, \dots, n]$. Then the functional $E = [E_1, \dots, E_n] : \mathcal{A} \rightarrow \mathbb{R}^n$ is λ -monotone in $\mathcal{A}$, where $\lambda := \min_i \lambda_i$.*

Proof of Lemma 3.28. We begin with the proof of (a). To prove that monotonicity implies convexity, select any $\pi \in \mathcal{A}$. Selecting a particular value of i , set $\rho_j^{(0)} = \rho_j^{(1)} = \pi_j$ for all $j \neq i$. Let $\rho_i^{(0)} = \pi_i$ and $\rho_i^{(1)}$ satisfy $\mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2 > 0$. First-order monotonicity results in

$$\int \langle x_i - y_i, \nabla_{x_i} \delta_{\rho_i} F_i[\rho^{(0)}](x_i) - \nabla_{x_i} \delta_{\rho_i} F_i[\rho^{(1)}](y_i) \rangle d\gamma_i(x_i, y_i) \geq \lambda \mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2,$$

which is a sufficient condition for λ -displacement convexity in $\mathcal{A}$. See [203, Proposition 16.2] for details on convexity inequalities.

Next we prove (b). Since each E_i is λ_i -displacement convex for all $\rho_i \in \mathcal{A}_i$, for all $i \in [1, \dots, n]$ it holds that

$$\int \langle x_i - y_i, \nabla_{x_i} \delta_{\rho_i} E_i[\rho_i^{(0)}](x_i) - \nabla_{x_i} \delta_{\rho_i} E_i[\rho_i^{(1)}](y_i) \rangle d\gamma_i(x_i, y_i) \geq \lambda_i \mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2,$$

for all $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ and $\gamma_i \in \Gamma^*(\rho_i^{(0)}, \rho_i^{(1)})$. Since there is no coupling among species, summing over all i results in

$$\int \langle \overline{\nabla \delta_\rho} E[\rho^{(0)}](x) - \overline{\nabla \delta_\rho} E[\rho^{(1)}](y), x - y \rangle d\gamma(x, y) \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2,$$

where $\lambda = \min_i \lambda_i$. The above is the monotonicity condition. □

In the following example, we illustrate a setting in which $F_i[\cdot, \rho_{-i}]$ is displacement convex for all i , but F is not monotone due to the nature of the coupling between the two species, showing why monotonicity is a stronger condition than convexity.

Example 3.29 (Convexity Does Not Imply Monotonicity). Consider a system $F : \overline{\mathcal{P}_2} \rightarrow \mathbb{R}^2$ given by

$$\begin{aligned} F_1[\rho_1, \rho_2] &= \frac{\lambda}{2} \int \|x_1\|^2 d\rho_1(x_1) + \iint x_1^\top A_1 x_2 d\rho_1(x_1) d\rho_2(x_2), \\ F_2[\rho_1, \rho_2] &= \frac{\lambda}{2} \int \|x_2\|^2 d\rho_2(x_2) + \iint x_1^\top A_2 x_2 d\rho_1(x_1) d\rho_2(x_2), \end{aligned}$$

where $A_1, A_2 \in \mathbb{R}^{d_1 \times d_2}$ and $\lambda > 0$. The second derivative of F_i along geodesics between $\rho^{(0)}, \rho^{(1)} \in \overline{\mathcal{P}_2}$ is

$$\frac{d^2}{ds^2} F_i[\rho_i^{(s)}, \rho_{-i}] = \lambda \int \|v_i^{(s)}(x_i)\|^2 d\rho_i^{(s)}(x_i) = \lambda W_2(\rho_i^{(0)}, \rho_i^{(1)})^2, \quad \forall i \in \{1, 2\},$$

where $(\rho_i^{(s)}, v_i^{(s)})$ solve (3.6), and therefore $F_1[\cdot, \rho_2]$ and $F_2[\rho_1, \cdot]$ are λ -displacement convex. The first-order monotonicity condition, where $A = \begin{bmatrix} 0 & A_1 \\ A_2^\top & 0 \end{bmatrix}$, is

$$\begin{aligned} & \int \left\langle x - y, \overline{\nabla}_{\delta_\rho} F[\rho^{(0)}](x) - \overline{\nabla}_{\delta_\rho} F[\rho^{(1)}](y) \right\rangle d\gamma(x, y) \\ &= \int \left\langle x - y, \lambda(x - y) + \int A(\hat{x} - \hat{y}) d\rho^{(0)}(\hat{x}) d\rho^{(1)}(\hat{y}) \right\rangle d\gamma(x, y) \\ &= \lambda \int \|x - y\|^2 d\gamma(x, y) + \mathbb{E}(x - y)^\top A \mathbb{E}(x - y), \end{aligned}$$

where $\mathbb{E}(x - y) = \int (x - y) d\rho^{(0)}(x) d\rho^{(1)}(y)$. Then for choices of A_1, A_2 such that $A + A^\top$ has one or more negative eigenvalues, F does not satisfy the monotonicity condition. However, selecting $A_1 = -A_2$ results in λ -monotonicity in $\overline{\mathcal{P}_2}$, since $A + A^\top = 0$.

3.4.3 Additivity of Monotonicity

Here we state the lemma which describes property (C2). This property means that the sum of two monotone systems is monotone, with the monotonicity coefficient at least the sum of the original coefficients.

Lemma 3.30 (Additivity of Monotonicity). Let $E = [E_1, \dots, E_n] : \mathcal{A} \rightarrow \mathbb{R}^n$ be λ -monotone in $\mathcal{A}$ and $F = [F_1, \dots, F_n] : \mathcal{B} \rightarrow \mathbb{R}^n$ be $\tilde{\lambda}$ -monotone in $\mathcal{B}$, and define $\hat{F} = [\hat{F}_1, \dots, \hat{F}_n] : \mathcal{A} \cap \mathcal{B} \rightarrow \mathbb{R}^n$ where $\hat{F}_i = \alpha E_i + \tilde{\alpha} F_i$, for some $\alpha, \tilde{\alpha} \in \mathbb{R}$. Then $\hat{F}$ is $\alpha\lambda + \tilde{\alpha}\tilde{\lambda}$ -monotone in $\mathcal{A} \cap \mathcal{B}$.

Proof. The proof follows by the linearity of the monotonicity condition in terms of the energy functionals. $\square$

The result from Lemma 3.28(b) shows that displacement convexity directly implies monotonicity when there is no coupling among different species. Lemma 3.30 is useful for computing monotonicity

of more intricate functionals because one can verify the monotonicity of only the inter-species terms to ensure overall monotonicity. For terms such as $\int \rho_i \log \rho_i$ which do not have any multispecies coupling and for which displacement convexity has been shown, one can instead check the monotonicity of F without these terms, and apply the previous lemmas, which allow for the addition of displacement convex terms with dependence on only one species to be added to monotone F , and the resulting game is monotone with an updated parameter.

Example 3.31 (Pairwise Zero-Sum Game with Nonlinear Diffusion: Application of Lemma 3.30).

In this example, let $\mathcal{A}_i = L^{m_i}(\mathbb{R}^{d_i}) \cap \mathcal{P}_2(\mathbb{R}^{d_i})$ with $m_i \geq 1$. Let $\hat{F}_i : \mathcal{A} \rightarrow \mathbb{R}$ take the form

$$\hat{F}_i[\rho] = \sum_{j \neq i} \iint f_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j) + \int V_i(x_i) d\rho_i(x_i) + \alpha_i \int U_{m_i}(\rho_i) dx_i$$

for $\alpha_i \geq 0$. Monotonicity for $\hat{F}$ is not defined in $\overline{\mathcal{P}_2}$ because of the nonlinear diffusion terms, for which the left-hand side of (3.4) is not defined on Diracs, but we will show that $\hat{F}$ is monotone in $\mathcal{A} \subset \overline{\mathcal{P}_2}$. Let $f_{ij} \in C^2(\mathbb{R}^{d_i} \times \mathbb{R}^{d_j}, \mathbb{R})$ satisfy $-f_{ij}(x_i, x_j) = f_{ji}(x_j, x_i)$, $\tilde{\lambda}_i \mathbf{I}_{d_i} \preceq \sum_{j \neq i} \nabla_{x_i}^2 f_{ij}(x_i, x_j)$ for all $x_i \in \mathbb{R}^{d_i}$, $x_j \in \mathbb{R}^{d_j}$, and $\int \sum_{j \neq i} \nabla_{x_i}^2 f_{ij}(x_i, x_j) d\rho_{-i}(x_{-i}) \in \mathbb{R}^{d_i \times d_i}$ for all $\rho_{-i} \in \mathcal{A}_{-i}$ and $x_i \in \mathbb{R}^{d_i}$. Let $V_i \in C^2(\mathbb{R}^{d_i}, \mathbb{R})$ satisfy $\nabla_{x_i}^2 V_i(x_i) \succeq \lambda_i \mathbf{I}_{d_i}$ for all $x_i \in \mathbb{R}^{d_i}$. Consider the set of functionals

$$F_i[\rho] = \sum_{j \neq i} \iint f_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j).$$

To check the reformulated monotonicity condition in Proposition 3.27, we compute the Hessian terms introduced in Definition 3.26:

$$\begin{aligned} (\overline{\nabla^2 \delta_\rho} F[\rho](x))_{ii} &= \sum_{j \neq i} \int \nabla_{x_i}^2 f_{ij}(x_i, x_j) d\rho_j(x_j) = \int \sum_{j \neq i} \nabla_{x_i}^2 f_{ij}(x_i, x_j) d\rho_{-i}(x_{-i}) \succeq \tilde{\lambda}_i \mathbf{I}_{d_i}, \\ (\overline{\nabla^2 \delta_\rho^2} F[\rho](x, \hat{x}))_{ij} &= \begin{cases} 0 & \text{if } i = j, \\ \nabla_{x_i \hat{x}_j}^2 f_{ij}(x_i, \hat{x}_j) & \text{if } i \neq j. \end{cases} \end{aligned}$$

Using a change of notation and $-f_{ij}(x_i, x_j) = f_{ji}(x_j, x_i)$,

$$\begin{aligned} & \int_0^1 \iint \left\langle x - y, \overline{\nabla^2 \delta_\rho^2} F[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) \cdot (\hat{x} - \hat{y}) \right\rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) ds \\ &= \frac{1}{2} \int_0^1 \iint \left\langle x - y, \left[(\overline{\nabla^2 \delta_\rho^2} F[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) + \overline{\nabla^2 \delta_\rho^2} F[\rho](I^{(s)}(\hat{x}, \hat{y}), I^{(s)}(x, y))^\top \right] \right. \\ & \quad \left. \cdot (\hat{x} - \hat{y}) \right\rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) ds \\ &= 0. \end{aligned}$$

The matrices J_1 and J_2 , given by

$$(J_1[\rho](x))_{ii} = -(\overline{\nabla^2 \delta_\rho} F[\rho](x))_{ii} = - \int \sum_{j \neq i} \nabla_{x_i}^2 f_{ij}(x_i, x_j) d\rho_{-i}(x_{-i}),$$

$$(J_2[\rho](x, \hat{x}))_{ij} = -(\overline{\nabla^2 \delta_\rho^2} F[\rho](x, \hat{x}))_{ij} = \begin{cases} 0 & \text{if } i = j, \\ -\nabla_{x_i \hat{x}_j}^2 f_{ij}(x_i, \hat{x}_j) & \text{if } i \neq j, \end{cases}$$

are well-defined ρ -a.e. due to $f_{ij} \in C^2$ and $|\int \sum_{j \neq i} \nabla_{x_i}^2 f_{ij}(x_i, x_j) d\rho_{-i}(x_{-i})| < \infty$, and so Assumption 3.13 holds. Applying Proposition 3.27 we conclude that F is $\tilde{\lambda}$ -monotone in $\overline{\mathcal{P}_2}$ with $\tilde{\lambda} := \min_i \tilde{\lambda}_i$. Further, let $\lambda = \min_i \lambda_i$. We can apply Lemma 3.30 for each player, using that $\int U_{m_i}(\rho_i) dx_i$ is 0-displacement convex [50] over $L^{m_i}(\mathbb{R}^{d_i})$ and $\int V_i(x_i) d\rho_i(x_i)$ is λ_i -displacement convex because $\nabla_{x_i}^2 V_i \succeq \lambda_i \mathbf{I}_{d_i}$, to obtain that $\hat{F}$ is $(\lambda + \tilde{\lambda})$ -monotone in $\mathcal{A}$, where $\lambda = \min_i \lambda_i$. Hence, Lemma 3.30 provides a method for using knowledge on the displacement-convexity of additive functionals to show monotonicity¹.

3.4.4 Monotonicity of Cross-Interaction Kernels

Next, we investigate monotonicity for a specific type of commonly-used functionals. In the PDE literature for biological models, interaction kernels are commonly employed to capture phenomena of interacting species. Our monotonicity condition provides convergence guarantees for systems driven by interaction kernels that have lower-bounded Hessians. To illustrate this, we draw a connection to the finite-dimensional contractivity literature; [75, Theorem 41] proves that when the convexity of cost functions dominate inter-species effects, the overall system is contractive. We extend this results the infinite-dimensional setting in the following lemma.

Lemma 3.32 (Cross-Interaction Kernels). *Consider the interaction term $\mathcal{W}_{ij} : \mathcal{A}_i \times \mathcal{A}_j \rightarrow \mathbb{R}$, parameterized by the interaction kernel $W_{ij} \in C^2(\mathbb{R}^{d_i} \times \mathbb{R}^{d_j}, \mathbb{R})$, given by*

$$\mathcal{W}_{ij}(\rho_i, \rho_j) = \iint W_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j).$$

Let $F_i = \sum_{j \neq i} \mathcal{W}_{ij}$ for all $i \in [1, \dots, n]$. If for all $i \neq j$, there exist $c_i \in \mathbb{R}$ and $\alpha_{ij} \geq 0$ such that for all $x_j, y_j \in \mathbb{R}^{d_j}$, $x_i \in \mathbb{R}^{d_i}$,

$$\|\nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(x_i, y_j)\| \leq \alpha_{ij} \|x_j - y_j\|, \quad \sum_{j \neq i} \nabla_{x_i}^2 W_{ij}(x_i, x_j) \succeq c_i \mathbf{I}_{d_i},$$

¹A potentially larger monotonicity coefficient $\min_i(\lambda_i + \tilde{\lambda}_i)$ can be obtained by including V_i in F_i instead of considering it separately via Lemma 3.30.

and the matrix

$$\Lambda := \begin{bmatrix} c_1 & -\alpha_{12} & \dots & -\alpha_{1n} \\ -\alpha_{21} & c_2 & & -\alpha_{2n} \\ \vdots & & \ddots & \vdots \\ -\alpha_{n1} & -\alpha_{n2} & \dots & c_n \end{bmatrix},$$

satisfies $\frac{1}{2}(\Lambda + \Lambda^\top) \succeq \lambda \mathbf{I}_d$, then F is λ -monotone in $\mathcal{A}$.

Proof. The monotonicity condition to check is: for all $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ and for all $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$,

$$\begin{aligned} \sum_{i,j \neq i} \int \left\langle \int \nabla_{x_i} W_{ij}(x_i, x_j) d\rho_j^{(0)}(x_j) - \int \nabla_{x_i} W_{ij}(y_i, y_j) d\rho_j^{(1)}(y_j), x_i - y_i \right\rangle d\gamma_i(x_i, y_i) \\ \geq \lambda \int \|x - y\|^2 d\gamma(x, y). \end{aligned}$$

We use the optimal plans to write the inner product as

$$\begin{aligned} \sum_{i,j \neq i} \int \left\langle \int \nabla_{x_i} W_{ij}(x_i, x_j) d\rho_j^{(0)}(x_j) - \int \nabla_{x_i} W_{ij}(y_i, y_j) d\rho_j^{(1)}(y_j), x_i - y_i \right\rangle d\gamma_i(x_i, y_i) \\ = \int \sum_{i,j \neq i} \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \rangle d\gamma(x, y), \end{aligned}$$

and apply the following exact Taylor expansion

$$\nabla_{x_i} W_{ij}(y_i, y_j) = \nabla_{x_i} W_{ij}(x_i, y_j) + (y_i - x_i)^\top \nabla_{x_i}^2 W_{ij}(\tilde{x}_i, y_j)$$

for some $\tilde{x}_i \in \mathbb{R}^{d_i}$. Then the integrand satisfies

$$\begin{aligned} \sum_{i,j \neq i} \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \rangle &= \sum_i (x_i - y_i)^\top \left(\sum_{j \neq i} \nabla_{x_i}^2 W_{ij}(\tilde{x}_i, y_j) \right) (x_i - y_i) \\ &\quad + \sum_{i,j \neq i} \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(x_i, y_j), x_i - y_i \rangle. \end{aligned}$$

Applying $\sum_{j \neq i} \nabla_{x_i}^2 W_{ij}(\tilde{x}_i, y_j) \succeq c_i \mathbf{I}_{d_i}$ for the first term and the Cauchy-Schwarz inequality plus the Lipschitz condition for the second term, we have

$$\begin{aligned} \sum_{i,j \neq i} \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \rangle \\ \geq \sum_i c_i \|x_i - y_i\|^2 - \sum_{i,j \neq i} \alpha_{ij} \|x_j - y_j\| \|x_i - y_i\| = z^\top \Lambda z = \frac{1}{2} z^\top (\Lambda + \Lambda^\top) z, \end{aligned} \tag{3.15}$$

where $z = z(x, y) \in \mathbb{R}^n$, $z_i := \|x_i - y_i\|$, and Λ is defined in the lemma statement. Plugging (3.15) into the monotonicity condition, we conclude

$$\begin{aligned} & \int \sum_{i,j \neq i} \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \rangle d\gamma(x, y) \\ & \geq \frac{1}{2} \int z(x, y)^\top (\Lambda + \Lambda^\top) z(x, y) d\gamma(x, y) \\ & \geq \lambda \int \|x - y\|^2 d\gamma(x, y). \end{aligned}$$

□

Remark 3.33. *The previous result can be adjusted to include nonlocal self-interaction terms $\iint W_i(x_i, \tilde{x}_i) d\rho_i(x_i) d\rho_i(\tilde{x}_i)$, using existing displacement convexity properties and the additivity property (Lemma 3.30). Moreover, if $-\Lambda$ is Hurwitz (i.e. all its eigenvalues have strictly negative real parts), even if $\Lambda + \Lambda^\top \not\leq 0$, we expect contraction in a reweighted joint $\mathcal{W}_2$ metric following the approach in [75, Thm 41].*

Lemma 3.32 shows that if the interaction kernels are strongly convex with large convexity coefficients c_i and small Lipschitz constants α_{ij} , which encodes how strongly the other species influence species i , then the overall system is monotone. For more general conditions that ensure monotonicity of F , see Lemma 3.39.

Example 3.34. *Examples of kernels with lower-bounded Hessians include the quadratic and Morse-like attractive-repulsive kernels, which are radially symmetric. In the setting where $\mathbb{R}^{d_i} = \mathbb{R}^{d_j}$, let $x = x_i - x_j$ and $W_{ij}(x_i, x_j) = W(x_i - x_j) = W(x)$ for any i, j ; some commonly-used kernels are*

$$\begin{aligned} W^{(k)}(x) &= \frac{\|x\|^k}{k}, \quad W^{(p)}(x) = \frac{\|x\|^a}{a} - \frac{\|x\|^b}{b}, \\ W^{(m)}(x) &= C_r \exp\left(-\frac{\|x\|^2}{l_r}\right) - C_a \exp\left(-\frac{\|x\|^2}{l_a}\right), \end{aligned}$$

where $k > 1 - d_i$, $l_a > l_r > 0$, $C_r > C_a > 0$ and $a > b > \max(2 - d_i, 2)$ with a, b even integers so that the second derivative is well-defined (see [15, 67] and references therein). The kernel $W^{(k)}$ is purely attractive for any power k , and satisfies the Lipschitz condition in Lemma 3.32 for $k = 2$. The kernel $W^{(m)}$ is a Morse-like attractive-repulsive potential which is repulsive for small x and attractive for large x ; this kernel also satisfies the Lipschitz condition. The power law potential $W^{(p)}$ is attractive-repulsive; as it does not satisfy the Lipschitz condition in Lemma 3.32, monotonicity is an open question, and may be tackled making use of the local Lipschitzness of this kernel. See the computations in section 3.B for details on the Hessian lower bounds of these kernels, and see [54]

for more details and references regarding these particular kernels. When the relevant conditions are satisfied, Lemma 3.32 can be applied to coupled PDEs such as

$$\begin{aligned}\partial_t \rho_1 &= \operatorname{div} (\rho_1 \nabla (H_1 * \rho_1 + K_1 * \rho_2)) \\ \partial_t \rho_2 &= \operatorname{div} (\rho_2 \nabla (H_2 * \rho_2 + K_2 * \rho_1)) ,\end{aligned}$$

for which existence of solutions was studied in [112, 102, 54], motivated by applications in biological systems.

3.5 Applications

In this section, we illustrate how our model can be applied to the service provider - population game, drawing on models from [198]. Then we draw connections between a n -player monotone game and a particular set of mean field games. Finally, we show a numeric example of a two-population game which is not monotone, but almost monotone, and we use diffusion to ensure sufficient probability mass over the monotone region to still observe convergence. Numerical simulations for these examples are implemented using the two-dimensional finite volume method analyzed in [47].

3.5.1 Multi-Learner System

In this example, we illustrate monotonicity results on a multi-learner setting with strategic users. In this setting, users in population $i \in [1, \dots, N]$ choose to interact with service provider $j \in [1, \dots, M]$, for a total of $n = N + M$ species in the system. Both the user populations and service providers update their parameters over time, evolving according to a gradient flow with respect to their own state. The user populations are denoted as in the rest of the paper by $\rho_i \in \mathcal{P}(\mathbb{R}^{d_\rho})$, and the service providers are assumed to be Diracs on the space of measures, that is, $\rho_{N+j} := \delta_{h_j}$, with $h_j \in \mathbb{R}^{d_h}$ for $j \in [1, \dots, M]$. The user populations have utility functions associated with each service provider, denoted $u_{ij} : \mathbb{R}^{d_\rho} \times \mathbb{R}^{d_h} \rightarrow \mathbb{R}$. The proportion of population i which purchases from service provider j is $a_{ij} \in [0, 1]$, such that $\sum_{j=1}^M a_{ij} = 1$ for all $i \in [1, \dots, N]$. The service provider j has a loss function for interacting with population i given by $l_{ij} : \mathbb{R}^{d_\rho} \times \mathbb{R}^{d_h} \rightarrow \mathbb{R}$. Additionally, each population nonlocally interacts with the other populations. The resulting energy functionals are, for all $i \in [1, \dots, n]$ and $j \in [1, \dots, M]$,

$$\begin{aligned}\text{user populations: } F_i[\rho, h] &= - \sum_{j=1}^M \int a_{ij} u_{ij}(x_i, h_j) d\rho_i(x_i) + \sum_{k=1}^N \int \rho_i W_{ik} * \rho_k + \alpha_i KL(\rho_i | \tilde{\rho}_i) \\ \text{service providers: } F_{N+j}[\rho, h] &= \sum_{i=1}^N \int a_{ij} l_{ij}(x_i, h_j) d\rho_i(x_i) + \beta_j \|h_j\|^2 ,\end{aligned}$$

where $\tilde{\rho}_i$ are given log-concave reference measures and $\alpha_i, \beta_j > 0$. The allocation weights a_{ij} update over time; populations will seek services from providers with the best utility. We select the

weight update from [198], given by

$$a_{ij}(t) := \frac{\exp[\eta \int u_{ij}(x_i, h_j(t)) d\rho_i(t, x_i)]}{\sum_{m=1}^M \exp[\eta \int u_{im}(x_i, h_m(t)) d\rho_i(t, x_i)]},$$

where $\eta \in \mathbb{R}_+$ parameterizes the weight update steepness. The system evolves according to (3.2). The system is monotone depending on the specific choice of u_{ij} , l_{ij} , W_{ik} , α_i , and β_j . In the setting where the finite-dimensional game

$$\begin{aligned} & - \left[\sum_{j=1}^M \int a_{1j} u_{1j}(x_1, h_j), \dots, \sum_{j=1}^M \int a_{Nj} u_{Nj}(x_N, h_j), \right. \\ & \quad \left. \sum_{i=1}^N \int a_{i1} l_{i1}(x_i, h_1), \dots, \sum_{i=1}^N \int a_{iM} l_{iM}(x_i, h_M) \right] \end{aligned}$$

is λ -monotone over Euclidean space, with $\lambda \geq 0$, and $W_{ik} = 0$ with $\alpha_i, \beta_j > 0$, the infinite-dimensional game F is monotone over $L^1(\tilde{\rho}_i)^{\otimes n}$ by Proposition 3.5, Lemma 3.28(b) and Lemma 3.30. However, if the kernel W_{ik} is, for example, a Morse-like kernel which is locally repulsive and has a negative convexity parameter, monotonicity depends on the condition in Lemma 3.32, in combination with sufficient monotonicity from u , l , and the regularizing terms $\beta_j \|h_j\|^2$ and $\alpha_i KL(\rho_i | \tilde{\rho}_i)$.

For a numeric example, we consider a system with $N = 2$ user populations, $M = 2$ service providers, $d_h = 2$, and $d_\rho = 2$. An example application for this setting was illustrated in [198] as a bank account fraud problem, in which a population putting forward valid requests to open new bank accounts, ρ_1 , and a population submitting fraudulent requests, ρ_2 , are distributions over features such as income, customer age, number of months at a current residence, and similarity between the email and name associated with the application. We initialize two log-loss classifiers with the same cost function but different initial conditions. We compare two settings: one, neither population has nonlocal interactions, and two, the fraudulent population has attractive-repulsive nonlocal interactions. The attractive-repulsive kernel models the fraudulent agents attempting to differentiate themselves from other fraudulent agents in order to be misclassified, while having limited ability to move away from the overall population. The functions used in the simulations are

$$\begin{aligned} W_{11} &= C_r \exp\left(-\frac{\|x\|_1}{l_r}\right) - C_a \exp\left(-\frac{\|x\|_1}{l_a}\right), \quad p(x, h) := (1 + \exp(-c\theta(h)^\top x + b(h)))^{-1}, \\ u_{ij}(x_i, h_j) &= 1 - p(x_i, h_i), \quad l_{1j}(x_1, h_j) = \log(p(x_1, h_j)), \quad l_{2j} = \log(1 - p(x_2, h_j)), \end{aligned}$$

where $\theta(h) := [\cos h(0), \sin h(0)]$ and $b(h) = h(1)$ where $h(0)$ and $h(1)$ are the first and second entries of $h \in \mathbb{R}^2$. The function $p(x, h)$ is the probability of being approved for a loan, and l_{ij} is a log-likelihood, which has a larger gradient (with respect to h) than p . The coefficients used are $\eta = 0.5$, $c = 2$, $\alpha_1 = \alpha_2 = 0.5$, $\beta_1 = \beta_2 = 2.4$, $C_r = 8$, $l_r = 0.5$, $C_a = 2$, and $l_a = 1$. Even though




Figure 3.2: (Multi-Learner System) The fraudulent population (blue) and valid population (orange) both aim for high probability of approval from the classifiers. The black lines labeled x_1 and x_2 mark the 50% classification probability threshold for h_1 and h_2 respectively, with the arrows pointing in the direction of increasing probability of loan approval. In this setting, the fraudulent population has an interaction kernel, and the fraudulent population splits into two modes, strongly preferring the blue classifier.




Figure 3.3: (Multi-Learner System) In the setting without any interaction kernels, the fraudulent population does not split. Additionally, there are more fraudulent agents obtaining services from bank 2 than in the setting with the interaction kernel.

this system is not λ -monotone in, for example, $\mathcal{P}_2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2) \times \mathcal{P}_2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2) \times \mathbb{R}^2 \times \mathbb{R}^2$, we empirically observe convergence.

	kernel	ρ_1 using bank 1	ρ_1 using bank 2	ρ_2 using bank 1	ρ_2 using bank 2
$T = 0$	yes	$a_{11} = 0.74$	$a_{12} = 0.26$	$a_{21} = 0.01$	$a_{22} = 0.99$
	no	$a_{11} = 0.74$	$a_{12} = 0.26$	$a_{21} = 0.01$	$a_{22} = 0.99$
$T = 6$	yes	$a_{11} = 0.73$	$a_{12} = 0.27$	$a_{21} = 0.03$	$a_{22} = 0.97$
	no	$a_{11} = 0.65$	$a_{12} = 0.35$	$a_{21} = 0.07$	$a_{22} = 0.93$
$T = 12$	yes	$a_{11} = 0.58$	$a_{12} = 0.42$	$a_{21} = 0.11$	$a_{22} = 0.89$
	no	$a_{11} = 0.52$	$a_{12} = 0.48$	$a_{21} = 0.25$	$a_{22} = 0.75$
$T = 20$	yes	$a_{11} = 0.50$	$a_{12} = 0.50$	$a_{21} = 0.35$	$a_{22} = 0.65$
	no	$a_{11} = 0.45$	$a_{12} = 0.55$	$a_{21} = 0.40$	$a_{22} = 0.60$

Table 3.1: (Multi-Learner System) At times $T = 6, 12$, bank 1 has a larger share of the fraudulent population (ρ_1) than bank 2, with a larger margin in the kernel setting. At time $T = 20$, the valid population (ρ_2) prefers bank 2 with or without the interaction kernel, but the interaction kernel changes the bank preference for the fraudulent population.

In comparing the trajectories from both settings, we observe two key differences. In fig. 3.3, the population never splits into a bimodal distribution, unlike in the setting with the nonlocal interaction

kernel shown in fig. 3.2. Secondly, as show in table 3.1, in the interaction kernel setting, bank 1 has a worse customer profile at time steps $T = 6$ and $T = 12$, as it has more fraudulent customers and less valid customers than in the setting with no interaction kernel. Likewise, bank 2, which has the better initial condition, has a better customer profile. Even though the final population distributions appear to be similar, the customer allocation between the two classifiers favors bank 2.

3.5.2 Team Mean Field Games

Consider a setting with a single energy $F : \mathcal{A}_1 \rightarrow \mathbb{R}$ and single species $\rho = \rho_1$. We would like to solve a mean-field game parameterized by the energy functional F , viewing ρ as the density of particles participating in the game. Each particle aims to minimize a cost function, which includes a velocity penalty, $\|v\|^2$, as well as a penalty on the entire distribution, given in terms of $F[\rho]$. For long time horizons, a discount factor ε is used to weigh the cost at earlier times more heavily than the cost at later times and a terminal cost is given by g_ε . This model is a generalization of the model in [41], which describes the density of a crowd of people with an intruder passing through. The optimization problem is

$$\inf_v \left\{ \int_0^T \int \left(\frac{\|v(t, y(t))\|^2}{2} e^{-t/\varepsilon} + f_\varepsilon(t, y(t), \rho(t)) \right) d\rho(y(t)) dt + g_\varepsilon(T, y(T), \rho(T)) \right\},$$

such that $dy(t) = v(t, y(t))dt$, $y(0) \sim \rho(0, \cdot)$, $\text{Law}\{y\} = \rho(t, \cdot)$

where $f_\varepsilon(t, y(t), \rho(t)) := \frac{e^{-t/\varepsilon}}{\varepsilon} \delta_\rho F[\rho(t)]$ and $g_\varepsilon(t, y(t), \rho(t)) = \varepsilon f_\varepsilon(t, y(t), \rho(t)) = e^{-t/\varepsilon} \delta_\rho F[\rho(t)]$. The particle $y(t)$ is sampled from the distribution $\rho(t)$, and driven by the velocity field v . This control problem is equivalent to the mean field game

$$\inf_{v, \rho} \left\{ \int_0^T \left(\int \frac{e^{-t/\varepsilon}}{2} \|v(t, x)\|^2 d\rho(t, x) + \frac{e^{-t/\varepsilon}}{\varepsilon} F[\rho(t)] \right) dt + F[\rho(T)] e^{-T/\varepsilon} \right\},$$

such that $\partial_t \rho(t) + \nabla \cdot (v(t) \rho(t)) = 0$

This mean field game is studied in [213], and it was shown that the optimal (ρ, v) results in a perturbation of a gradient flow [213, Theorem 14, Remark 15]. Here, we consider an extension to the setting where different particles have different cost functionals depending on their type, or species. Collecting the particles by species, the optimization problem becomes

$$\inf_{v_i, \rho_i} \left\{ \int_0^T \left(\int \frac{e^{-t/\varepsilon}}{2} \|v_i(t, x_i)\|^2 d\rho_i(t, x_i) + \frac{e^{-t/\varepsilon}}{\varepsilon} F_i[\rho(t)] \right) dt + F_i[\rho(T)] e^{-T/\varepsilon} \right\},$$

such that $\partial_t \rho_i(t) + \nabla_{x_i} \cdot (v_i(t) \rho_i(t)) = 0$, $\forall i \in [1, \dots, n]$

Following a similar approach as the proof in [213] (equations 74-75), the optimal solution takes the form

$$\begin{aligned} \partial_t \rho_i &= \operatorname{div}_{x_i} (\rho_i \nabla_{x_i} \delta_{\rho_i} F_i[\rho]) + \varepsilon \operatorname{div}_{x_i} \left(\rho_i \nabla_{x_i} (\partial_t U_i - \frac{1}{2} |\nabla U_i|^2) \right) \\ -\varepsilon \partial_t U_i + U_i + \frac{\varepsilon}{2} |\nabla_{x_i} U_i|^2 &= \delta_{\rho_i} F_i[\rho]. \end{aligned} \quad (3.16)$$

For monotone F , (3.16) is a perturbation of coupled gradient flows with monotone energies. It is an interesting open problem to determine the sufficient conditions on $F_i[\rho_i, \cdot]$, such as smoothness, under which we expect that the result in [213, Theorem 14] for the limit as $\varepsilon \rightarrow 0$ can be shown in this multispecies setting. The long-time behavior of the solution in the limit $\varepsilon \rightarrow 0$ can then be analyzed using our monotonicity framework, since it follows the evolution (3.2).

3.5.3 Degenerately Monotone Game with Diffusion

For an energy functional that is "almost" strongly convex, exponential convergence of the corresponding gradient flow can still be shown if the energy functional also contains a linear entropy term. The notion of convexity used is called *degenerate* convexity. A potential function $V : \mathbb{R}^d \rightarrow \mathbb{R}$ is degenerately convex if $\nabla^2 V \geq \psi(x)$ for some $\psi \in C(\mathbb{R}^d)$ with $\psi(x) > 0$ if $x \neq 0$, and $\psi(x)$ is uniformly bounded from below by some positive constant as $\|x\| \rightarrow \infty$. Exponential convergence for gradient flows with energies of the form

$$F_i[\rho_i] = \int V_i(x_i) d\rho_i(x_i) + \int W_i(x_i - y_i) d\rho_i(x_i) d\rho_i(y_i) + \int \log \rho_i(x_i) d\rho_i(x_i),$$

where V_i and W_i are degenerately convex, are given in [50]; this occurs because the entropy term ensures ρ_i has mass in regions of the $\mathbb{R}^{d_i}$ where the Hessians of V_i and W_i are strongly convex. In finite dimensions, a convex potential that is not strongly convex results in a polynomial rate, but not an exponential rate. Here we apply the results from [50], Rule #5, to obtain a rate in combination with the effects of coupling between the species. From a PDE analysis perspective, it is interesting that the strength of the inter-species coupling does not necessarily change the exponential convergence rate. From a game theory perspective, our framework expands the set of games typically considered by including entropy in the cost function, and a framework for understanding why adding entropy, even though it is not strongly convex, enables an exponential rate. The energy functionals are $F_1 = F$ and $F_2 = -F$, with F given by

$$\begin{aligned} F[\rho_1, \rho_2] &= \frac{a}{2} \int \|x_1\|^4 d\rho_1(x_1) - \frac{a}{2} \int \|x_2\|^4 d\rho_2(x_2) + \frac{b}{2} \iint x_1^\top x_2 d\rho_1(x_1) d\rho_2(x_2) \\ &\quad + c(\mathcal{H}(\rho_1) - \mathcal{H}(\rho_2)), \end{aligned}$$

where $\mathcal{H}(\rho_i) = \int \log \rho_i(x_i) d\rho_i(x_i)$, and $a, b, c > 0$. In [50], the authors illustrate how functionals with degenerate convexity can still have exponential convergence due to diffusion, which keeps mass

in an area of the domain where the energy can be lower-bounded by a strongly convex function. In our example, the quartic term is not strongly convex; however, the diffusion term allows us to obtain a rate and we call this setting *degenerately monotone*. The dynamics are given by

$$\partial_t \rho_1 = -\nabla \mathcal{W}_{2,\rho_1} F[\rho_1, \rho_2] \quad \partial_t \rho_2 = \nabla \mathcal{W}_{2,\rho_2} F[\rho_1, \rho_2].$$

In fig. 3.4, the energy is plotted as a function of time, with the initial energy normalized to $F[\rho(0)] = 1$. Fixing $a = 15$ and $c = 0.5$, we see that the value of b has little effect on the convergence rate but does affect the oscillation frequency. This is because the interaction term is linear for each player. The simulation indicates that convergence can still occur without strong




Figure 3.4: (Degenerately Monotone Game with Diffusion) The upper-bound for the convergence rate across the tested values of b is $\lambda = 0.31$. As the coupling between the two players strengthens, the oscillations of F increase in frequency, but still decay at roughly the same rate.

monotonicity, as expected. In fig. 3.5, the densities of ρ_1 and ρ_2 are plotted for $t \in [0, 4]$, and the oscillating behavior of both species is demonstrated. Eventually the densities have a shape




Figure 3.5: (Degenerately Monotone Game with Diffusion) In the setting where $b = 150$, the density of each species is plotted from $t = 0$ to $t = 4$, during which time the energy has oscillated a few times. The densities for ρ_1 (blue) and ρ_2 (orange) spread out due to diffusion, eventually minimizing a quartic-shaped basin. The species ρ_2 aims to overlap with ρ_1 while ρ_1 aims to be in the opposite quadrant as ρ_2 , due to the coupling term $\iint x_1^\top x_2 d\rho_1(x_1) d\rho_2(x_2)$, which ρ_1 aims to minimize and ρ_2 aims to maximize.

consistent with quartic tails over $\mathbb{R}^2$. An interesting area of future work would be to explore not only the setting of degenerate monotonicity, but also more general non-monotone settings.

Acknowledgments

The authors are grateful for discussions with José A. Carrillo about displacement convexity, with Giuseppe Savaré about accretivity and recent literature, and with Xianjin Yang about mean field games, as well as to Dejan Slepčev, Eitan Levin, and Matthieu Darcy. Giulia Cavagnari provided a correction for the assumption of Theorem 3.2. LC is funded by an NDSEG fellowship from the AFOSR and a PIMCO fellowship, FH is supported by start-up funds at the California Institute of Technology and by NSF CAREER Award 2340762, EM is supported in part from NSF award 2240110, LJR is supported in part by NSF 1844729, NSF 2312775, and ONR YIP N000142012571.

3.A Auxiliary Definitions and Results

We start with an overview of different characterizations of the Wasserstein-2 metric used throughout the manuscript. Then we detail how to derive the equivalence between first order and second order monotonicity conditions in section 3.A.1.

Definition 3.35 (Wasserstein-2 Metric). *The Wasserstein-2 metric between two probability measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ is given by*

$$\mathcal{W}_2(\mu, \nu)^2 = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|z - z'\|_2^2 d\gamma(z, z')$$

where Γ is the set of all joint probability distributions with marginals μ and ν , i.e. $\mu(dz) = \int \gamma(dz, z') dz'$ and $\nu(dz') = \int \gamma(z, dz') dz$.

If $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ are such that at least one transport map T pushing forward μ to ν exists (written as $\nu = T_{\#}\mu$), we obtain the *Monge formulation* of the Wasserstein-2 distance:

$$\mathcal{W}_2(\mu, \nu)^2 = \inf_{T : \nu = T_{\#}\mu} \int \|x - T(x)\|^2 d\mu(x).$$

Since this is not true for all pairs $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$, the Kantorovich formulation Definition 3.35 is more general than the Monge formulation. Another way of writing the Wasserstein-2 Metric is via the *Benamou-Brenier formulation* (also known as the *dynamic formulation* of $\mathcal{W}_2$), expressing the distance between two probability measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ as the path solving a continuity equation with minimal kinetic energy:

$$\mathcal{W}_2(\mu, \nu)^2 = \inf_{(\mu_t, w_t)} \left\{ \int_0^1 \int \|w(t, x)\|^2 d\mu(t, x) dt ; \partial_t \mu_t + \nabla \cdot (\mu_t w_t) = 0 \text{ s.t. } \mu_0 = \mu, \mu_1 = \nu \right\}.$$

The Benamou-Brenier formulation is useful for geodesic calculations, as the optimal (μ, w) are also solutions to the geodesic equation (3.6).

3.A.1 Second order monotonicity condition

In what follows, we prove the equivalence of the first and second order monotonicity conditions (see Definition 3.1 and Proposition 3.27). Recall the displacement interpolant between $\rho_i^{(0)}$ and $\rho_i^{(1)}$ is defined as $\rho_i^{(s)} = (I_i^{(s)})_{\#} \gamma_i$, where $I_i^{(s)}(x_i, y_i) = (1-s)x_i + sy_i$, for all $s \in [0, 1]$ and $\gamma_i \in \Gamma^*(\rho_i^{(0)}, \rho_i^{(1)})$.

Lemma 3.36 (Second-Order Monotonicity, Geodesics). *Let Assumption 3.13 hold and $\mathcal{A} \subset \overline{\mathcal{P}_2}$ be geodesically convex. Then $v : \mathcal{A} \rightarrow \mathbb{R}^n$ is λ -monotone in $\mathcal{A}$ if and only if*

$$\begin{aligned} & - \int_0^1 \left(\iint \langle x - y, J_2[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) (\hat{x} - \hat{y}) \rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) \right. \\ & \quad \left. + \int \langle x - y, J_1[\rho^{(s)}](I^{(s)}(x, y)) (x - y) \rangle d\gamma(x, y) \right) ds \geq \lambda \int \|x - y\|^2 d\gamma(x, y), \end{aligned}$$

holds for all $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ and $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$, where $(\rho^{(s)})_{s \in [0, 1]}$ denotes the geodesic between $\rho^{(0)}$ and $\rho^{(1)}$.

Proof. Select $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ and $\gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$. We rewrite Definition 3.1 in terms of the derivatives of the interpolants evaluated at $s = 0$ and $s = 1$:

$$\begin{aligned} & \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(s)}]((1-s)x_i + sy_i) \rangle d\gamma_i(x_i, y_i) \right) \Big|_{s=0} = \int \langle y - x, v[\rho^{(0)}](x) \rangle d\gamma(x, y), \\ & \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(s)}]((1-s)x_i + sy_i) \rangle d\gamma_i(x_i, y_i) \right) \Big|_{s=1} = \int \langle y - x, v[\rho^{(1)}](y) \rangle d\gamma(x, y). \end{aligned}$$

Noting that

$$\begin{aligned} & \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(s)}]((1-s)x_i + sy_i) \rangle d\gamma_i(x_i, y_i) \right) \Big|_{s=1} \\ & - \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(s)}]((1-s)x_i + sy_i) \rangle d\gamma_i(x_i, y_i) \right) \Big|_{s=0} \\ & = \int_0^1 \frac{d}{d\tau} \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(\tau)}]((1-\tau)x_i + \tau y_i) \rangle d\gamma_i(x_i, y_i) \right) d\tau, \end{aligned}$$

monotonicity according to Definition 3.1 can be equivalently written as

$$- \int_0^1 \frac{d}{d\tau} \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(\tau)}]((1-\tau)x_i + \tau y_i) \rangle d\gamma_i(x_i, y_i) \right) d\tau \geq \lambda \int \|x - y\|^2 d\gamma(x, y).$$

We will use derivatives of the interpolants, which are computed in the weak form via

$$\int \partial_s \rho_i^{(s)}(x_i) \phi(x_i) dx_i = \int \partial_s \phi(I^{(s)}(x_i, y_i)) d\gamma_i(x_i, y_i) = \int \langle y_i - x_i, \nabla \phi(I^{(s)}(x_i, y_i)) \rangle d\gamma_i(x, y) \\ \forall \phi \in C_c^1(\mathbb{R}^{d_i}).$$

Writing out the derivative results in

$$\begin{aligned} & \frac{d}{ds} \left(\sum_{i=1}^n \int \langle y_i - x_i, v_i[\rho^{(s)}](I_i^{(s)}(x_i, y_i)) \rangle d\gamma_i(x_i, y_i) \right) \\ &= \sum_{i=1}^n \int \langle x_i - y_i, \nabla_{x_i} v_i[\rho^{(s)}](I_i^{(s)}(x_i, y_i)) (x_i - y_i) \rangle d\gamma_i(x_i, y_i) \\ &+ \sum_{i,j=1}^n \iint \langle \hat{x}_j - \hat{y}_j, \nabla_{x_j} \delta_{\rho_j} v_i[\rho^{(s)}](I_i^{(s)}(x_i, y_i), I_i^{(s)}(\hat{x}_j, \hat{y}_j)) (x_i - y_i) \rangle d\gamma_i(x_i, y_i) d\gamma_j(\hat{x}_j, \hat{y}_j) \\ &= \iint \langle x - y, J_2[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) (\hat{x} - \hat{y}) \rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) \\ &+ \int \langle x - y, J_1[\rho^{(s)}](I^{(s)}(x, y)) (x - y) \rangle d\gamma(x, y). \end{aligned}$$

Now we have shown that, for $\rho^{(s)}$ the geodesic interpolant between $\rho^{(0)}$ and $\rho^{(1)}$,

$$\begin{aligned} & \int_0^1 \left(\iint \langle x - y, J_2[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) (\hat{x} - \hat{y}) \rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y}) \right. \\ & \quad \left. + \int \langle x - y, J_1[\rho^{(s)}](I^{(s)}(x, y)) (x - y) \rangle d\gamma(x, y) \right) ds \\ &= \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y), \end{aligned}$$

and therefore the second-order geodesic notion of monotonicity is equivalent to the first-order definition. $\square$

Since $\overline{W}$ has only a formal Riemannian structure, we cannot directly apply results for monotonicity in the Riemannian manifold setting, but we can use an appropriate notion of tangent vectors instead of exponential maps, following the proof techniques which are outlined in [203, Proposition 16.2].

Lemma 3.37 (Second-Order Monotonicity, Velocity Fields). *Let Assumptions 3.13 to 3.15 hold and $\mathcal{A} \subset \overline{\mathcal{P}_2}$ be geodesically convex. Then $v : \mathcal{A} \rightarrow \mathbb{R}^n$ is λ -monotone in $\mathcal{A}$ if and only if*

$$\begin{aligned} & - \int \langle w(x), J_1[\rho^{(0)}](x) w(x) \rangle d\rho^{(0)}(x) - \iint \langle w(x), J_2[\rho^{(0)}](x, \hat{x}) w(\hat{x}) \rangle d\rho^{(0)}(x) d\rho^{(0)}(\hat{x}) \\ & \geq \lambda \int \|w(x)\|^2 d\rho^{(0)}(x). \end{aligned}$$

for all $\rho^{(0)} \in \mathcal{A}$ and $w_i \in Y_i(\rho_i^{(0)}, \mathcal{A}_i)$ for all $i \in [1, \dots, n]$.

Proof. We will show this local notion of monotonicity using ideas from the proof of [203, Proposition 16.2]. Fix $\rho \in \mathcal{A}$ and a set of velocity fields $\{w_i \in Y_i(\rho_i, \mathcal{A}_i)\}_{i=1}^n$. Since $w_i \in Y_i(\rho_i, \mathcal{A}_i)$ by Assumption 3.15,

$$\tilde{\rho}_i := (w_i/c + \text{id})_{\#}\rho_i \in \mathcal{A}_i \quad \text{and} \quad (\text{id}, w_i/c + \text{id})_{\#}\rho_i \in \Gamma^*(\rho_i, \tilde{\rho}_i)$$

for some $c > 0$. The geodesic interpolant between ρ_i and $\tilde{\rho}_i$ is given by

$$\rho_i^{(\tau)} := [(1 - \tau) \text{id} + \tau(w_i/c + \text{id})]_{\#}\rho_i = [\text{id} + \tau w_i/c]_{\#}\rho_i, \quad \tau \in [0, 1].$$

Since both $\rho_i, \tilde{\rho}_i \in \mathcal{A}_i$, we can use Lemma 3.36 applied to ρ and $\rho^{(\tau)}$. Noting that $\gamma_i^{(\tau)} := (\text{id}, \text{id} + \tau w_i/c)_{\#}\rho_i$ is an optimal plan between ρ_i and $\rho_i^{(\tau)}$, and $I^{(s)}(x, x + \tau w(x)/c) = x + s\tau w(x)/c$, we have that the geodesic at $s \in [0, 1]$ between ρ_i and $\rho_i^{(\tau)}$ is given by $\rho^{(s\tau)}$, and

$$\begin{aligned} & - \int_0^1 \left(\iint \langle x - y, J_2[\rho^{(s\tau)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y})) (\hat{x} - \hat{y}) \rangle d\gamma^{(\tau)}(x, y) d\gamma^{(\tau)}(\hat{x}, \hat{y}) \right. \\ & \quad \left. + \int \langle x - y, J_1[\rho^{(s\tau)}](I^{(s)}(x, y)) (x - y) \rangle d\gamma^{(\tau)}(x, y) \right) ds \\ & = -\frac{\tau^2}{c^2} \int_0^1 \left(\iint \langle w(x), J_2[\rho^{(s\tau)}](x + s\tau w(x)/c, \hat{x} + s\tau w(\hat{x})/c) \cdot w(\hat{x}) \rangle d\rho(x) d\rho(\hat{x}) \right. \\ & \quad \left. + \frac{\tau^2}{c^2} \int \langle w(x), J_1[\rho^{(s\tau)}](x + s\tau w(x)/c) \cdot w(x) \rangle d\rho(x) \right) ds \geq \lambda \frac{\tau^2}{c^2} \int \|w(x)\|^2 d\rho(x). \end{aligned}$$

Since $\mathcal{A}$ is convex, $\rho^{(\tau)} \in \mathcal{A}$ for all $\tau \in [0, 1]$, and the monotonicity condition holds as $\tau \rightarrow 0$. Thanks to the continuity properties in Assumption 3.14, dividing both sides by $\frac{\tau^2}{c^2}$ and taking the limit as $\tau \rightarrow 0$ gives

$$\begin{aligned} & - \iint \langle w(x), J_2[\rho](x, \hat{x}) w(\hat{x}) \rangle d\rho(x) d\rho(\hat{x}) - \int \langle w(x), J_1[\rho](x) w(x) \rangle d\rho(x) \\ & \geq \lambda \int \|w(x)\|^2 d\rho(x). \end{aligned}$$

□

Proof of Proposition 3.27. The proof follows by applying Lemma 3.36 and Lemma 3.37 with $v_i[\rho] = -\nabla_{x_i} \delta_{\rho_i} F_i[\rho]$. □

3.A.2 Monotonicity vs Convexity and Measures vs Vectors

In tables 3.2 and 3.3, the differences between monotonicity and convexity and between games in Euclidean spaces and games in measure spaces over Euclidean spaces are detailed, where the finite-dimensional games are defined as in Definition 3.25. For both convexity and monotonicity, a key difference between finite and infinite dimensions is the role of the Onsager operator's dependence

on ρ for the $\mathscr{W}_2$ metric. For simplicity, we will illustrate this in the single species setting, the multi-species setting follows similarly. Let (U, d) be a metric space, and let $E : U \rightarrow \mathbb{R} \cup \{+\infty\}$ be an energy functional. Consider the single species (informal) gradient flow

$$\dot{u} = -\mathbb{K}(u)DE(u), \quad (3.17)$$

where $\mathbb{K}(u) : (T_u U)^* \rightarrow T_u U$ is the Onsager operator corresponding to the metric d , mapping from the dual of the tangent space at u to the tangent space at u . D is a differential operator on the energy E . We would like a condition under which the functional

$$V_2(x) := \frac{1}{2} \langle DE(u), \mathbb{K}(u)DE(u) \rangle$$

converges exponentially along solutions to (3.17) (see the Lyapunov functionals defined in (3.5) for comparison). Taking a time derivative of V_2 ,

$$\begin{aligned} \dot{V}_2(u) &= \frac{1}{2} \left\langle DE(u), \frac{d}{du} (\mathbb{K}(u)DE(u)) [\dot{u}] \right\rangle + \frac{1}{2} \langle \dot{u}D^2E(u), \mathbb{K}(u)DE(u) \rangle \\ &= \frac{1}{2} \left\langle DE(u), \frac{d}{du} (\mathbb{K}(u)) [\dot{u}]DE(u) \right\rangle + \langle \dot{u}D^2E(u), \mathbb{K}(u)DE(u) \rangle, \end{aligned}$$

where the last equality above holds if $\mathbb{K}(u)$ is a linear operator, and $\frac{d}{du} (\mathbb{K}(u)) [\dot{u}]DE(u)$ denotes the derivative of $\mathbb{K}$ with respect to u in the direction $\dot{u}$, acting on $DE(u)$. A convexity condition allows to conclude that

$$\dot{V}_2(u) \leq -\lambda \langle DE(u), \mathbb{K}(u)DE(u) \rangle = -2\lambda V_2(u),$$

guaranteeing exponential convergence of V_2 . For more details on the above general framework for gradient flows, see [167]. Next, we instantiate this general framework for the Euclidean case, where $U = \mathbb{R}^d$ and $d = \|\cdot\|$, and in the case of measures, selecting $U = \mathcal{P}_2$ and $d = \mathscr{W}_2$. In the Euclidean setting, we identify $DE(u)$ as $\nabla E(u)$ and $D^2E(u)$ as $\nabla^2 E(u)$; in the measure setting, we identify $DE(u)$ as $\delta E(u)$ and $D^2E(u)$ as $\delta^2 E(u)$. We have

Euclidean:	$\mathbb{K}(u) = \text{id},$	$\frac{d}{du} (\mathbb{K}(u)) [\dot{u}] = 0,$
Measures:	$\mathbb{K}(u) = -\text{div}(u\nabla \cdot),$	$\frac{d}{du} (\mathbb{K}(u)) [\dot{u}] = -\text{div}(\dot{u}\nabla \cdot).$

In the Euclidean setting,

$$\dot{V}_2(u) = -\langle \nabla E(u), \nabla^2 E(u) \cdot \nabla E(u) \rangle \leq -\lambda \langle \nabla E(u), \nabla E(u) \rangle = -2\lambda V_2(u),$$

due to the convexity condition $\nabla^2 E(u) \succeq \lambda \mathbf{I}_d$ for all $u \in \mathbb{R}^d$. In the measure setting, we have

$$\begin{aligned}
\dot{V}_2(u) &= -\frac{1}{2} \langle \delta E(u), \operatorname{div}(\dot{u} \nabla \delta E(u)) \rangle - \langle \operatorname{div}(u \nabla \delta E(u)) \delta^2 E(u), \operatorname{div}(u \nabla \delta E(u)) \rangle \\
&= \frac{1}{2} \int \|\nabla \delta E[u](x)\|^2 \operatorname{div}(u(x) \nabla \delta E[u](x)) \, dx \\
&\quad - \iint \langle \nabla \delta E[u](x), \nabla^2 \delta^2 E[u](x, \hat{x}) \nabla \delta E[u](\hat{x}) \rangle \, du(x) du(\hat{x}) \\
&= - \int \langle \nabla \delta E[u](x), \nabla^2 \delta E[u](x) \nabla \delta E[u](x) \rangle \, du(x) \\
&\quad - \iint \langle \nabla \delta E[u](x), \nabla^2 \delta^2 E[u](x, \hat{x}) \nabla \delta E[u](\hat{x}) \rangle \, du(x) du(\hat{x}).
\end{aligned}$$

We recognize the above expression as the left-hand side in the second-order convexity condition (single-species monotonicity in Lemma 3.37), which results in

$$\dot{V}_2(u) \leq -2\lambda V_2(u).$$

In summary, since in the Euclidean case $\mathbb{K}$ does not depend on u , and so $\frac{d}{du}(\mathbb{K}(u))[\dot{u}]DE(u) = 0$ and there is only one Hessian-like term in the second-order monotonicity condition in the finite-dimensional case (see Definition 3.25 or table 3.2). In contrast, this term is not zero in the measure setting (see Proposition 3.27 or table 3.3) as the derivative of the Onsager operator results in $\int \langle \nabla \delta E[u](x), \nabla^2 \delta E[u](x) \nabla \delta E[u](x) \rangle \, du(x)$ in the $\mathcal{W}_2$ case. If the Onsager operator was nonlinear in the object on which it acts (for instance when there is a cosh structure [181] instead of a quadratic dissipation structure), then there could be three Hessian-like terms, with the third coming from the derivative of $\mathbb{K}$ with respect to DE .

Monotonicity and Convexity over $\mathbb{R}^d$		
$\forall x, y \in \mathbb{R}^d$	First Order	Second Order
Convexity	$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \lambda \ x - y\ ^2$	$\frac{1}{2}(\nabla^2 f(x) + \nabla^2 f(x)^\top) \succeq \lambda \mathbf{I}_d$
Monotonicity	$\langle \bar{\nabla} V(x) - \bar{\nabla} V(y), x - y \rangle \geq \lambda \ x - y\ ^2$	$\frac{1}{2}(\bar{\nabla}^2 V(x) + \bar{\nabla}^2 V(x)^\top) \succeq \lambda \mathbf{I}_d$

Table 3.2: (Coupled Gradient Flow Setting) The key difference between monotonicity and convexity is the operator $\bar{\nabla}$ which is *not* a gradient operator; it is an entry-wise operator which takes a gradient with respect to different states.

3.B Interaction Kernel Computations

Under appropriate conditions we prove that the nonlocal cross-interaction terms appearing in section 3.4.4 are geodesically displacement convex Lemma 3.38, and provide Hessian lower bounds

Monotonicity and Convexity over $\mathcal{P}(\mathbb{R}^d)$

	First Order
Convexity	$\int \left\langle \nabla \delta_{\rho_1} F[\rho_1^{(0)}](x) - \nabla \delta_{\rho_1} F[\rho_1^{(1)}](y), x - y \right\rangle d\gamma_1(x) \geq \lambda \mathcal{W}_2(\rho_1^{(0)}, \rho_1^{(1)})^2$ $\forall \rho_1^{(0)}, \rho_1^{(1)} \in \mathcal{A}_1, \forall \gamma_1 \in \Gamma^*(\rho_1^{(0)}, \rho_1^{(1)})$
Monotonicity	$\int \left\langle \overline{\nabla \delta_{\rho}} F[\rho^{(0)}](x) - \overline{\nabla \delta_{\rho}} F[\rho^{(1)}](y), x - y \right\rangle d\gamma(x, y) \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2$ $\forall \rho^{(0)}, \rho^{(1)} \in \mathcal{A}, \forall \gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$
	Second Order
Convexity	$\int (\iint \langle x - y, \nabla^2 \delta_{\rho_1}^2 F[\rho_1^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y}))(\hat{x} - \hat{y}) \rangle d\gamma_1(x, y) d\gamma_1(\hat{x}, \hat{y})$ $+ \int \langle x - y, \nabla^2 \delta_{\rho_1} F[\rho_1^{(s)}](I^{(s)}(x, y))(x - y) \rangle d\gamma_1(x, y) ds \geq \lambda \overline{W}(\rho_1^{(0)}, \rho_1^{(1)})^2$ $\forall \rho_1^{(0)}, \rho_1^{(1)} \in \mathcal{A}_1, \gamma_1 \in \Gamma^*(\rho_1^{(0)}, \rho_1^{(1)})$
Monotonicity	$\int (\iint \langle x - y, \overline{\nabla^2 \delta_{\rho}} F[\rho^{(s)}](I^{(s)}(x, y), I^{(s)}(\hat{x}, \hat{y}))(\hat{x} - \hat{y}) \rangle d\gamma(x, y) d\gamma(\hat{x}, \hat{y})$ $+ \int \langle x - y, \overline{\nabla^2 \delta_{\rho}} F[\rho^{(s)}](I^{(s)}(x, y))(x - y) \rangle d\gamma(x, y) ds \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2$ $\forall \rho^{(0)}, \rho^{(1)} \in \mathcal{A}, \gamma \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$

Table 3.3: (Coupled Gradient Flow Setting) Similar to the $\mathbb{R}^d$ setting, the difference between monotonicity and convexity is the entry-wise operators. The second order inequalities have two Hessian-like terms in the measure setting, unlike the $\mathbb{R}^d$ setting.

kernels such as the attractive-repulsive kernel and power law kernel. Next, going beyond the condition in Lemma 3.32, we show monotonicity under more general criteria in Lemma 3.39.

Lemma 3.38 (Convexity of $\mathcal{W}$). *Consider the functional $\mathcal{W}$ defined as*

$$\mathcal{W}_{ij}(\rho_i, \rho_j) = \iint W_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j).$$

Let the interaction kernel $W_{ij} \in C^2(\mathbb{R}^{d_i} \times \mathbb{R}^{d_j}; \mathbb{R})$ satisfy $\text{Hess}(W_{ij}) \succeq \tilde{\lambda} \mathbf{I}_{d_i+d_j}$. Then the second derivative of $\mathcal{W}_{ij}$ along geodesics between $\rho^{(0)}$ and $\rho^{(1)}$ satisfies

$$\frac{d^2}{ds^2} \mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) \geq \tilde{\lambda} \left(\mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2 + \mathcal{W}_2(\rho_j^{(0)}, \rho_j^{(1)})^2 \right) \quad \forall s \in [0, 1].$$

Proof. Select $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ and consider a geodesic $(\rho^{(s)}, \nu^{(s)})$ between them, solving (3.6). The

first derivative is given by

$$\begin{aligned} \frac{d}{ds} \mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) &= \int \left\langle \nabla_{x_i} \delta_{\rho_i} \mathcal{W}_{ij}[\rho_i^{(s)}, \rho_j^{(s)}](x_i), v_i^{(s)}(x_i) \right\rangle d\rho_i^{(s)}(x_i) \\ &\quad + \int \left\langle \nabla_{x_j} \delta_{\rho_j} \mathcal{W}_{ij}[\rho_i^{(s)}, \rho_j^{(s)}](x_j), v_j^{(s)}(x_j) \right\rangle d\rho_j^{(s)}(x_j) \\ &= \iint \left\langle \begin{bmatrix} \nabla_{x_i} W_{ij}(x_i, x_j) \\ \nabla_{x_j} W_{ij}(x_i, x_j) \end{bmatrix}, \begin{bmatrix} v_i^{(s)}(x_i) \\ v_j^{(s)}(x_j) \end{bmatrix} \right\rangle d\rho_i^{(s)}(x_i) d\rho_j^{(s)}(x_j). \end{aligned}$$

Differentiating again, we have

$$\frac{d^2}{ds^2} \mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) = \iint \begin{bmatrix} v_i^{(s)}(x_i) \\ v_j^{(s)}(x_j) \end{bmatrix}^\top H(x_i, x_j) \begin{bmatrix} v_i^{(s)}(x_i) \\ v_j^{(s)}(x_j) \end{bmatrix} d\rho_i^{(s)}(x_i) d\rho_j^{(s)}(x_j)$$

where

$$H(x_i, x_j) = \text{Hess}(W_{ij}) := \begin{bmatrix} \nabla_{x_i}^2 W(x_i, x_j) & \nabla_{x_j, x_i}^2 W_{ij}(x_i, x_j) \\ \nabla_{x_i, x_j}^2 W(x_i, x_j) & \nabla_{x_j}^2 W_{ij}(x_i, x_j) \end{bmatrix}.$$

The Hessian H is well-defined because $W \in C^2$. Under the condition for the Hessian $\text{Hess}(W) \succeq \tilde{\lambda} \mathbf{I}_{d_i+d_j}$, we have that

$$\begin{aligned} \frac{d^2}{ds^2} \mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) &\geq \tilde{\lambda} \iint \left\| \begin{bmatrix} v_i^{(s)}(x_i) \\ v_j^{(s)}(x_j) \end{bmatrix} \right\|^2 d\rho_i^{(s)}(x_i) d\rho_j^{(s)}(x_j) \\ &= \tilde{\lambda} (\mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2 + \mathcal{W}_2(\rho_j^{(0)}, \rho_j^{(1)})^2). \end{aligned}$$

□

In the setting where $\mathbb{R}^{d_i} = \mathbb{R}^{d_j}$ and $x = x_i - x_j$, we compute the lower bounds for the Hessians of $W^{(k)}$, $W^{(m)}$, and $W^{(p)}$, given by

$$\begin{aligned} W^{(k)}(x) &= \frac{\|x\|^k}{k}, \quad W^{(m)}(x) = C_r \exp\left(-\frac{\|x\|^2}{l_r}\right) - C_a \exp\left(-\frac{\|x\|^2}{l_a}\right), \\ W^{(p)}(x) &= \frac{\|x\|^a}{a} - \frac{\|x\|^b}{b}, \end{aligned}$$

where $l_a > l_r > 0$, $C_r > C_a > 0$ and $C_r/C_a < (l_r/l_a)^{-d_i}$ and $a > b$ with a, b even integers. These lower bounds allow us to apply Lemma 3.32 to systems with interaction kernels of this form. The kernel $W^{(k)} = \frac{\|x\|^k}{k}$ has a Hessian given by $\text{Hess}(W^{(k)}) \succeq 0 \mathbf{I}_{d_i}$ for integers $k \geq 1$. The Hessian of the Morse-like potential $W^{(m)}$ is more involved; the gradient is given by

$$\nabla W^{(m)}(x) = 2x \left(-\frac{C_r}{l_r} \exp(-\|x\|^2/l_r) + \frac{C_a}{l_a} \exp(-\|x\|^2/l_a) \right),$$

and the Hessian is

$$\begin{aligned} \text{Hess} \left(W^{(m)} \right) &= 2 \mathbf{I}_{d_i} \left(-\frac{C_r}{l_r} \exp(-\|x\|^2 / l_r) + \frac{C_a}{l_a} \exp(-\|x\|^2 / l_a) \right) \\ &\quad + 4xx^\top \left(\frac{C_r}{l_r^2} \exp(-\|x\|^2 / l_r) - \frac{C_a}{l_a^2} \exp(-\|x\|^2 / l_a) \right). \end{aligned}$$

The smallest eigenvalues of each matrix are given by

$$\begin{aligned} \lambda_{m1} &:= 2 \inf_x \left(-\frac{C_r}{l_r} \exp(-\|x\|^2 / l_r) + \frac{C_a}{l_a} \exp(-\|x\|^2 / l_a) \right) \\ \lambda_{m2} &:= 4 \min \left\{ \inf_x \left(\|x\|^2 \left(\frac{C_r}{l_r^2} \exp(-\|x\|^2 / l_r) - \frac{C_a}{l_a^2} \exp(-\|x\|^2 / l_a) \right) \right), 0 \right\}, \end{aligned}$$

and since $l_a, l_r > 0$, the smallest eigenvalues satisfy $\lambda_{m1}, \lambda_{m2} > -\infty$ and so there exists $\tilde{\lambda} = \lambda_{m1} + \lambda_{m2}$ such that $\text{Hess} \left(W^{(m)} \right) \succeq \tilde{\lambda} \mathbf{I}_{d_i}$.

For the third interaction kernel $W^{(p)}$, the Hessian is given by

$$\text{Hess} \left(W^{(p)} \right) = \left((a-2) \|x\|^{a-4} - (b-2) \|x\|^{b-4} \right) xx^\top + (\|x\|^{a-2} - \|x\|^{b-2}) \mathbf{I}_{d_i}.$$

Because a is even and $a > b$, the Hessian has a uniform lower bound on the smallest eigenvalue for fixed a, b . When $b = 2$, $\text{Hess} \left(W^{(p)} \right) \succeq -\mathbf{I}_{d_i}$. When $b > 2$, both $(a-2) \|x\|^{a-4} - (b-2) \|x\|^{b-4}$ and $\|x\|^{a-2} - \|x\|^{b-2}$ have uniform lower-bounds because $a > b$, and therefore $\text{Hess} \left(W^{(p)} \right)$ has a finite minimum eigenvalue $\tilde{\lambda}$. See [16] for more details on the estimates for $W^{(p)}$. Below, we present a general condition on interaction kernels that results in monotonicity.

Lemma 3.39 (Cross-Interaction Kernels). *Consider the interaction term $\mathcal{W}_{ij} : \mathcal{A}_i \times \mathcal{A}_j \rightarrow \mathbb{R}$, parameterized by the interaction kernel $W_{ij} \in C^2(\mathbb{R}^{d_i} \times \mathbb{R}^{d_j}, \mathbb{R})$, given by*

$$\mathcal{W}_{ij}(\rho_i, \rho_j) = \iint W_{ij}(x_i, x_j) d\rho_i(x_i) d\rho_j(x_j),$$

Let $F_i = \sum_{j=1}^n \mathcal{W}_{ij}$ for all $i \in [1, \dots, n]$. When $\text{Hess} \left(W_{ij} \right) \succeq \lambda_{ij} \mathbf{I}_{d_i+d_j}$ and

$$\begin{aligned} &\alpha \sum_{i,j=1}^n \left\langle \nabla_{x_i} [W_{ij}(x_i, x_j) + W_{ji}(x_j, x_i)] - \nabla_{x_i} [W_{ij}(y_i, y_j) + W_{ji}(y_j, y_i)], x_i - y_i \right\rangle \\ &\leq \sum_{i,j=1}^n \left\langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \right\rangle \quad \forall x_i, y_i \in \mathbb{R}^{d_i}, x_j, y_j \in \mathbb{R}^{d_j}, \end{aligned} \tag{3.18}$$

then F is $\alpha\lambda$ -monotone in $\mathcal{A}$, where $\lambda_i = \sum_{j=1}^n \lambda_{ij} + \lambda_{ji}$ and $\lambda = \min_i \lambda_i$.

Proof. Using the exact form of the Taylor expansion for $\mathcal{W}_{ij}$, and the bound for the second derivative from Lemma 3.38, $\frac{d^2}{ds^2}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) \geq \lambda_{ij}(\mathcal{W}_2(\rho_i^{(0)}, \rho_i^{(1)})^2 + \mathcal{W}_2(\rho_j^{(0)}, \rho_j^{(1)})^2)$, we have along geodesics between $(\rho_i^{(0)}, \rho_j^{(0)})$ and $(\rho_i^{(1)}, \rho_j^{(1)})$,

$$\begin{aligned}\mathcal{W}_{ij}(\rho_i^{(1)}, \rho_j^{(1)}) &= \mathcal{W}_{ij}(\rho_i^{(0)}, \rho_j^{(0)}) + \frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=0} \\ &\quad + \int_0^1 (1-t) \left(\frac{d^2}{ds^2}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) \right)\Big|_{s=t} dt \\ &\geq \mathcal{W}_{ij}(\rho_i^{(0)}, \rho_j^{(0)}) + \frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=0} + \frac{\lambda_{ij}}{2} \sum_{k=i,j} \mathcal{W}_2(\rho_k^{(0)}, \rho_k^{(1)})^2, \\ \mathcal{W}_{ij}(\rho_i^{(0)}, \rho_j^{(0)}) &= \mathcal{W}_{ij}(\rho_i^{(1)}, \rho_j^{(1)}) - \frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=1} - \int_1^0 t \left(\frac{d^2}{ds^2}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) \right)\Big|_{s=t} dt \\ &\geq \mathcal{W}_{ij}(\rho_i^{(1)}, \rho_j^{(1)}) - \frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=1} + \frac{\lambda_{ij}}{2} \sum_{k=i,j} \mathcal{W}_2(\rho_k^{(0)}, \rho_k^{(1)})^2.\end{aligned}$$

The term $\frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})$, using geodesics (3.6), is given by

$$\begin{aligned}\frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)}) &= \int \left(\left\langle \nabla_{x_i} W_{ij}(x_i, x_j), w_i^{(s)}(x_i) \right\rangle \right. \\ &\quad \left. + \left\langle \nabla_{x_j} W_{ij}(x_i, x_j), w_j^{(s)}(x_j) \right\rangle \right) d\rho_i^{(s)}(x_i) d\rho_j^{(s)}(x_j).\end{aligned}$$

Recall $\rho_i^{(s)} = (I_i^{(s)})_{\#}\gamma_i$ where $I_i^{(s)}(x_i, y_i) = (1-s)x_i + sy_i$, and that $w_i^{(s)}(I_i^{(s)}(x_i, y_i)) = y_i - x_i$. Evaluating $\frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})$ at $s = 0$ and $s = 1$,

$$\begin{aligned}\frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=0} &= \iint \left\langle \begin{bmatrix} \nabla_{x_i} W_{ij}(x_i, x_j) \\ \nabla_{x_j} W_{ij}(x_i, x_j) \end{bmatrix}, \begin{bmatrix} y_i - x_i \\ y_j - x_j \end{bmatrix} \right\rangle d\gamma_i(x_i, y_i) d\gamma_j(x_j, y_j) \\ \frac{d}{ds}\mathcal{W}_{ij}(\rho_i^{(s)}, \rho_j^{(s)})\Big|_{s=1} &= \iint \left\langle \begin{bmatrix} \nabla_{x_i} W_{ij}(y_i, y_j) \\ \nabla_{x_j} W_{ij}(y_i, y_j) \end{bmatrix}, \begin{bmatrix} y_i - x_i \\ y_j - x_j \end{bmatrix} \right\rangle d\gamma_i(x_i, y_i) d\gamma_j(x_j, y_j).\end{aligned}$$

Adding the two Taylor expansion inequalities results in

$$\iint \left\langle \begin{bmatrix} \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j) \\ \nabla_{x_j} W_{ij}(x_i, x_j) - \nabla_{x_j} W_{ij}(y_i, y_j) \end{bmatrix}, \begin{bmatrix} x_i - y_i \\ x_j - y_j \end{bmatrix} \right\rangle d\gamma_i(x_i, y_i) d\gamma_j(x_j, y_j) \quad (3.19)$$

$$\geq \lambda_{ij} \sum_{k=i,j} \mathcal{W}_2(\rho_k^{(0)}, \rho_k^{(1)})^2. \quad (3.20)$$

Summing (3.19) over all $i, j \in [1, \dots, n]$,

$$\begin{aligned}\sum_{i,j=1}^n \iint \langle \nabla_{x_i} [W_{ij}(x_i, x_j) + W_{ji}(x_j, x_i)] - \nabla_{x_i} [W_{ij}(y_i, y_j) + W_{ji}(y_j, y_i)], x_i - y_i \rangle d\gamma(x, y) \\ \geq \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2,\end{aligned}$$

where $\lambda_i = \sum_{j=1}^n \lambda_{ij} + \lambda_{ji}$ and $\lambda = \min_i \lambda_i$. Then multiplying both sides by α and applying (3.18),

$$\sum_{i,j=1}^n \iint \langle \nabla_{x_i} W_{ij}(x_i, x_j) - \nabla_{x_i} W_{ij}(y_i, y_j), x_i - y_i \rangle d\gamma(x, y) \geq \alpha \lambda \overline{W}(\rho^{(0)}, \rho^{(1)})^2.$$

Therefore F is $\alpha\lambda$ -monotone in $\mathcal{A}$. $\square$

3.C Optimal and Admissible Plans in the Monotonicity Inequality

In this section, we prove the equivalence between the monotonicity condition using optimal plans and using admissible plans, in the setting where v is monotone over $\overline{\mathcal{P}}_2$. The proofs use techniques from [55], where the authors first prove the equivalence for plans which have finite support, and then extend this result to any probability measure with finite second moment.

Recall that $\Gamma(\rho_i^{(0)}, \rho_i^{(1)})$ is the set of all plans with marginals $\rho_i^{(0)}, \rho_i^{(1)}$, and $\Gamma^*(\rho_i^{(0)}, \rho_i^{(1)}) \subset \Gamma(\rho_i^{(0)}, \rho_i^{(1)})$ is the set of optimal plans with marginals $\rho_i^{(0)}, \rho_i^{(1)}$. Let $\mathcal{P}^{(f)}(\mathbb{R}^{d_i}) \subset \mathcal{P}_2(\mathbb{R}^{d_i})$ be the set of measures on $\mathbb{R}^{d_i}$ with finite support, defined as

$$\mathcal{P}^{(f)} = \{\rho_i \in \mathcal{P}_2(\mathbb{R}^{d_i}) \mid \text{supp } \rho_i = \{x_i^{(1)}, \dots, x_i^{(M_i)}\}, M_i < \infty\},$$

and let $\overline{\mathcal{P}}^{(f)} := \mathcal{P}^{(f)}(\mathbb{R}^{d_1}) \times \dots \times \mathcal{P}^{(f)}(\mathbb{R}^{d_n})$. We define $I_i^{(t)}(x_i, y_i) = (1-t)x_i + ty_i$.

Theorem 3.40 (Local optimality of discrete interpolations (Theorem 6.2, [55])). *Let $\rho_i^{(0)}, \rho_i^{(1)} \in \mathcal{P}^{(f)}(\mathbb{R}^{d_i})$, select any $\gamma_i \in \Gamma(\rho_i^{(0)}, \rho_i^{(1)})$, and define $\rho_i^{(\tau)} := (I_i^{(\tau)})_{\#} \gamma_i$. Then there exists a finite number of points $\tau_0^{(i)} = 0 < \tau_1^{(i)} < \dots < \tau_{K_i}^{(i)} = 1$ such that for every $k = 1, \dots, K_i$, $\rho_i^{(\tau)}$ for $\tau \in [\tau_{k-1}^{(i)}, \tau_k^{(i)}]$ is a minimal constant-speed geodesic and*

$$\mathcal{W}_2^2(\rho_i^{(t_1)}, \rho_i^{(t_2)}) = |t_1 - t_2|^2 \int \|x - y\|^2 d\gamma_i(x, y), \quad \forall t_1, t_2 \in [\tau_{k-1}^{(i)}, \tau_k^{(i)}].$$

Proposition 3.41 (Admissible Couplings for Finitely-Supported Marginals). *If $\overline{\mathcal{P}}^{(f)} \subseteq \mathcal{A}$ and v is λ -monotone in $\mathcal{A}$, i.e. for any $\rho^{(0)}, \rho^{(1)} \in \overline{\mathcal{P}}^{(f)}$ with $\gamma_i^* \in \Gamma^*(\rho_i^{(0)}, \rho_i^{(1)})$, $\gamma^* = \prod_{i=1}^n \gamma_i^*$, it holds*

$$-\int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma^*(x, y) \geq \lambda \int \|x - y\|^2 d\gamma^*(x, y),$$

then for any $\gamma_i \in \Gamma(\rho_i^{(0)}, \rho_i^{(1)})$, $\gamma = \prod_{i=1}^n \gamma_i$ such that $\rho_i^{(\tau)} := (I_i^{(\tau)})_{\#} \gamma_i$ satisfies $\rho^{(\tau)} \in \mathcal{A}$ for all $\tau \in [0, 1]$, it holds

$$-\int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \int \|x - y\|^2 d\gamma(x, y).$$

Proof. We select $\gamma_i \in \Gamma(\rho_i^{(0)}, \rho_i^{(1)})$ for all i . First, we will construct a set of time increments such that this admissible plan is locally optimal for all species in each increment. By Theorem 3.40,

for each species i there exists a set of points $(\tau_j^{(i)})_{j=1}^{K_i}$ such that $\rho_i^{(\tau)}$ is a constant-speed geodesic for $\tau \in [\tau_{k-1}^{(i)}, \tau_k^{(i)}]$. Taking the union of all such time steps for each species and renaming them in strictly increasing order,

$$\bigcup_{i=1}^n (\tau_j^{(i)})_{j=1}^{K_i} = \{\bar{\tau}_0 < \bar{\tau}_1 < \dots < \bar{\tau}_{\bar{K}}\},$$

the local optimality of $\gamma_i^{(\tau)}$ holds for all species in each interval $[\bar{\tau}_{k-1}, \bar{\tau}_k]$. The set of time steps $(\bar{\tau}_k)$ has size $\bar{K} \leq \sum_{i=1}^n K_i - 2n + 2$, where the $-2n + 2$ comes from the same initial and final time steps for all species. Also by Theorem 3.40, the optimal coupling from $\rho_i^{(\bar{\tau}_{k-1})}$ to $\rho_i^{(\bar{\tau}_k)}$ is given by $\gamma_i^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}$, where

$$\gamma_i^{(s,t)} := [I_i^{(s)}, I_i^{(t)}]_{\#} \gamma_i.$$

Hence, by λ -monotonicity,

$$- \int \langle x - y, v[\rho^{(\bar{\tau}_{k-1})}](x) - v[\rho^{(\bar{\tau}_k)}](y) \rangle d\gamma^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}(x, y) \geq \lambda \int \|x - y\|^2 d\gamma^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}(x, y).$$

Next, we apply the pushforward definition of $\gamma^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}$ to the terms in the monotonicity inequality,

$$\int \|x - y\|^2 d\gamma^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}(x, y) = (\bar{\tau}_k - \bar{\tau}_{k-1})^2 \int \|x - y\|^2 d\gamma(x, y),$$

and

$$\begin{aligned} & \int \langle x - y, v[\rho^{(\bar{\tau}_{k-1})}](x) - v[\rho^{(\bar{\tau}_k)}](y) \rangle d\gamma^{(\bar{\tau}_{k-1}, \bar{\tau}_k)}(x, y) = \\ & (\bar{\tau}_k - \bar{\tau}_{k-1}) \int \langle x - y, v[\rho^{(\bar{\tau}_{k-1})}](I^{(\bar{\tau}_{k-1})}(x, y)) - v[\rho^{(\bar{\tau}_k)}](I^{(\bar{\tau}_k)}(x, y)) \rangle d\gamma(x, y). \end{aligned}$$

The monotonicity condition becomes, for every $k \in [1, \dots, \bar{K}]$,

$$\begin{aligned} & - \int \langle x - y, v[\rho^{(\bar{\tau}_{k-1})}](I^{(\bar{\tau}_{k-1})}(x, y)) - v[\rho^{(\bar{\tau}_k)}](I^{(\bar{\tau}_k)}(x, y)) \rangle d\gamma(x, y) \\ & \geq (\bar{\tau}_k - \bar{\tau}_{k-1}) \lambda \int \|x - y\|^2 d\gamma(x, y). \end{aligned}$$

Thanks to the integration against the same measure γ for every k , summing from $k = 1$ to $\bar{K}$ results in

$$\begin{aligned} \lambda \int \|x - y\|^2 d\gamma(x, y) & \leq - \int \langle x - y, v[\rho^{(0)}](I^{(0)}(x, y)) - v[\rho^{(1)}](I^{(1)}(x, y)) \rangle d\gamma(x, y) \\ & = - \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y). \end{aligned}$$

□

In order to extend Proposition 3.41 to all measures in $\overline{\mathcal{P}}_2$, we adopt the approximation method from [56], where $\rho^{(0)}, \rho^{(1)} \in \overline{\mathcal{P}}_2$ are approximated by $\rho^{(n,0)}, \rho^{(n,1)} \in \overline{\mathcal{P}}^{(f)}$ such that $\overline{W}(\rho^{(n,0)}, \rho^{(0)}) \rightarrow 0$ and $\overline{W}(\rho^{(n,1)}, \rho^{(1)}) \rightarrow 0$ as $n \rightarrow \infty$. In order to pass to the $n \rightarrow \infty$ limit in the monotonicity condition, we require that the velocity fields $v[\rho](x)$ satisfy an appropriate notion of continuity with respect to ρ and x .

Definition 3.42 (Strong-Weak Convergence). *For $x^{(n)}, v^{(n)} \in \mathbb{R}^d$, a sequence $(x^{(n)}, v^{(n)})$ converges to (x, v) in the strong-weak topology, denoted $(x^{(n)}, v^{(n)}) \xrightarrow{sw} (x, v)$, if*

$$x^{(n)} \rightarrow x, \quad f(v^{(n)}) \rightarrow f(v) \quad \forall f \in \mathcal{B}(\mathbb{R}^d),$$

where $\mathcal{B}(\mathbb{R}^d)$ is the set of linear bounded functionals on $\mathbb{R}^d$.

The set C_2^{sw} is given by

$$C_2^{sw} := \{ \zeta : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R} \mid \lim_{n \rightarrow \infty} \zeta(x^{(n)}, v^{(n)}) = \zeta(x, v) \quad \forall (x^{(n)}, v^{(n)}) \xrightarrow{sw} (x, v); \forall \varepsilon > 0 \\ \exists A_\varepsilon \geq 0 \text{ such that } |\zeta(x, v)| \leq A_\varepsilon(1 + \|x\|^2) + \varepsilon \|v\|^2, \forall (x, v) \in \mathbb{R}^d \times \mathbb{R}^d \}.$$

Definition 3.43 (Demicontinuity of v). *The velocity field $v : \mathcal{A} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is demicontinuous at $\rho \in \mathcal{A}$ if for any $\rho^{(n)} \rightarrow \rho$ in $(\mathcal{A}, \overline{W})$, it holds that*

$$\lim_{n \rightarrow \infty} \int \zeta(x, v[\rho^{(n)}](x)) d\rho^{(n)}(x) = \int \zeta(x, v[\rho](x)) d\rho(x), \quad \forall \zeta \in C_2^{sw}.$$

Assumption 3.44 (Demicontinuity of v). *The velocity field v is demicontinuous in $\mathcal{A}$.*

Corollary 3.45 (Admissible Couplings). *Let Assumption 3.44 hold. If $\overline{\mathcal{P}}^{(f)} \subseteq \mathcal{A}$, and for any $\rho^{(0)}, \rho^{(1)} \in \mathcal{A}$ with $\gamma^* \in \Gamma^*(\rho^{(0)}, \rho^{(1)})$ it holds*

$$- \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma^*(x, y) \geq \lambda \int \|x - y\|^2 d\gamma^*(x, y),$$

then for any $\gamma \in \Gamma(\rho^{(0)}, \rho^{(1)})$ such that $\rho^{(\tau)} \in \mathcal{A}$ for all $\tau \in [0, 1]$,

$$- \int \langle x - y, v[\rho^{(0)}](x) - v[\rho^{(1)}](y) \rangle d\gamma(x, y) \geq \lambda \int \|x - y\|^2 d\gamma(x, y).$$

Proof. The proof follows exactly as the approximation procedure in the proof of [55, Theorem 7.6], with the convergent sequences defined for each species. $\square$

In particular, if $\mathcal{A} = \overline{\mathcal{P}}_2$, then $\overline{\mathcal{P}}^{(f)} \subseteq \mathcal{A}$ and so Theorem 3.2 holds.

Remark 3.46. Corollary 3.45 does not necessarily hold for sets $\mathcal{A}$ which do not contain $\overline{\mathcal{P}}^{(f)}$. For example, consider the single-species setting where $v[\rho] = -\nabla \log \rho$. The domain of $v : L^1(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ excludes measures with finite support. Writing $v[\rho] = -\nabla \delta_\rho \mathcal{H}(\rho)$, where $\mathcal{H} = \int \rho \log \rho$, it holds that

$$\int \langle x - y, \nabla \log \rho(x) - \nabla \log \tilde{\rho}(y) \rangle d\gamma^*(x, y) \geq 0$$

for $\gamma^* \in \Gamma^*(\rho, \tilde{\rho})$, because $\mathcal{H}$ is 0-displacement convex over absolutely continuous measures. We check if the inequality still holds for the admissible plan $\gamma(x, y) = \rho(x)\tilde{\rho}(y)$:

$$\begin{aligned} \int \langle x - y, \nabla \log \rho(x) - \nabla \log \tilde{\rho}(y) \rangle d\rho(x)d\tilde{\rho}(y) &= - \int \langle x, \nabla \rho(x) \rangle dx - \int \langle y, \nabla \tilde{\rho}(y) \rangle dy \\ &\leq -2d < 0, \end{aligned}$$

so the monotonicity inequality does not hold for general $\gamma \in \Gamma(\rho, \tilde{\rho})$.

Chapter 4

DISSIPATIVE VARIATIONAL EVOLUTIONS IN METRIC SPACES

L. Conger, F. Hoffmann, and G. Savaré. Dissipative variational evolutions in metric spaces, 2026

ABSTRACT

We study the evolution variational inequality (EVI) driven by a bifunction $b : X \times X \rightarrow \mathbb{R}$ on a metric space (X, d_X) , providing conditions for existence of solutions, establishing solution properties, and characterizing asymptotic behavior. A motivating application is the setting of multispecies and minimax coupled gradient flows, in which each of N species evolves in the direction of steepest descent of its own energy functional and the joint system is not a gradient flow of any single energy. To construct solutions, we introduce a variational movement scheme (VMS), a time discretization scheme in which each update is a saddle point of a penalized bifunction, generalizing the classical minimizing movement scheme for single species gradient descent. In the coupled gradient flow case, the VMS reduces to a Nash equilibrium problem. We prove existence of discrete VMS solutions via two approaches: a primal approach using a metric-space version of the KKM theorem and compactness of superlevel sets, and a dual approach using two structural conditions—*barycentric convexity* and *interaction dissipativity*—which together provide a new zeroth-order notion of game-theoretic monotonicity applicable to general metric spaces without linear structure.

4.1 Introduction

The theory of gradient flows on a metric space (X, d_X) driven by an energy $\varphi : D \subset X \rightarrow \mathbb{R}$ with complete sublevels is by now classical. Among the various formulations, the strongest one involves a system of Evolution Variational Inequalities (EVI): for a locally Lipschitz curve $x : (0, \infty) \rightarrow D$ we ask that

$$\frac{1}{2} \frac{d}{dt} d_X^2(y, x(t)) + \frac{\kappa}{2} d_X^2(y, x(t)) \leq \varphi(y) - \varphi(x(t)), \quad \forall t \geq 0, y \in D, \quad (4.1)$$

where $\kappa \in \mathbb{R}$ is a parameter which is tailored to the second order derivative of φ along geodesics.

Sufficient conditions to solve (4.1) can be expressed through the generating function

$$\Phi(\tau, \bar{x}; x) := \frac{1}{2\tau} d_X^2(\bar{x}, x) + \varphi(x), \quad \tau > 0, \bar{x}, x \in D, \quad (4.2)$$

which describes the celebrated JKO (or Minimizing Movement) scheme [130, 9],

$$X_\tau^n \in \arg \min_{x \in D} \Phi(\tau, X_\tau^{n-1}; x). \quad (4.3)$$

If

$$\begin{aligned} &\text{every pair } x_1, x_2 \in D \text{ can be connected by a curve in } D \text{ along which} \\ &\text{the functional } x \mapsto \Phi(\tau, \bar{x}; x) \text{ is } (\tau^{-1} + \kappa) \text{ convex} \end{aligned} \quad (4.4)$$

then the MM scheme (4.3) is solvable and its solution converge to a locally Lipschitz curve solving (4.1) as the time step τ goes to 0. EVI and JKO have become canonical machinery for evolutions driven by a single energy [9]; condition (4.4) provides a purely metric criterion encompassing several relevant cases, including NPC spaces and functionals that are convex along generalized geodesics in Wasserstein spaces.

Many problems of interest, however, are not driven by a single energy. Saddle-point evolutions, coupled multispecies systems, and Nash-type games are described in product spaces $X = \prod X^{(i)}$ by an antisymmetric Lagrangian $F(y^{(1)}, x^{(2)}) - F(x^{(1)}, y^{(2)})$ or by N energies $(F^{(i)})$; they give rise to a **bifunction** $b : D \times D \rightarrow \mathbb{R}$, which reduces to the gradient flow case above when $b(y, x) = \varphi(y) - \varphi(x)$.

Taking also inspiration by the theory of contraction semigroups generated by monotone operators, this paper aims to extend the metric EVI / MM scheme machinery to bifunction-driven evolutions, where the bifunction b takes the place of φ and a new variational movement scheme takes the place of (4.3). We propose a new framework using a natural extension of (4.1) that tries to capture dissipativity at a metric level without any underlying Hilbertian structure, both providing existence of solutions and estimates for long-time behavior.

To that end, given a driving bifunction $b : D \times D \rightarrow \mathbb{R}$, we look for curves $x : (0, \infty) \rightarrow D$ that satisfy the EVI

$$\frac{d}{dt} \frac{1}{2} d_X^2(y, x(t)) + \frac{\kappa}{2} d_X^2(y, x(t)) \leq b(y, x(t)), \quad \text{for all } t \geq 0, y \in D. \quad (\text{EVI})$$

The value $\kappa \in \mathbb{R}$ depends on the (lack of) convexity of $b(\cdot, x)$.

The aim of this work is to provide general conditions under which for every initial datum $x_0 \in D$ there exists some solution to (EVI) starting from x_0 , determine its properties, and characterize its asymptotic behavior.

Several familiar settings fit into (EVI) through specific choices of b . We list them in order of increasing complexity.

(i) Single-species gradient flow. We have already seen that for an energy $\varphi : D \rightarrow \mathbb{R}$, setting

$$b(y, x) = \varphi(y) - \varphi(x), \quad \text{for all } x, y \in X,$$

(EVI) coincides with the usual EVI characterization of the metric gradient flow of φ (4.1).

(ii) Two-species min-max. $X = X^{(1)} \times X^{(2)}$ and we consider a Lagrangian $F : X^{(1)} \times X^{(2)} \rightarrow \mathbb{R}$, with one species minimizing and the other maximizing F , setting

$$b(y, x) = F(y^{(1)}, x^{(2)}) - F(x^{(1)}, y^{(2)}), \quad \text{for all } x = (x^{(1)}, x^{(2)}), y = (y^{(1)}, y^{(2)}) \in X.$$

Then (EVI) describes a saddle-point evolution. In a Hilbert space and for convex–concave F this corresponds to the evolution driven by the monotone operator associated with a convex–concave Lagrangian [190]. When F is regular, we have the system

$$\dot{x}^{(1)}(t) = -\nabla_1 F(x^{(1)}(t), x^{(2)}(t)), \quad \dot{x}^{(2)}(t) = \nabla_2 F(x^{(1)}(t), x^{(2)}(t))$$

whose solutions are characterized by (EVI).

(iii) Multispecies coupled gradient flow. In a system of N species, each species i is equipped with a metric space structure $(X^{(i)}, d_{X^{(i)}})$ and we consider the product space $X := \prod_{i=1}^N X^{(i)}$. Each species aims to minimize its own energy functional $F^{(i)} : X \rightarrow \mathbb{R}$, which may depend on the other species, and evolves in the direction of steepest descent for the species i . We use the notation $F^{(j)}(x^{(i)}, y^{(-i)}) := F^{(j)}(y^{(1)}, \dots, y^{(i-1)}, x^{(i)}, y^{(i+1)}, \dots, y^{(n)})$ to denote that the i^{th} entry in $F^{(j)}$ is $x^{(i)}$, and all other entries are the entries of y . Selecting the bifunction

$$b(y, x) = \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}), \quad \text{for all } x, y \in X, \quad (4.5)$$

results in coupled gradient flow dynamics for any curve $x(t)$ solving the EVI (see Lemma 4.58). Here, while each species evolves in the direction of maximal slope, the joint system is not necessarily evolving according to a gradient flow, because there is not necessarily a single energy that parameterizes the evolution. Additionally, the individual energies can increase, decrease, or oscillate in time even while the entire system contracts to a unique steady state. Moreover, standard notions of convexity for each $F^{(i)}(\cdot, y^{(-i)})$ are not sufficient for guaranteeing existence of joint critical points, which are candidates for steady states of the system.

The sense in which convexity is defined is not necessarily along geodesics or generalized geodesics as in the Wasserstein-2 gradient flow case; the choice of curves depends on the convexity properties of both $(F^{(i)})_{i=1}^N$ and $d_X^2(\cdot, x)$.

Interaction dissipativity. It is clear that stability properties of solutions to (EVI) should be related to suitable metric dissipativity conditions on b involving $b(x, y) + b(y, x)$. One can easily check that the bifunction b of the previous examples (i) and (ii) satisfies the structural condition

$$b(x, y) + b(y, x) = 0 \quad \text{for every } x, y \in D. \quad (4.6)$$

Such a condition is particularly useful to prove an exponential stability estimate among two solutions x, y of (EVI) of the form

$$d_X(x(t), y(t)) \leq e^{\lambda(t-s)} d_X(x(s), y(s)) \quad 0 \leq s \leq t, \quad (4.7)$$

for $\lambda := -\kappa$. We will systematically adopt a more general and natural condition, which reveals sufficiently flexible to cover interesting examples such as (iii); it can be expressed by the η -interaction dissipativity condition

$$b(x, y) + b(y, x) \leq \eta d_X^2(x, y) \quad \text{for every } x, y \in D, \quad (4.8a)$$

which yields (4.7) for $\lambda := \eta - \kappa$. We will also show that such a condition, coupled with the diagonal property

$$b(x, x) = 0 \quad \text{for every } x \in D, \quad (4.8b)$$

is useful to study the solvability of the discrete scheme we propose to solve (EVI).

The Variational Movement Scheme and the Discrete EVI. In order to discretize (EVI), we look for a time-stepping scheme whose iterates satisfy a discrete analogue of the EVI. We have seen that in the case (i) of the single-species gradient flow $b(y, x) = \varphi(y) - \varphi(x)$, the JKO scheme is obtained by a recursive minimization associated to the generating function Φ (4.2).

For a general bifunction b , minimization of a single functional is no longer available; we propose to replace it with a system of variational inequalities associated with the generating bifunction

$$B(\tau, \bar{x}; y, x) := \frac{1}{2\tau} d_X^2(\bar{x}, y) - \frac{1}{2\tau} d_X^2(\bar{x}, x) + b(y, x), \quad (4.9)$$

which in case (i) reduces to $B(\tau, \bar{x}; y, x) = \Phi(\tau, \bar{x}; y) - \Phi(\tau, \bar{x}; x)$. The scheme reads:

Variational movement scheme (VMS).

Given $X_\tau^{n-1} \in D$, find $X_\tau^n \in D$ such that

$$0 \leq B(\tau, X_\tau^{n-1}; Y, X_\tau^n) \quad \text{for every } Y \in D.$$

We will show that, under general assumptions on b and on (X, d_X) , involving a reinforced convexity condition of the map $y \mapsto B(\tau, \bar{x}; y, x)$, which can be considered as a natural counterpart of (4.4), the scheme admits at least one solution at every step. Moreover, under the η -interaction dissipativity conditions (4.8a), (4.8b), solutions to (VMS) can be equivalently formulated in terms of the dual variational inequality

$$B(\tau, X_\tau^{n-1}; X_\tau^n, X) \leq 0 \quad \text{for every } X \in D, \quad (4.10)$$

so that (X_τ^n, X_τ^n) is a saddle point of $(x, y) \mapsto B(\tau, X_\tau^{n-1}; y, x)$ over $D \times D$.

Still assuming that $y \mapsto B(\tau, \bar{x}; y, x)$ satisfies the $(1/\tau + \kappa)$ -convexity condition as in (4.4), we can show that any solution $(X_\tau^n)_n$ of (VMS) satisfies the discrete EVI

$$\frac{1}{2\tau} \left(d_X^2(X_\tau^n, Y) - d_X^2(X_\tau^{n-1}, Y) \right) + \frac{1}{2\tau} d_X^2(X_\tau^n, X_\tau^{n-1}) + \frac{\kappa}{2} d_X^2(X_\tau^n, Y) \leq b(Y, X_\tau^n), \quad \text{for every } Y \in D, \quad n \geq 1. \quad (4.11)$$

As in the gradient flow setting, piecewise constant interpolant constructed from (VMS) converges, as $\tau \rightarrow 0$, to a solution of (EVI).

The Variational Movement Scheme has a nice interpretation in the reference cases (i)-(ii)-(iii). First, it is easy to check that in case (i) the scheme reduces to a JKO step. In the saddle-Lagrangian case (ii), every solution $X_\tau^n = (X_\tau^{n,(1)}, X_\tau^{n,(2)})$ is a saddle point of the penalized Lagrangian

$$(x_1, x_2) \mapsto \frac{1}{2\tau} d_{X^{(1)}}^2(X_\tau^{n-1,(1)}, x_1) - \frac{1}{2\tau} d_{X^{(2)}}^2(X_\tau^{n-1,(2)}, x_2) + F(x_1, x_2).$$

When b is selected to parameterize a multispecies coupled gradient flow (4.5) as in (iii), the variational inequality is, in fact, a Nash equilibrium problem. A Nash equilibrium x_* of a set of energy functionals $(F^{(i)})$ satisfies

$$F^{(i)}(x_*^{(i)}, x_*^{(-i)}) \leq F^{(i)}(y^{(i)}, x_*^{(-i)}) \quad \text{for every } y \in X, \quad \text{for every } i \in [1, \dots, N].$$

A Nash equilibrium $x_* \in X$ is such that no species can decrease its energy, given that all the others are fixed.

We can then form the corresponding distance-penalized functions as in (4.2)

$$\Phi^{(i)}(\tau, \bar{x}; x^{(i)}, x^{(-i)}) := F^{(i)}(x^{(i)}, x^{(-i)}) + \frac{1}{2\tau} d_{X^{(i)}}^2(x^{(i)}, \bar{x}^{(i)}) \quad (4.12)$$

It is immediate to check that X_τ^n is a solution to one step of (VMS) starting from X_τ^{n-1} if and only if

$$X_\tau^n \text{ is a Nash equilibrium of } \left\{ x \mapsto \Phi^{(i)}(\tau, X_\tau^{n-1}; x^{(i)}, x^{(-i)}) \right\}_i,$$

Existence of Discrete VMS Solutions. A number of challenges arise in proving existence of solutions to VMS. If the metric space has some linear structure and $y \mapsto b(y, x)$ has compact convex sublevels, one can hope to apply general existence theorems for solutions of systems of variational inequalities (see, e.g., [14]). In the multispecies coupled gradient flow setting with an unbounded domain, from the finite-dimensional game theory literature we expect such an equilibrium point to exist under *strong monotonicity*, which is a first-order (gradient) condition on the energy functionals that guarantees contraction of coupled gradient descent dynamics [64]. However, in general metric spaces, we prefer a zeroth-order notion of monotonicity, since the appropriate notion of gradient is often difficult to determine.

To address these challenges, we employ a notion of *barycentric convexity*, inspired by (4.4), that provides a surrogate of a linear structure, and *interaction dissipativity*, which enables contraction of the dynamics.

Barycentric convexity is a generalization of convexity along a curve which connects two points in X to an interpolation among J points in X . We thus consider systems $\mathcal{B}_J$ of barycentric maps $\mathbf{x} : \Delta_J \rightarrow D$ from the $J - 1$ dimensional simplex Δ_J to $D \subset X$ such that for every choice of J points $(x_j)_{j=1}^J$ in D there exists a map $\mathbf{x} \in \mathcal{B}$ such that $x_j = \mathbf{x}(e_j^J) =: \mathbf{x}_j$, where $(e_j^J)_j$ is the canonical basis of $\mathbb{R}^J$. We say that $\varphi : D \rightarrow \mathbb{R}$ is *barycentrically κ -convexlike* with respect to a system $(\mathcal{B}_J)_{J \in \mathbb{N}_{\geq 2}}$ of barycentric maps if for every integer $J \in \mathbb{N}_{\geq 2}$, every $\mathbf{x} \in \mathcal{B}_J$, and every $\alpha \in \Delta_J$,

$$\varphi(\mathbf{x}(\alpha)) \leq \sum_{j=1}^J \alpha_j \varphi(\mathbf{x}_j) - \frac{\kappa}{4} \sum_{j,k=1}^J \alpha_j \alpha_k d_X^2(\mathbf{x}_j, \mathbf{x}_k).$$

In the case where $J = 2$, the convexity condition becomes the standard two-point convexity inequality along a curve parameterized by $\alpha = (\alpha, 1 - \alpha)$, and if the curve parameterized by α is a geodesic, the inequality corresponds to geodesic κ -convexity.

The interaction dissipativity condition, in combination with barycentric κ -convexity of $y \mapsto b(y, x)$, provides a definition of game-theoretic monotonicity for games in general metric spaces. We say that a game with energies $(F^{(i)})$ is $-\lambda$ -monotone with respect to $\mathcal{B}$ with $\lambda = \eta - \kappa$ if

- (κ convexity) $y \mapsto b(y, x)$ is barycentrically κ -convexlike with respect to a system $\mathcal{B}$.
- (η interaction dissipativity) $b(y, x) + b(x, y) \leq \eta d_X^2(x, y)$.

We show that, for particular choices of $\mathcal{B}$, this notion of game-theoretic monotonicity implies the classical definitions in Euclidean space and in the space of measures equipped with the Wasserstein-2 metric, as defined in [64, 206]. In [64], monotonicity is defined with respect to geodesic curves, whereas in [206] it is defined along linear interpolations of probability measures. Using different choices of $\mathcal{B}$ allows for the zeroth order definition to imply each of these first-order conditions, and thus generalizes previous notions of monotonicity.

Equipped with barycentric convexity and interaction dissipativity, we take two approaches for showing existence of the discrete VMS: one via compactness, and one via interaction dissipativity. The first provides a primal solution to the variational equilibrium problem (VMS), which requires compactness of superlevel sets of $b(y, \cdot)$. The linear structure of the simplex Δ_J allows us to prove a version of the Knaster–Kuratowski–Mazurkiewicz (KKM) theorem in the metric space setting. While the required compactness condition for the first approach is more general than the assumptions in the second approach, it can be difficult to show. The second approach uses the dual formulation of the variational equilibrium problem, and relies on barycentric κ -convexity of $y \mapsto b(y, x)$ and η interaction dissipativity, both of which replace the compactness assumption needed in the primal formulation.

Discrete to Continuous Limit. In order to take $\tau \rightarrow 0$ for the sequence $(X_\tau^n)_n$, the continuous-time piecewise constant interpolant $\bar{X}_\tau : [0, \infty) \rightarrow X$ is constructed,

$$\bar{X}_\tau(0) := X_\tau^0, \quad \bar{X}_\tau(t) := X_\tau^n \quad \text{if } t \in ((n-1)\tau, n\tau].$$

We show a Cauchy condition for piecewise-constant interpolations of (X_τ^n) which results in the existence of a limit $x(t)$ as $\tau \rightarrow 0$. This estimate is achieved adapting the doubling integration method from [175] to this setting. See also [69, 70, 135] for the Crandall-Liggett error estimation method for Banach spaces. Then, taking the limit of the discrete EVI, the continuous EVI is recovered. The interaction dissipativity inequality in combination with the continuous EVI enables the $\lambda = \eta - \kappa$ dissipativity estimate

$$d_X(y(t), x(t)) \leq e^{\lambda(t-s)} d_X(y(s), x(s)) \quad \text{whenever } 0 \leq s \leq t.$$

For any curves $x(t), y(t)$ satisfying the continuous EVI, when $\eta < \kappa$ and (D, d) is complete, the contraction enables us to establish the existence of a unique steady state. In the multispecies gradient flow setting, we can also show that the steady state is the unique Nash equilibrium of $(F^{(i)})$ (Corollary 4.54).

To understand the roles of η and κ in determining λ dissipativity, consider the following three examples for b associated with a min-max flow, a monotone flow, and a general multispecies coupled gradient flow in a Hilbert space:

$$\begin{aligned} b(y, x) &= F(y^{(1)}, x^{(2)}) - F(x^{(1)}, y^{(2)}), & \eta &= 0, \lambda = -\kappa \\ b(y, x) &= \langle A(x), y - x \rangle, & \kappa &= 0, \lambda = \eta \\ b(y, x) &= \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}), & \lambda &= \eta - \kappa. \end{aligned}$$

In the first example, we have that $b(x, y) + b(y, x) = 0$, so $\eta = 0$ and the dissipativity is entirely dependent on the convexity of $y \mapsto b(y, x)$. In the case of a convex-concave Lagrangian F we also have $\kappa = 0$.

In contrast, for the second example, $y \mapsto b(y, x)$ is 0-convex, so the dissipativity comes entirely from the interaction dissipativity

$$b(x, y) + b(y, x) = -\langle A(x) - A(y), x - y \rangle.$$

In the last example, both κ convexity and η interaction dissipativity contribute to the overall system dissipativity. In this paper, we prove that in the multispecies coupled gradient flow setting, $\eta \geq 0$. This means that, for an uncoupled gradient flow system of N species, adding interaction terms to the energies cannot improve the convergence rate, only making it worse (Proposition 4.27).

Notice that only the second case need a linear Hilbertian structure.

EVI: $\frac{d}{dt}d_X^2(y, x(t)) + \frac{\kappa}{2}d_X^2(y, x(t)) \leq b(y, x(t))$		Metric Spaces
Parameters:	Hilbert Spaces	
	– accretivity = – monotonicity = dissipativity	
1. η interaction dissipativity $b(x, y) + b(y, x) \leq \eta d_X^2(x, y)$.	Example: $X = (H, \ \cdot\)$, $b(y, x) = \langle A(x), y - x \rangle$.	Example: $X = (H, \ \cdot\)$, $b(y, x) = \varphi(y) - \varphi(x)$.
2. κ convexity $y \mapsto b(y, x)$ is κ -convex.	1. $\langle A(x) - A(y), x - y \rangle \geq -\eta \ x - y\ ^2$.	1. $b(x, y) + b(y, x) = 0$; $\eta = 0$.
3. $\lambda = \eta - \kappa$ dissipativity $d_X(y(t), x(t)) \leq e^{\lambda t} d_X(y(0), x(0))$.	2. $y \mapsto b(y, x)$ is $\kappa = 0$ convex.	2. φ is κ -convex.
	3. A is $-\lambda = -\eta$ accretive.	3. The gradient flow system is $-\lambda = \kappa$ accretive.

Figure 4.1: The nomenclature and coefficients for convexity, monotonicity, accretivity, and dissipativity as used in this work are illustrated in particular settings.

Dissipativity, monotonicity, accretivity, and convexity. The solution x to the EVI is $\lambda = \kappa - \eta$ dissipative if b is η -interaction dissipative and $b(\cdot, x)$ is barycentrically κ -convexlike; for $\lambda < 0$, the solution is contractive. This is a generalization of several well-established notions. In Hilbert spaces, the negative of dissipativity is called monotonicity, that is, a system that is λ dissipative is $-\lambda$ monotone. This sign convention is also used in game theory, where a monotone system is contracting if it has a positive monotonicity parameter. A positive κ convexity parameter corresponds to a strongly convex function, carrying the standard sign convention for convexity. A summary of this nomenclature is provided in Figure 4.1, along with the examples illustrating the various convexity and dissipativity parameters. In this work, we will use the term *dissipativity* to describe the contractivity. In the particular setting of multispecies coupled gradient flows, we will use the term *game-theoretic monotonicity* for notions parallel with monotonicity in the game theory literature, using a consistent sign for λ throughout: $\lambda < 0$ means a contracting system.

Running Examples. Throughout this paper, we illustrate the intermediate results on two examples. The first example is a PDE in a Hilbert space driven by an operator $A : H \rightarrow H$, and the second example is a multispecies coupled Wasserstein-2 gradient flow system. A number of other examples are included throughout the paper, and Section 4.6 contains specific choices of energy functionals for multispecies gradient flows in Euclidean and Wasserstein-2 spaces.

Example (I). *PDE in Hilbert space:* $X = H$, $d_X(x, y) = \|x - y\|$, $A : H \rightarrow H$. The solution $x(t)$ satisfies (weakly)

$$\frac{d}{dt}x(t) + A(x(t)) = 0, \quad b(y, x) := \langle A(x), y - x \rangle.$$

For instance, with $H = W_0^{1,2}(\mathbb{R}^d)$ a suitable Sobolev space, and $A = -\Delta$ the Laplacian operator, we obtain the heat equation.

Example (II). *Coupled Wasserstein-2 Gradient Flows:* $X^{(i)} = \mathcal{P}_2(\mathbb{R}^d)$, $d_{X^{(i)}}(x^{(i)}, y^{(i)}) = \mathcal{W}_2(x^{(i)}, y^{(i)})$. The solution $x(t)$ satisfies (weakly) the coupled Wasserstein-2 gradient flow system

$$\begin{aligned} \frac{d}{dt}x(t) &= - \begin{bmatrix} \nabla_{\mathcal{W}_2, x^{(1)}} F^{(1)}(x) \\ \vdots \\ \nabla_{\mathcal{W}_2, x^{(n)}} F^{(n)}(x) \end{bmatrix} = \begin{bmatrix} \operatorname{div} \left(x^{(1)} \nabla_{\delta_{x^{(1)}}} F^{(1)}(x) \right) \\ \vdots \\ \operatorname{div} \left(x^{(n)} \nabla_{\delta_{x^{(n)}}} F^{(n)}(x) \right) \end{bmatrix}, \\ b(y, x) &= \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}). \end{aligned}$$

Throughout this example, we will suppose that $(F^{(i)})_i$ are $-\lambda$ -monotone according to [64] in some subset $\mathcal{A} \subset \mathcal{P}_2(\mathbb{R}^d)$ of probability measures with bounded second moment. $-\lambda$ -monotonicity




Figure 4.2: A multispecies coupled gradient flow system in the Wasserstein-2 metric is not necessarily a joint gradient flow system. Particular choices of energies $F^{(i)}$ result in min-max flows or a single-species gradient flow.

implies λ -dissipativity in the sense that $\mathcal{W}_2(x(t), y(t)) \leq e^{\lambda t} \mathcal{W}_2(x(0), y(0))$ for any two solutions x, y (if they exist) to the coupled gradient flow problem. This setting includes the min-max flow setting and the single-species setting, as illustrated in Figure 4.2. We show formally in Lemma 4.58 that solutions to coupled Wasserstein-2 dynamics satisfy the EVI for the choice of b above.

4.1.1 Contributions

The main contributions of this paper are:

1. Existence of solutions to the EVI driven by bifunction b , which includes multispecies coupled gradient flows.
2. Existence of solutions to variational equilibrium problems in general metric spaces using barycentric convexity.
3. A zeroth-order definition of game-theoretic monotonicity for games in general metric spaces.

The existence of solutions results are summarized in the following theorem.

Theorem 4.1 (Existence and Contraction Estimate). *Let $\tau_o > 0$ such that $\tau_o^{-1} + \kappa > 2\eta$ and suppose that*

- *for every $\tau \in (0, \tau_o)$, every $\bar{x}, x \in D$, the functions $y \mapsto B(\tau, \bar{x}; y, x)$, given by*

$$B(\tau, \bar{x}; y, x) = \frac{1}{2\tau} d_X^2(\bar{x}, y) - \frac{1}{2\tau} d_X^2(\bar{x}, x) + b(y, x),$$

are barycentrically $(\tau^{-1} + \kappa)$ convexlike with respect to some barycentric system $\mathcal{B}$.

- *$x \mapsto B(\tau, \bar{x}; y, x)$ is upper hemicontinuous with respect to $\mathcal{B}_2$, and*
- *b is η interaction dissipative according to (4.8a)-(4.8b).*

Then for any initial condition $x_0 \in D$, there exists a solution $(X_\tau^n)_{n \geq 0}$ to (VMS), which satisfies the discrete EVI (4.11). There exists a continuous curve $x : [0, +\infty) \rightarrow D$ with $x(0) = x_0$ that satisfies the EVI

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), y) + \frac{\kappa}{2} d_X^2(x(t), y) \leq b(y, x(t)) \quad t \in (0, +\infty), \quad \text{for every } y \in D.$$

Any continuous curves x, y which satisfy the EVI above also satisfy the contraction estimate

$$d_X(x(t), y(t)) \leq e^{\lambda(t-s)} d_X(x(s), y(s)), \quad \text{for all } t \geq s > 0.$$

We present a zeroth-order definition of game-theoretic monotonicity using barycentric convexity and interaction dissipativity; a larger convexity parameter κ increases monotonicity whereas a larger interaction dissipativity parameter η decreases monotonicity. We show in Lemmas 4.55 and 4.57 that Definition 4.2 implies Euclidean and Wasserstein-2 monotonicity.

Definition 4.2 (Game-Theoretic Monotonicity). *A set of energy functionals $(F^{(i)})$ are $-\lambda$ -monotone, where $\lambda = \eta - \kappa$, if*

- κ convexity: $y \mapsto b(y, x)$ is barycentrically κ -convexlike with respect to a system $\mathcal{B}$.
- η dissipativity: $b(y, x) + b(x, y) \leq \eta d_X^2(x, y)$.

Equivalently, each $F^{(i)}(\cdot, \cdot) - F^{(i)}(x^{(i)}, \cdot)$ is barycentrically κ -convex and for every $x, y \in D$,

$$\sum_{i=1}^N F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}) + F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) \leq \eta d_X^2(x, y).$$

4.1.2 Related Literature

This work extends results for accretive and dissipative evolutions to the metric space setting, using variational equilibrium problem tools. An important application of the results is the setting of multispecies gradient flows, which can be studied with a game theory point of view, and for which general existence results were left as an open question in [64]. Related literature from each of these fields is discussed in this section.

Accretive and Dissipative Evolutions. The theory of nonlinear evolutions driven by accretive operators in Banach and Hilbert spaces provides a classical foundation for the single-species special case of our framework. The foundational Crandall–Liggett generation theorem [71] established that accretive operators on general Banach spaces generate semigroups, with solutions obtained as limits of implicit time-discretizations. The connection between accretive operators and evolution

inequalities in Hilbert spaces was further developed by Brézis [32]. For the case of λ -dissipative systems (corresponding to strongly accretive operators), contraction and long-time behavior estimates analogous to our dissipativity inequality were obtained by Kobayashi [135]; see also Crandall [69] for a broader survey of nonlinear semigroup theory and its connections to accretive operators. The optimal error estimates between piecewise-constant interpolants in the $\tau \rightarrow 0$ limit that we rely upon are adapted from those in [175], which extended the earlier error analysis of Crandall and Evans [70] to derive optimal a posteriori error estimates; the method of doubling the time variable used therein originates in [135].

Variational Equilibrium Problems. The existence theory for saddle points and minimax problems in convex-concave settings is classical, with foundational contributions due to von Neumann [205]. The connection between the sub-super differential of a convex-concave Lagrangian and monotonicity has been developed by Rockafellar [189]. The general theory of variational inequalities was developed in the works of Browder [34] and later unified with fixed-point and KKM theory; see Fan [106] for the Fan–KKM theorem, which underlies our primal existence proof (Theorem 4.16).

Gradient Flows and Systems of Coupled Gradient Flows. The development of the theory of gradient flows in Wasserstein and metric spaces received a decisive impetus from [130, 177] which introduced what is now called the JKO scheme and provided many tools and geometric insights to interpret a large class of diffusion equations as gradient flows in the Wasserstein-2 metric. In the broader metric setting, [9] provides a rigorous theory for gradient flows in general metric spaces, revisiting De Giorgi’s approach to curves of maximal slope [76] and introducing the notion of EVI (4.1) and convexity along generalized geodesics, motivated by applications in the Wasserstein spaces inspired by [130]. A systematic study of EVI properties for gradient flows has been developed by [171].

When more than one species evolves according to the direction of steepest descent of its own energy, the joint system is not necessarily a gradient flow. This occurs, for example, in cross-diffusion systems. The work [112] proved existence of solutions for two-species aggregation systems using a semi-implicit JKO scheme in which each species minimizes its own energy with the other held fixed, an approach equivalent to the implicit-explicit update introduced above rather than the fully implicit Nash update. The extension to diffusive systems was treated in [21]. The coupling-induced loss of contractivity is a central theme in [19], which studies two Wasserstein gradient flows coupled through a cross-diffusion term; while coupling destroys the individual contractivity, exponential convergence to a unique steady state is recovered for sufficiently small coupling. Other coupled multispecies Wasserstein-2 gradient flow well-posedness works include [86, 49, 54]; see also references within. For settings with an infinite number of species, well posedness was studied in [80, 43].

In other cases, long-time behavior of coupled systems has been shown while leaving existence as an open problem. In the min-max setting, [63] shows conditions under which a min-max Wasserstein gradient flow converges to a unique steady state. The convergence of two-timescale gradient descent-ascent in a nonconvex setting is analyzed in [10], which establishes long-time convergence. Similarly, in the setting of measure optimization, [97] studies a class of functional saddle-point problems, modeling the dynamics as a Fisher-Rao-RKHS gradient flow, proposing a primal-dual kernel mirror prox algorithm and providing convergence guarantees. Existence of solutions in the min-max setting was then addressed in [127], and our work generalizes their results by moving beyond the min-max setting. The long-time behavior of general multispecies setting in the Wasserstein-2 metric was analyzed in [64], and this paper provides conditions under which solutions exist to that system. Similarly, [206] considers a multispecies coupled Wasserstein-2 gradient flow with potential functions and entropy, and our work provides conditions under which solutions exist to the evolution studied there. Other coupled gradient flow systems that fall within our more general framework include Cahn-Hilliard/Allen-Cahn system with cross-kinetic coupling, with existence studied in [37] using Galerkin approximations.

There are coupled gradient flow settings where our results cannot be applied, either due to convexity issues with the metric or with the energies. Two recent works address spherical Hellinger gradient flows, also known as Fisher-Rao gradient flows. In [46], the authors present convexity properties of f -divergence energy functionals along geodesics and in [143], the authors analyze the long-time behavior of min-max flows, both leaving existence of solutions as future work. However, we cannot directly apply our theory to these systems due to the positive curvature of the spherical Hellinger metric. In cases where the joint energy is not jointly convex—such as cross-diffusion systems in which the pressure depends on the total density—our approach cannot be directly applied. In [99], a cross-diffusion system for two interacting populations is studied in which diffusion is governed by the aggregate density through a fast-diffusion pressure law; the non-convexity of the pressure energy is a central obstruction to a direct gradient flow treatment, and global weak solutions are obtained under a mixing condition on the initial data. The existence theory for this same system was subsequently completed in full generality, without restrictions on the initial data, in [165]. The present framework does not directly apply to these settings without convexity along some set of curves for each of the individual energies.

Game Theory. In finite dimensional games, the existence of Nash equilibria follows from compactness and continuity assumptions via the Kakutani-Brouwer fixed-point theorem [174, 82], and first-order monotonicity conditions on cost functions in games have also been used to prove existence and uniqueness [191]. For a summary of Nash equilibrium existence results and their connections to variational inequalities, see [103]; our notion of game-theoretic $-\lambda$ -monotonicity

extends these definitions to metric spaces. In infinite-dimensional settings, results typically rely on compactness assumptions and linear structure in the metric space [120] and our aim is to provide results for spaces without such a structure. For an extensive treatment of Nash equilibria in unbounded, strongly monotone Euclidean settings, see [104]. In mean field games, a notion of *displacement monotonicity* was recently proposed in [113], with many works such as [164] generalizing their results; this is a convexity property along Wasserstein-2 geodesics. The earlier *Lasry-Lion monotonicity* [145] is a convexity property along linear interpolants. Both of these definitions are for a single-species mean field game, with dynamics that are not a coupled gradient flow system.

4.1.3 Plan of the paper

In Section 4.2, we provide preliminary information such as notation and definitions used throughout and introduce the two running examples. Next, Section 4.3 introduces the main variational equilibrium problems and results for existence of solutions to those problems. In Section 4.4, we define notions of solutions and the corresponding EVIs; the proof of the Cauchy inequality needed for taking the time step to zero to obtain the continuous EVI from the discrete EVI is postponed to Appendix 4.A. Section 4.5 applies the equilibrium problem results to the discrete EVI. Finally, Section 4.6 contains example systems to which our results can be applied and shows how our zeroth-order notion of game theoretic monotonicity implies the classical first-order condition in Euclidean space and recent notions in the Wasserstein-2 space.

List of main notation

The notation used in this paper is listed in the table below.

A	monotone operator
A	subset of D
α, β	$\alpha \in \Delta_J$ point in the simplex
$\mathfrak{B}$	Borel sets
b	(bi)function $b : X \times X \rightarrow \mathbb{R}$
b_A	$b_A(y) := \sup_{x \in A} b(y, x)$
$\mathcal{B}$	set of $\mathcal{B}_J$ for all $J \in \mathbb{N}_{\geq 2}$
$\mathcal{B}_J$	barysystem with J reference points
$\mathcal{B}_{\bar{x}, 2}$	barysystem with 2 reference points and dependent on $\bar{x}$
B	bifunction in variational movement scheme
B	subset of $[1, \dots, J]$
D	subset of metric space X
D^J	J -product space of subset of metric space X : $D^J = D \times \dots \times D$

$ D^-\varphi (x),$	descending and ascending metric slopes
$ D^+\varphi (x)$	
d_X	metric on X
δ	first variation
Δ_J	simplex for J points
$\Delta^\lambda(z)$	bifunction slope
$(e_j^i)_{i=1}^J$	canonical basis of $\mathbb{R}^J$
E	evaluation map from $\mathcal{B}_J$ to D_J
$F^{(i)}$	energy for i^{th} species, maps from X to $\mathbb{R}$
φ	function from X to $\mathbb{R}$
$\mathcal{G}_2$	set of geodesic curves, particular choice of $\mathcal{B}_2$
$\Gamma(\mu, \nu), \Gamma^*(\mu, \nu)$	set of joint probability measures with marginals μ and ν (with $*$, Wasserstein-2 optimal plans)
H	Hilbert space
J	positive integer, number of barysystem reference points
K	interaction kernel
κ_i, κ	convexity parameter for i^{th} species, convexity of b
λ	dissipativity
$\mathcal{O}_{\bar{\mu}, J}$	collection of $J + 1$ objects in a Hilbert space with reference point $\bar{\mu}$
ℓ, ℓ_τ	time interpolators
η	interaction dissipativity
N	number of species in multispecies setting
$\mathcal{P}, \mathcal{P}_p$	space of probability measures, subscript p restricts to measures with finite p^{th} moment
$\mathcal{P}$	a particular probability measure
σ_-, σ^+	sublevel and superlevel sets of b
S, S_t	semigroup operators
$\mathcal{S}$	subset of simplex
S, T	finite time horizon
τ	time step in VMS
$\vartheta \in [0, 1]$	interpolation variable
$ \dot{u}^\pm (t), \dot{u}_\pm (t)$	left/right, upper/lower metric derivatives
μ, ν	measures
$\bar{\mu}, \bar{x}$	barycenter (in measure spaces or in general metric spaces, respectively)
v, v_i, w	elements in vector space
V	vector space

$\mathcal{W}$	Wasserstein metric
X	joint product space over which metric space is defined
$X^{(i)}$	metric space over which i^{th} species is defined
x	curve in X , $x(t) \in X$
$\mathbf{x}$	barycentric map from simplex to D
$\mathbf{x}, (x_i)_{i=1}^m$	collection of J (m respectively) points in X
$x, x^{(i)}, x^{(-i)}$	point in X , point in $X^{(i)}$, point in $\prod_{j \neq i} X^{(j)}$
x^*, x_*	primal and dual variational problem solutions
Y, X_τ^n	object in discrete time, sequence in X parameterized by τ with index n
$\bar{X}_\tau(t), \underline{X}_\tau(t)$	piecewise constant right/left continuous interpolants of X_τ^n
(Z, d_Z)	nonpositively curved metric space

4.2 Preliminaries: metric spaces, slopes, and convexlike functions

In this section, we establish definitions of convexity in general metric spaces, along with various notions of slopes. We illustrate the convexity notions in the settings of Examples I and II. The convexity definitions center around the idea of barycentric convexity, which generalizes convexity along a curve connecting two points to an interpolation between an arbitrary number of points. The primary motivation for this formulation is to show the existence of solutions to the equilibrium problem that appears in the time-discretized approximation of flow satisfying the EVI (EVI). In particular, to prove existence of the equilibrium solution, we need to show a finite intersection property of superlevel sets $\sigma^+(y) := \{x \in D : b(y, x) \geq 0\}$, which we will achieve using interpolations of an arbitrary number J of points in X . By showing convexity of $b(y, x)$ with respect to y over an interpolation of J points, the finite intersection property can be shown (see the proof of Theorem 4.22). The secondary motivation for this formulation, just as in the classical gradient flow setting, is that the notion of convexity must be compatible with 1-convexity of the distance function, requiring flexibility for the choice of curves along which convexity is defined. In the case of Wasserstein-2 gradient flows, generalized geodesics are used to obtain $\frac{1}{\tau} + \kappa$ convexity of an energy functional plus the $\frac{1}{2\tau}$ -weighted distance term. However, there is not necessarily a notion of generalized geodesics for other metric spaces. In these settings, we identify barycentric convexity as a suitable, flexible notion of convexity.

4.2.1 Barycentric Convexity and Related Notions

Let (X, d_X) be a metric space. A geodesic in X is a curve $x : [0, 1] \rightarrow X$ such that

$$d_X(x(s), x(t)) = |t - s|d_X(x(0), x(1)) \quad \text{for every } s, t \in [0, 1]. \quad (4.13)$$

Given $x_0, x_1 \in X$ and $\vartheta \in (0, 1)$, we say that x_ϑ is a ϑ -intermediate point between x_0 and x_1 if

$$d_X(x_0, x_\vartheta) = \vartheta d_X(x_0, x_1) \quad d_X(x_\vartheta, x_1) = (1 - \vartheta) d_X(x_0, x_1), \quad \text{for all } \vartheta \in [0, 1].$$

We denote by $(e_j^i)_{i=1}^J$ the canonical basis of $\mathbb{R}^J$ and by Δ_J the $J - 1$ -dimensional probability simplex in $\mathbb{R}^J$:

$$\Delta_J := \text{co}\{e_j^1, e_j^2, \dots, e_j^J\} = \left\{ \alpha = (\alpha_1, \dots, \alpha_J) \in \mathbb{R}^J : \alpha_j \geq 0, \sum_{j=1}^J \alpha_j = 1 \right\}. \quad (4.14)$$

A J -barycentric map in $D \subset X$ is a map $\mathbf{x} : \Delta_J \rightarrow D$. We set $\mathbf{x}_i := \mathbf{x}(e_j^i)$. A system $\mathcal{B}_J$ of J -barycentric maps is a collection of J -barycentric maps $\mathbf{x}$ such that the evaluation map

$$E : \mathcal{B}_J \rightarrow D^J, \quad E[\mathbf{x}] := (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_J) \quad \text{is surjective} \quad (4.15)$$

i.e. for every collection of J points in D , denoted $\mathbf{x} \in D^J$, there exists $\mathbf{x} \in \mathcal{B}_J$ such that $x_j = \mathbf{x}_j = \mathbf{x}(e_j^i)$, $j = 1, \dots, J$. A system $\mathcal{B}$ of barycentric maps is the union of J -systems $\mathcal{B}_J$, for every $J \in \mathbb{N}_{\geq 2}$.

Definition 4.3 (Convex and Convexlike Functions). *Let $\varphi : X \rightarrow (-\infty, +\infty]$.*

1. *We say that $\varphi : X \rightarrow (-\infty, +\infty]$ is κ -convex if for all $\alpha \in \Delta_2$*

$$\varphi(\alpha_1 x_1 + \alpha_2 x_2) \leq \alpha_1 \varphi(x_1) + \alpha_2 \varphi(x_2) - \frac{\kappa}{2} \alpha_1 \alpha_2 d_X^2(x_1, x_2).$$

2. *We say that φ is κ -convexlike with respect to a system $\mathcal{B}_2$ of 2-barycentric maps if for every $\mathbf{x} \in \mathcal{B}_2$ and $\alpha \in \Delta_2$,*

$$\varphi(\mathbf{x}(\alpha)) \leq \alpha_1 \varphi(\mathbf{x}_1) + \alpha_2 \varphi(\mathbf{x}_2) - \frac{\kappa}{2} \alpha_1 \alpha_2 d_X^2(\mathbf{x}_1, \mathbf{x}_2). \quad (4.16)$$

3. *We say that φ is **barycentrically κ -convexlike** with respect to a system $\mathcal{B}$ of barycentric maps if for every integer $J \in \mathbb{N}_{\geq 2}$, every $\mathbf{x} \in \mathcal{B}$, and every $\alpha \in \Delta_J$*

$$\varphi(\mathbf{x}(\alpha)) \leq \sum_{j=1}^J \alpha_j \varphi(\mathbf{x}_j) - \frac{\kappa}{4} \sum_{j,k=1}^J \alpha_j \alpha_k d_X^2(\mathbf{x}_j, \mathbf{x}_k) \quad (4.17)$$

4. *We say that φ is **barycentrically quasi-convexlike** with respect to a system $\mathcal{B}$ of barycentric maps if for every integer $J \in \mathbb{N}_{\geq 2}$, every $\mathbf{x} \in \mathcal{B}$, and every $\alpha \in \Delta_J$*

$$\varphi(\mathbf{x}(\alpha)) \leq \max \left\{ \varphi(\mathbf{x}_j) : j \in \{1, \dots, J\}, \alpha_j > 0 \right\} \quad (4.18)$$

5. A collection $(\varphi_a)_{a \in A}$ of functions is **jointly** κ -convexlike (resp. barycentrically κ convexlike) if every function φ_a is κ -convexlike (resp. barycentrically κ -convexlike) with respect to a common system $\mathcal{B}_2$ of 2-barycentric maps (resp. a common system $\mathcal{B}$ of barycentric maps)

Definition 4.4. A geodesic space (X, d_X) is a nonpositively curved (NPC) space in the sense of Alexandrov if it is a CAT(0) space, which is equivalent to the following convexity-nonnegativity inequality

$$d_X^2(\mathbf{x}(0.5, 0.5), z) \leq \frac{1}{2}d_X^2(x, z) + \frac{1}{2}d_X^2(y, z) - \frac{1}{4}d_X^2(x, y) \quad \text{for all } x, y, z \in X, \quad (\text{CN})$$

where $\mathbf{x}$ is a geodesic connecting x and y . We use the notation (Z, d_Z) to denote a NPC space.

The equivalence of CAT(0) spaces and spaces which satisfy (CN) holds specifically for geodesic spaces. Note that (CN) implies that $z \mapsto d_Z^2(z, \bar{z})$ is 1-convex along geodesics for all $z \in Z$. See [33, 142] for more details.

Definition 4.5. We say that φ is upper hemicontinuous with respect to a system $\mathcal{B}_2$ of 2-barycentric maps if for every $\mathbf{x} \in \mathcal{B}_2$, $\varphi \circ \mathbf{x}$ is upper semicontinuous at e_2^1 , i.e. $\varphi(\mathbf{x}(1, 0)) \geq \limsup_{t \downarrow 0} \varphi(\mathbf{x}(1 - t, t))$.

Definition 4.6. We say that a function $\varphi : X \rightarrow (-\infty, +\infty]$ has complete sublevels if for every $c \in \mathbb{R}$ the corresponding sublevel set $\{x \in X : \varphi(x) \leq c\}$ is complete (possibly empty). In particular φ is lower semicontinuous.

$\psi : X \rightarrow [-\infty, +\infty)$ has complete superlevels if for every $c \in \mathbb{R}$ the corresponding superlevel set $\{x \in X : \psi(x) \geq c\}$ is complete. In particular ψ is upper semicontinuous.

Remark 4.7. If X is complete, then φ has complete sublevels (resp. superlevels) iff φ is lower (resp. upper) semicontinuous.

Lemma 4.8. If $\varphi_\alpha : X \rightarrow (-\infty, +\infty]$, $\alpha \in A$, is a collection of functions with complete sublevels, also $\varphi := \sup_{\alpha \in A} \varphi_\alpha$ has complete sublevels. Similarly, if $\psi_\alpha : X \rightarrow [-\infty, +\infty)$, $\alpha \in A$, is a collection of functions with complete superlevels, then $\psi := \inf_{\alpha \in A} \psi_\alpha$ has complete superlevels.

Proof. For every $c \in \mathbb{R}$ we have

$$\{x \in X : \varphi(x) \leq c\} = \bigcap_{\alpha \in A} \{x \in X : \varphi_\alpha(x) \leq c\},$$

so that all the sublevels are intersection of complete subsets of X and are therefore complete as well. The same follows for ψ with $\bigcup_{\alpha \in A}$ instead of $\bigcap_{\alpha \in A}$. $\square$

If $\varphi : X \rightarrow \overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$, we set

$$\text{Dom}(\varphi) := \{x \in X : \varphi \in \mathbb{R}\} \quad (4.19)$$

and we say that φ is **proper** if $\text{Dom}(\varphi)$ is not empty.

Proposition 4.9. *Let $\varphi : X \rightarrow (-\infty, +\infty]$ be a proper function bounded from below with complete sublevels. If φ is κ -convexlike for some $\kappa > 0$ with respect to some system $\mathcal{B}_2$ of 2-barycentric maps, then*

1. *there exists a unique minimizing point $\bar{x}$ such that $\varphi(\bar{x}) = \inf_{x \in X} \varphi(x)$;*
2. *every infimizing sequence $(x_n)_{n \in \mathbb{N}}$ in X converges to $\bar{x}$.*

Proof. **1.** follows by [9, Lemma 2.4.8].

To show **2.**, Let $\bar{\varphi} := \inf_{x \in X} \varphi(x) \in \mathbb{R}$. Fix any $\varepsilon > 0$. We argue as in the proof of Lemma 2.4.8 [9] using the estimate

$$\varphi(x_0) \leq \bar{\varphi} + \varepsilon, \quad \varphi(x_1) \leq \bar{\varphi} + \varepsilon \quad \Rightarrow \quad d_X^2(x_0, x_1) \leq \frac{8\varepsilon}{\kappa}, \quad \text{for all } x_0, x_1 \in X, \quad (4.20)$$

which holds due to the convexity inequality, for any $\mathbf{x} \in \mathcal{B}$ such that $x_0 = \mathbf{x}(e_2^1)$ and $x_1 = \mathbf{x}(e_2^2)$,

$$\bar{\varphi} \leq \varphi(\mathbf{x}(1/2, 1/2)) \leq \frac{1}{2}\varphi(x_0) + \frac{1}{2}\varphi(x_1) - \frac{\kappa}{8}d_X^2(x_0, x_1)^2 \quad \Rightarrow \quad \frac{\kappa}{8}d_X(x_0, x_1) \leq \varepsilon.$$

(4.20) immediately shows that an infimizing sequence satisfies the Cauchy condition and therefore converges to a point $\bar{x}$, which is a minimizing point for φ thanks to its lower semicontinuity. Following Lemma 4.8, (4.20) also shows that minimizing points are unique. $\square$

4.2.2 Metric Derivatives

Definition 4.10. *Let $u : (a, b) \rightarrow X$. The upper/lower right/left metric derivatives and the metric derivative of u are defined by*

$$\begin{aligned} |\dot{u}^+|(t) &:= \limsup_{h \downarrow 0} \frac{d_X(u(t+h), u(t))}{h}, & |\dot{u}_+|(t) &:= \liminf_{h \downarrow 0} \frac{d_X(u(t+h), u(t))}{h}, \\ |\dot{u}^-|(t) &:= \liminf_{h \downarrow 0} \frac{d_X(u(t), u(t-h))}{h}, & |\dot{u}_-|(t) &:= \limsup_{h \downarrow 0} \frac{d_X(u(t), u(t-h))}{h}, \\ |\dot{u}|(t) &:= \lim_{h \rightarrow 0} \frac{d_X(u(t+h), u(t))}{|h|}. \end{aligned} \quad (4.21)$$

A curve u is absolutely continuous if there exists a positive function $m \in L^1(a, b)$ such that

$$d_X(u(t), u(s)) \leq \int_s^t m(r) dr \quad \text{for every } a \leq s \leq t \leq b. \quad (4.22)$$

Let $\mathcal{L}^1$ be the Lebesgue measure on $\mathbb{R}$. An absolutely continuous function $u : (a, b) \rightarrow X$ has metric derivative $\mathcal{L}^1$ -a.e. in (a, b) and $|\dot{u}|(t) \leq m(t)$ for $\mathcal{L}^1$ -a.e. $t \in (a, b)$.

4.2.3 Inequalities with upper and lower Dini derivatives

Definition 4.11. For a real function $\zeta : (a, b) \rightarrow \mathbb{R}$ and $t \in (a, b)$ we consider the upper and lower right/left Dini derivatives

$$\begin{aligned} \frac{d^+}{dt} \zeta(t) &:= \limsup_{h \downarrow 0} \frac{\zeta(t+h) - \zeta(t)}{h}, & \frac{d}{dt_+} \zeta(t) &:= \liminf_{h \downarrow 0} \frac{\zeta(t+h) - \zeta(t)}{h} \\ \frac{d^-}{dt} \zeta(t) &:= \limsup_{h \downarrow 0} \frac{\zeta(t) - \zeta(t-h)}{h}, & \frac{d}{dt_-} \zeta(t) &:= \liminf_{h \downarrow 0} \frac{\zeta(t) - \zeta(t-h)}{h} \end{aligned} \quad (4.23)$$

Lemma 4.12. Let $\zeta : (a, b) \rightarrow [0, \infty)$ be continuous and let $v : (a, b) \rightarrow \mathbb{R}$ be a function satisfying

1. $v(t) = 0$ in $Z := \{t \in (a, b) : \zeta(t) = 0\}$,
2. v is upper semicontinuous in $(a, b) \setminus Z$,
3. v is locally bounded from above in (a, b) .

If

$$\frac{d}{dt_-} \zeta(t) \leq v(t) \quad \text{for every } t \in (a, b) \setminus Z, \quad (4.24)$$

then v is locally integrable in (a, b) and for every $a_0 \in (a, b)$ the function

$$\tilde{\zeta}(t) := \zeta(t) - \int_{a_0}^t v(s) ds \quad \text{is nonincreasing and continuous.} \quad (4.25)$$

In particular (4.24) is equivalent to the distributional inequality

$$\frac{d}{dt} \zeta \leq v \quad \text{in } \mathcal{D}'(a, b), \quad (4.26)$$

$\mathcal{D}'(a, b)$ is the space of distributions dual to $C_c^\infty((a, b), \mathbb{R})$.

Finally, if ζ is locally absolutely continuous (in particular if ζ is locally Lipschitz) then the above conditions are equivalent to

$$\frac{d}{dt} \zeta(t) \leq v(t) \quad \text{for a.e. } t \in (a, b) \quad (4.27)$$

Proof. The proof can be obtained as an adaptation of [171, Lemma A.1]. □

4.2.4 Examples

Examples 4.13. 1. In a geodesic metric space X there is a natural system $\mathcal{G}_2$ of 2-barycentric maps given by all the minimal, constant speed geodesic paths $\mathbf{x} : [0, 1] \rightarrow X$ satisfying (4.13). A function φ is geodesically κ -convex if it is κ -convexlike with respect to a subsystem $\mathcal{B}_2 \subset \mathcal{G}_2$.

2. In a vector space V we can use the canonical convex combinations $\mathbf{x}(\vec{\alpha}) = \sum_{j=1}^J \alpha_j v_j$ parametrized by the points $\vec{\alpha} = (v_1, \dots, v_J) \in V^J$. Notice that the family of shifted norms $\varphi_w(x) := \|x - w\|_V$, $w \in V$, is jointly convex. If the corresponding squared family $\psi_z(x) := \frac{1}{2}\|x - z\|_V^2$ with $z \in V$ is jointly 1-convex then V is a Hilbert space. See [129] for how the parallelogram law ensures that a Banach space is induced by an inner product, $\langle x, y \rangle = \frac{1}{4}(\|x + y\|^2 - \|x - y\|^2)$.

3. If (Z, d_Z) is a NPC space then for every $z \in Z^J$ we can define

$$\mathbf{x}(z; \alpha) := \arg \min_{z \in Z} \sum_{j=1}^J \alpha_j d_Z^2(z, z_j);$$

then the unique minimizer defines a J -barycentric map and the collections of all these maps is a barycentric system. The family $z \mapsto \frac{1}{2}d_Z^2(z, \bar{z})$, $\bar{z} \in Z$, is jointly 1-convex for every choice of $\bar{z} \in Z$.

4. Let $\mathcal{X} = \mathcal{P}_p(X)$, $p \geq 1$, X being a complete and separable metric space. We endow $\mathcal{X}$ with the p -Wasserstein metric W_p . As for the first example, we can consider the collection $\mathcal{B}_J$ of maps parametrized by $\mu = (\mu_1, \dots, \mu_J) \in \mathcal{X}^J$: $\mathbf{x}(\mu; \alpha) = \sum_{j=1}^J \alpha_j \mu_j$. The family $\mu \mapsto W_p(\mu, \nu)$, $\nu \in \mathcal{X}$, is jointly convex.

5. Let us suppose that X is a separable Banach space, with $\mathcal{X} = \mathcal{P}_p(X)$, we can consider a reference standard Borel probability space $(\Omega, \mathcal{B}, \mathcal{P})$ with an atomless Borel probability measure $\mathcal{P}$. We introduce the collection $\mathcal{B}_J$ of maps parametrized by $\mathbf{X} = (X_1, \dots, X_J) \in (L^p(\Omega, \mathcal{P}, \mathcal{X}))^J$: $\mathbf{x}(\mathbf{X}; \alpha) = \left(\sum_{j=1}^J \alpha_j X_j \right)_{\#} \mathcal{P}$. A function which is jointly convex with respect to $\mathcal{B}$ is also called **totally convex**.

6. Let us now consider $\mathcal{X} = \mathcal{P}_2(H)$, H a Hilbert space. For a given reference measure $\bar{\mu} \in \mathcal{X}$ we consider the sets

$$\mathcal{O}_{\bar{\mu}, J} := \left\{ (\bar{X}, X_1, X_2, \dots, X_J) \in (L^p(\Omega, \mathcal{P}, \mathcal{X}))^{J+1} : \bar{X}_{\#} \mathcal{P} = \bar{\mu}, (\bar{X}, X_j)_{\#} \mathcal{P} \text{ optimal in } \mathcal{X} \right\}.$$

For every $\mathbf{X} \in \mathcal{O}_{\bar{\mu}, J}$ we set $\mathbf{x}(\mathbf{X}; \alpha) := \left(\sum_{j=1}^J \alpha_j X_j \right)_{\#} \mathcal{P}$. The collection $\mathcal{B}_{\bar{\mu}}$ (depending on the reference measure $\bar{\mu}$) of these maps provides a natural extension to the notion of

generalized geodesics. Notice that $\mu \mapsto \frac{1}{2}\mathcal{W}_2^2(\mu, \bar{\mu})$ is 1-convex with respect to $\mathcal{B}_{\bar{\mu}}$ and the examples considered in [9, Sec. 9.2] are also barycentrically convexlike with respect to $\mathcal{B}_{\bar{\mu}}$ for every choice of $\bar{\mu}$.

7. Consider the product space $X = \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d$. Such a space was used for the system describing the dynamics of strategic agents interacting with algorithms, studied in [62]. We can define a different interpolation for the component of $\mathcal{P}_2(\mathbb{R}^d)$ than that in $\mathbb{R}^d$, resulting many different interpolation options. We can select the barycentric maps defined in 6. for $\mathcal{P}_2(\mathbb{R}^d)$ above with reference measure $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ and the barycentric maps defined in 1. for $\mathbb{R}^d$,

$$\mathbf{x}(\alpha) := \left(\sum_{j=1}^J \alpha_j \mathbf{x}_j^{(1)} \right)_{\#} \mathcal{P} \times \sum_{j=1}^J \alpha_j \mathbf{x}_j^{(2)}, \quad (\mathbf{x}_j, \rho)_{\#} \mathcal{P} \text{ optimal in } W_2.$$

Then $x \mapsto \frac{1}{2}W_2^2(x^{(1)}, \rho) + \frac{1}{2}\|x^{(2)}\|^2$ is 1-convex with respect to $\mathcal{B}$ parameterized by $\mathbf{x}$ above. Alternatively, we could select 4. for $\mathcal{P}_2(\mathbb{R}^d)$ and 3. for $\mathbb{R}^d$,

$$\mathbf{x}(\alpha) := \sum_{j=1}^J \alpha_j \mathbf{x}_j^{(1)} \times \arg \min_{x^{(2)} \in \mathbb{R}^d} \sum_{j=1}^J \alpha_j \left\| \mathbf{x}_j^{(2)} - x_j^{(2)} \right\|_V^2,$$

where $\|\cdot\|_V$ is any Euclidean norm. However, while this interpolant results in 1-convexity for the family $x^{(2)} \mapsto \frac{1}{2}\|x^{(2)} - \bar{x}^{(2)}\|^2$ for any $\bar{x}^{(2)} \in \mathbb{R}^d$, the family $x^{(1)} \mapsto \mathcal{W}_2^2(x^{(1)}, \rho)$ for any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ is not 1-convex, which precludes the conditions required for taking $\tau \rightarrow 0$ in the movement scheme.

Given a barycentric system $\mathcal{B}$, the convexity of $\alpha \mapsto d_X^2(\mathbf{x}(\alpha), \bar{x})$ for any $\mathbf{x} \in \mathcal{B}$ will be critical for taking the time step to zero; since the choice of $\mathcal{B}$ determines the notion of convexity, selection of the appropriate set of interpolants is important. Additionally, the selection of $\mathcal{B}$ changes the convexity parameter of energy functionals. This is illustrated in the following example.

Example 4.14 (Selection of $\mathcal{B}$ and Energies). Consider the setting $X = \mathcal{P}_2(\mathbb{R}^d)$ endowed with $\mathcal{W}_2$ and energies $E_1(\rho), E_2(\rho) : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{+\infty\}$ defined by

$$E_1(\rho) = \int V(z) d\rho(z), \quad E_2(\rho) = \begin{cases} \int \rho(z) \log \rho(dz) & \text{if } \rho \in L^1(\mathbb{R}^d), \\ +\infty & \text{otherwise} \end{cases},$$

with $V \in C^2(\mathbb{R}^d, \mathbb{R})$ satisfying $\nabla^2 V \succeq \kappa I$. We set $\mathbf{x}$ according to Wasserstein-2 geodesics $\mathcal{G}_2$, and $\mathbf{y}$ according to linear interpolation. Checking the convexity conditions with respect to each of these

barysystems,

$$\begin{aligned}
 E_1(\mathbf{x}(\alpha)) &\leq \alpha_1 E_1(x_1) + \alpha_2 E_1(x_2) - \frac{\kappa}{2} \alpha_1 \alpha_2 W_2^2(x_1, x_2) \\
 E_1(\mathbf{y}(\alpha)) &= \alpha_1 E_1(y_1) + \alpha_2 E_1(y_2) \\
 E_2(\mathbf{x}(\alpha)) &\leq \alpha_1 E_2(x_1) + \alpha_2 E_2(x_2) \\
 E_2(\mathbf{y}(\alpha)) &\leq \alpha_1 E_2(y_1) + \alpha_2 E_2(y_2).
 \end{aligned}$$

If $\kappa > 0$, then for E_1 $\mathbf{x}$ has a "better" (more positive) barycentric convexity parameter; if $\kappa < 0$, then since $\mathbf{y}$ results in 0-convexity, it is a better choice. E_1 is barycentrically quasi-convexlike with respect to both $\mathbf{x}$ and $\mathbf{y}$. For E_2 , geodesics and linear interpolations both result in 0-barycentric convexity.

4.2.4.1 Example (I): Convexity for monotone operators

Let $X = H$ be a Hilbert space and D be a convex subset of H . Consider a bifunction $b : D \times D \rightarrow \mathbb{R}$ of the form $b(x, y) := \langle A(x), y - x \rangle$, where $A : D \rightarrow H$ is an operator. Then $y \mapsto b(y, x) = \langle A(x), y - x \rangle$ barycentrically 0-convexlike for the canonical choice of $\mathcal{B}$ given by Example 4.13(2).

4.2.4.2 Example (II): Convexity for multispecies coupled gradient flows

In the coupled gradient flow setting, any κ -barycentrically convexlike $y \mapsto b(y, x)$ implies that $F^{(i)}(\cdot, \rho^{(-i)})$ is κ -barycentrically convexlike. For any $\rho, \mu \in \mathcal{P}_2(\mathbb{R}^d)^N$ the κ -convexity condition using the definition of $b(\rho, \mu)$ is, for every integer $J \in \mathbb{N}_{\geq 2}$

$$\begin{aligned}
 &\sum_{i=1}^N F^{(i)}(\rho^{(i)}(\alpha), \rho^{(-i)}(\alpha)) - F^{(i)}(\mu^{(i)}, \rho^{(-i)}(\alpha)) \\
 &\leq \sum_{j=1}^J \sum_{i=1}^N \alpha_j \left(F^{(i)}(\rho_j^{(i)}, \rho_j^{(-i)}) - F^{(i)}(\mu^{(i)}, \rho_j^{(-i)}) \right) \\
 &- \frac{\kappa}{4} \sum_{j,k=1}^J \sum_{i=1}^N \alpha_j \alpha_k d_{X^{(i)}}^2(\rho_k^{(i)}, \rho_j^{(i)}), \quad \text{for every } \rho \in \mathcal{B}, \alpha \in \Delta_J.
 \end{aligned}$$

To extract the convexity of $F^{(i)}(\cdot, \mu^{(-i)})$, we select $\rho_j^{(-i)} = \rho_k^{(-i)} =: \mu^{(-i)}$ for all $j, k \in [1, \dots, J]$. Then, the convexity condition becomes

$$F^{(i)}(\rho^{(i)}(\alpha), \rho^{(-i)}) \leq \sum_{j=1}^J \alpha_j F^{(i)}(\rho_j^{(i)}, \rho^{(-i)}) - \frac{\kappa}{4} \sum_{k,j=1}^J \alpha_k \alpha_j d_{X^{(i)}}^2(\rho_k^{(i)}, \rho_j^{(i)}), \quad (4.28)$$

which shows that $F^{(i)}(\cdot, \rho^{(-i)})$ is κ -barycentrically convex for all $i \in [1, \dots, N]$. In the Wasserstein-2 space, two different choices of $J = 2$ barysystems result in classical types of convexity. Let $\Gamma^*(\mu, \nu)$

be the set of Wasserstein-2 optimal transport maps with marginals μ and ν . If we select $\mathcal{B} = \mathcal{G}_2$ to be Wasserstein-2 geodesics, then the interpolant $\rho(\alpha)$ is given by

$$\rho^{(i)}(\alpha) = [\alpha_1 z_1 + \alpha_2 z_2]_{\#} \gamma^{(i)}, \quad \text{where } \gamma^{(i)} \in \Gamma^*(\rho_1^{(i)}, \rho_2^{(i)}) \quad \text{and } 1 - \alpha_1 = \alpha_2.$$

Then the convexity inequality (4.28) is

$$F^{(i)}(\rho^{(i)}(\alpha), \rho^{(-i)}) \leq \alpha_1 F^{(i)}(\rho_1^{(i)}, \rho^{(-i)}) + \alpha_2 F^{(i)}(\rho_2^{(i)}, \rho^{(-i)}) - \frac{\kappa}{2} \alpha_1 \alpha_2 \mathcal{W}_2^2(\rho_1^{(i)}, \rho_2^{(i)}),$$

which is the condition for κ displacement convexity of $F^{(i)}(\cdot, \rho^{(-i)})$ in the sense of McCann [163].

Later, we will see in Example 4.5.1.3 that one must consider convexity with respect to generalized geodesics, since the Wasserstein-2 distance function is not geodesically 1-convex. In this case, the barysystem will be dependent on a reference measure $\bar{\rho}$:

$$\begin{aligned} \rho^{(i)}(\alpha) &= [\alpha_1 z_1 + \alpha_2 z_3]_{\#} \hat{\gamma}^{(i)}, \quad \alpha_1 = 1 - \alpha_2, \\ \text{where } \hat{\gamma} &\in \mathcal{P}_2((\mathbb{R}^d)^3) \text{ satisfies } [z_1, z_2]_{\#} \hat{\gamma}^{(i)} \in \Gamma^*(\rho_1^{(i)}, \bar{\rho}^{(i)}), \quad [z_2, z_3]_{\#} \hat{\gamma} \in \Gamma^*(\bar{\rho}^{(i)}, \rho_2^{(i)}). \end{aligned}$$

Convexity along generalized geodesics implies geodesic convexity [9, Remark 9.2.8]; the reverse implication does not necessarily hold. It does hold, however, for common choices such as potential, self-interaction, and internal energy functionals; see [9, Section 9.3] for details.

4.3 Equilibrium problems

We now present the equilibrium problems that are the key tools for ensuring existence of solutions to the discrete b-EVI (4.11), strongly inspired by the exposition of [14].

A first approach to existence of solutions to the so-called primal formulation relies on the notion of barycentric quasi-convexity and on compactness. If the bifunction also satisfies an interaction dissipativity condition and a stronger barycentric convexity, we will also show how to obtain existence without any compactness assumption by employing a dual formulation.

We state our results for a general bifunction $B : D \times D \rightarrow \mathbb{R}$ defined in a subset D of a metric space (X, d_X) . Even if they are completely general, we use the symbol B since their main applications will concern the generating bifunction introduced by (4.9).

We will always assume that B is nonnegative along the diagonal:

$$B(x, x) \geq 0 \quad \text{for every } x \in D. \quad (4.29)$$

Let us first study a one-sided equilibrium problem induced by B .

Problem 4.15. *We say that $x^* \in D$ is a solution of the Primal Equilibrium Problem if*

$$B(y, x^*) \geq 0 \quad \text{for every } y \in D. \quad (\text{EP})$$

Theorem 4.16. *Let us suppose $B : D \times D \rightarrow \mathbb{R}$ satisfies (4.29) and that*

1. *the functions $y \mapsto B(y, x)$, $x \in D$, are jointly barycentrically quasi-convexlike with respect to a system of **continuous** barycentric maps $\mathcal{B}$ in D .*
2. *the superlevels $\sigma^+(y) := \{x \in D : B(y, x) \geq 0\}$ are closed in D and at least one of them is compact.*

Then Problem 4.15 has a solution.

Proof. It is sufficient to check that the collection of sets $\sigma^+(y)$, $y \in D$, have the finite intersection property. Let $y_1, \dots, y_J \in D$, let $\mathbf{y} \in \mathcal{B}_J$ such that $y_k = \mathbf{y}(e_j^k)$, and consider the sets

$$\mathcal{S}(k) \subset \Delta_J, \quad \mathcal{S}(k) := \left\{ \beta \in \Delta_J : B(y_k, \mathbf{y}(\beta)) \geq 0 \right\}.$$

Since $\mathbf{y}$ is continuous and the superlevels $\sigma^+(y_k)$ are closed, every $\mathcal{S}(k)$ is closed. Let us check that $\{\mathcal{S}(k)\}_{k=1}^J$ satisfies the assumption of the KKM lemma (see Lemma 4.17 below): if $B = \{k_1, \dots, k_m\}$ is a subset of $\{1, \dots, J\}$ we need to prove that $\text{co}(e_j^{k_1}, \dots, e_j^{k_m}) \subset \bigcup_{h=1}^m \mathcal{S}(k_h)$. By contradiction, if $\Delta_J \ni \alpha = \sum_{k=1}^J \alpha_k e_j^k \notin \bigcup_{h=1}^m \mathcal{S}(k_h)$, for some coefficients α_j with $\alpha_j = 0$ if $j \notin B$, then $B(y_{k_h}, \mathbf{y}(\alpha)) < 0$ for every $h = 1, \dots, m$ and therefore, since $y \mapsto B(y, \mathbf{y}(\alpha))$ is quasi-convexlike with respect to $\mathcal{B}_J$, $B(\mathbf{y}(\alpha), \mathbf{y}(\alpha)) \leq \max_{k=1, \dots, m} B(y_{k_h}, \mathbf{y}(\alpha)) < 0$, a contradiction with (4.29). Hence, $\{\mathcal{S}(k)\}_{k=1}^J$ is a KKM covering and there exists $\beta^* \in \bigcap_{k=1}^J \mathcal{S}(k)$. Let $x^* = \mathbf{y}(\beta^*)$. Then $B(y_k, x^*) \geq 0$ for all $k \in [1, \dots, J]$, i.e. $x^* \in \bigcap_{k=1}^J \sigma^+(y_k)$. $\square$

Lemma 4.17 (KKM Lemma [134]). *Let Δ_J be the $J - 1$ -dimensional simplex with J vertices. Let $C_1, \dots, C_J \subset \Delta_J$ be closed subsets such that every $\alpha \in \Delta_J$ satisfies $\alpha \in \bigcup_{j=1}^J C_j$. Then $\bigcap_{j=1}^J C_j \neq \emptyset$.*

We want to study now the relation between Problem 4.15 and its dual formulation.

Problem 4.18. *We say that y_* is a solution of the Dual Equilibrium Problem if*

$$B(y_*, x) \leq 0 \quad \text{for every } x \in D. \quad (\text{EP}')$$

The following metric version of Minty's trick shows that solutions to Problem 4.18 are also solution to Problem 4.15.

Lemma 4.19 (A metric Minty's lemma). *Let us suppose that*

1. *the functions $y \mapsto B(y, x)$, $x \in D$, are jointly quasi-convexlike with respect to a continuous system of 2-barycentric maps $\mathcal{B}_2$ in D .*

2. for every $y \in D$ the map $x \mapsto B(y, x)$ is upper-hemicontinuous with respect to $\mathcal{B}_2$.

Then every solution of Problem 4.18 is also a solution of Problem 4.15.

Proof. Let y_* be a solution to Problem 4.18, let $y \in D$ and let $\mathbf{x} \in \mathcal{B}_2$ satisfying $\mathbf{x}_1 = y_*$, $\mathbf{x}_2 = y$. We set $x_t := \mathbf{x}((1 - t), t)$. Since y_* is a solution to Problem 4.18 we have

$$B(y_*, x_t) \leq 0.$$

On the other hand (4.29) yields

$$B(x_t, x_t) \geq 0.$$

Since $s \mapsto B(x_s, x_t)$ satisfies $B(x_s, x_t) \leq \max(B(y_*, x_t), B(y, x_t))$ and using $s = t$ we conclude that $B(y, x_t) \geq 0$. Passing to the limit as $t \downarrow 0$ and using the upper-semicontinuity of $t \mapsto B(y, x_t)$ at $t = 0$ we obtain $B(y, y_*) \geq 0$. $\square$

Remark 4.20. If B satisfies the dissipativity condition

$$B(x, y) + B(y, x) \leq 0 \quad \text{for every } x, y \in D \quad (4.30)$$

then clearly every solution to Problem 4.15 is also a solution to Problem 4.18. If in addition the mild regularity assumptions of Lemma 4.19 hold, the two problems are equivalent.

Notice that (4.30) and (4.29) force

$$B(x, x) = 0 \quad \text{for every } x \in D. \quad (4.31)$$

Remark 4.21. $x_* = x^*$ is a joint solution of the two equilibrium problems 4.15 and 4.18 if and only if x_* is a diagonal saddle point of B , i.e.

$$B(y, x_*) \geq 0 = B(x_*, x_*) \geq B(x_*, x) \quad \text{for every } x, y \in D. \quad (\text{Saddle})$$

Let us observe that we are looking for saddle points of B on the diagonal of $D \times D$. Therefore, if we introduce the marginal functions

$$f(x) := \inf_{y \in D} B(y, x), \quad g(y) := \sup_{x \in D} B(y, x) \quad (4.32)$$

which satisfy

$$\max_{x \in D} f(x) = \max_{x \in D} \inf_{y \in D} B(y, x), \quad \min_{y \in D} g(y) = \min_{y \in D} \sup_{x \in D} B(y, x), \quad (4.33)$$

it is immediate to observe that (4.31) yields

$$g(y) \geq 0 \geq f(x) \quad \text{for every } x, y \in D. \quad (4.34)$$

(Saddle) is thus equivalent to

$$g(x_*) = f(x_*). \quad (4.35)$$

We now state our second existence result, when B satisfies a more restrictive dissipativity condition and a uniform barycentric convexity. In this case, we can remove the compactness assumption and relax the upper semicontinuity condition as well. For every $x \in D$ we set

$$\sigma_-(x) := \{y \in D : B(y, x) \leq 0\}. \quad (4.36)$$

Theorem 4.22. *Let $\xi > \eta_+$ and let $B : D \times D \rightarrow \mathbb{R}$ be a function satisfying (4.29) and the following properties:*

1. *η interaction dissipativity:*

$$B(x, y) + B(y, x) \leq \eta d_X^2(y, x) \quad \text{for every } x, y \in D. \quad (4.37)$$

2. *The maps $y \mapsto B(y, x)$, $x \in D$, have complete sublevels $\sigma_-(x)$ and are jointly barycentrically ξ -convexlike with respect to a system of continuous barycentric maps $\mathcal{B}$.*

3. *For some $x_0 \in D$ we have $\inf_{y \in D} B(y, x_0) > -\infty$.*

4. *For every $y \in D$ the map $x \mapsto B(y, x)$ are upper-hemicontinuous with respect to $\mathcal{B}_2$.*

Then Problem 4.18 has a unique solution y_* which is also a solution to Problem 4.15.

Proof. It is not restrictive to assume $\eta \geq 0$.

a) *The sets $\sigma_-(x)$, $x \in D$, have the finite intersection property.*

Let $x_1, x_2, \dots, x_J \in D$ and consider the bilinear function $q(\alpha, \beta) := \sum_{j,k} \alpha_j \beta_k B(x_j, x_k)$ defined in $\Delta_J \times \Delta_J$. When $\alpha = \beta$ we have

$$\begin{aligned} q(\alpha, \alpha) &= \sum_{j,k} \alpha_j \alpha_k B(x_j, x_k) = \frac{1}{2} \left(\sum_{j,k} \alpha_j \alpha_k B(x_j, x_k) + \sum_{j,k} \alpha_k \alpha_j B(x_k, x_j) \right) \\ &= \frac{1}{2} \sum_{j,k} \alpha_j \alpha_k (B(x_j, x_k) + B(x_k, x_j)) \leq \frac{\eta}{2} \sum_{j,k} \alpha_j \alpha_k d_X^2(x_j, x_k). \end{aligned}$$

thanks to (4.37). It follows that

$$\min_{\beta \in \Delta_J} q(\alpha, \beta) \leq q(\alpha, \alpha) \leq \frac{\eta}{2} \sum_{j,k} \alpha_j \alpha_k d_X^2(x_j, x_k) \leq \frac{\eta}{2} Q(\alpha), \quad Q(\alpha) := \max_{\alpha \in \Delta_J} \sum_{j,k} \alpha_j \alpha_k d_X^2(x_j, x_k).$$

Taking a maximum over α ,

$$\max_{\alpha \in \Delta_J} \min_{\beta \in \Delta_J} q(\alpha, \beta) \leq \frac{\eta}{2} Q(\alpha).$$

By the Von Neumann minimax theorem [136], we deduce

$$\min_{\beta \in \Delta_J} \max_{\alpha \in \Delta_J} q(\alpha, \beta) \leq \frac{\eta}{2} Q(x),$$

so that there exists $\beta^* \in \Delta_J$ such that

$$q(\alpha, \beta^*) \leq \frac{\eta}{2} Q(x) \quad \text{for every } \alpha \in \Delta_J.$$

Selecting $\alpha = [1/J, \dots, 1/J]$,

$$\sum_j \beta_j^* B(x_j, x_k) \leq \frac{\eta}{2} Q(x) \quad \text{for every } k = 1, \dots, J. \quad (4.38)$$

Since $y \mapsto B(y, x_k)$ are jointly ξ barycentrically convexlike, selecting $y_* := \mathbf{x}(\beta^*)$,

$$\begin{aligned} B(y_*, x_k) &\leq \sum_{j=1}^J \beta_j^* B(x_j, x_k) - \frac{\xi}{2} \sum_{h,j=1}^J \beta_h^* \beta_j^* d_X^2(x_h, x_j) \\ &\leq \frac{\eta}{2} \left(Q(x) - \sum_{h,j=1}^J \beta_h^* \beta_j^* d_X^2(x_h, x_j) \right) \leq 0 \quad \text{for every } k = 1, \dots, J, \end{aligned}$$

using the estimate (4.38). Therefore $y_* \in \bigcap_{k=1}^J \sigma_-(x_k)$.

b) $\bigcap_{x \in D} \sigma_-(x)$ is nonempty and contains a unique element.

For every $A \subset D$ with $A \ni x_0$, where $\inf_y (y, x_0) > -\infty$ from the third assumption, we set

$$B_A(y) := \sup_{x \in A} B(y, x), \quad \sigma_-(A) := \left\{ y \in D : B_A(y) \leq 0 \right\}. \quad (4.39)$$

The function B_A is ξ -convexlike and uniformly bounded from below since $B_A \geq B(\cdot, x_0) \geq \inf_y b(y, x_0) > -\infty$, and together with (4.37), we conclude that $B_A(\cdot)$ is proper. Using the assumption that $B(\cdot, x)$ has complete sublevel set, B_A has a unique minimizer y_A by Proposition 4.9. In the particular case when $A = D$ we recover the marginal function g of (4.39) $g(y) = B_D(y) = \sup_{x \in D} B(y, x) \geq B(x, x) = 0$ therefore $\sigma_-(D)$ is nonempty if $\min B_D = 0$.

Let $\Lambda = \{A \subset D : x_0 \in A\}$, restricted such that $A, B \in \Lambda$ implies $A \subseteq B$ or $B \subseteq A$, be the collection of subsets with a finite number of elements of D containing x_0 and ordered by inclusion. Λ is a directed set. Every set $\sigma_-(A)$, $A \in \Lambda$, is nonempty by Claim a) because each $A \in \Lambda$ is finite and it is included in $\sigma_-(x_0)$, which is a complete set. We have $m_A := B_A(y_A) \leq 0$ by Claim a); in particular $y_A \in \sigma_-(A)$ for every $A \in \Lambda$.

Claim b) then follows if we show that $A \mapsto y_A$ is a Cauchy net. We know that $A \mapsto m_A$ is bounded from above and increasing: if $A_1 \leq A_2$,

$$m_{A_1} = \min_{y \in D} B_{A_1}(y) \leq \min_{y \in D} B_{A_2}(y) = m_{A_2}.$$

From this, $\lim_{A \in \Lambda} m_A = \sup_{A \in \Lambda} m_A =: m \leq 0$, where $\liminf$ is taken in the sense of nets. Given $\varepsilon > 0$ we can find $A^\varepsilon \in \Lambda$ such that $m_{A^\varepsilon} \geq m - \varepsilon$ and therefore for every A_1, A_2 containing A^ε with $y_k := y_{A_k}$ we have

$$m_{A^\varepsilon} \leq B_{A^\varepsilon}(y^{(k)}) \leq B_{A_k}(y^{(k)}) = m_{A_k} \leq m \leq m_{A^\varepsilon} + \varepsilon.$$

We now use that $y \mapsto B(y, x)$ is ξ -barycentrically convexlike, selecting $J = 2$ and test points y_1, y_2 :

$$B(\mathbf{y}(0.5, 0.5), x) \leq \frac{1}{2}B(y_1, x) + \frac{1}{2}B(y_2, x) - \frac{\kappa}{8}d_X^2(y_1, y_2).$$

Using that for all $x \in A^\varepsilon$, $B(y_1, x) \leq m_{A^\varepsilon} + \varepsilon$ and $B(y_2, x) \leq m_{A^\varepsilon} + \varepsilon$, along with $B(\mathbf{y}(0.5, 0.5), x) \geq B_{A^\varepsilon}(y_{A^\varepsilon})$, we have

$$d_X^2(y_1, y_2) \leq \frac{8\varepsilon}{\xi}.$$

The above estimate in combination with $y_A \in \sigma_-(x_0)$, a complete set, allows us to conclude that $A \mapsto y_A$ is a Cauchy net converging to a limit $\bar{y} \in \sigma_-(x_0)$. Since $A \cup \{x\} \mapsto y_{A \cup \{x\}}$ is a subnet of $A \mapsto y_A$ for any $x \in D$, if $A \mapsto y_A$ converges to $\bar{y}$, then $A \cup \{x\} \mapsto y_{A \cup \{x\}}$ also converges to $\bar{y}$. Therefore, for every $x \in D$, $B(\bar{y}, x) \leq \liminf_{A \in \Lambda} B(y_A, x) \leq \liminf_{A \in \Lambda} B_A(y_A) = m \leq 0$ so that $\bar{y} \in \bigcap_{x \in D} \sigma_-(x)$. Note that lower semicontinuity of $B(\cdot, x)$ follows from the assumption that $\sigma_-(x)$ are complete. $\square$

4.3.1 Saddle points of antysymmetric bifunctions

The antisymmetric case when

$$B(x, y) + B(y, x) = 0 \quad \text{for every } x, y \in D, \quad (4.40)$$

deserves a specific analysis, since we can considerably simplify the convexity assumptions on B . It is clear that in this case Problems 4.15 and 4.18 are equivalent and any (common) solution is a saddle point according to Remark 4.21.

Theorem 4.23. *Let $B : D \times D \rightarrow \mathbb{R}$ be a function satisfying the antisymmetric identity (4.40). If*

1. *The maps $y \mapsto B(y, x)$, $x \in D$, have complete sublevels and are jointly ξ -barycentrically convexlike with respect to some family $\mathcal{B}$ for some $\xi > 0$.*
2. *For some $x_0 \in D$ we have $-I := \inf_{y \in D} B(y, x_0) > -\infty$.*
3. *For every $y \in D$ the maps $x \mapsto B(y, x)$ are upper-hemicontinuous with respect to $\mathcal{B}_2$.*

then Problems 4.15 and 4.18 have a unique saddle solution $x^* = y_*$.

Proof. We adapt the argument of [32, Chapter 1, Theorem 1.1] and [98, Chap. 6, Theorem 2.1] due to Shiffman. Let us define $g(y) = \sup_{x \in D} B(y, x)$ as in (4.32). Notice that (4.40) yields $f(x) = -g(x)$. We observe that the function $y \mapsto g(y)$ is ξ -convexlike, it has complete sublevels, and it is bounded from below since $g(y) \geq B(y, y) = 0$. Moreover

$$B(x_0, x) = -B(x, x_0) \leq I \quad \text{for every } x \in D,$$

so that $g(x_0) \leq I < +\infty$ and therefore g is not identically $+\infty$ it attains its minimum at some point $\bar{y} \in D$. We want to show that $g(\bar{y}) = 0$. Let us fix $y \in D$ and let $\mathbf{y} \in \mathcal{B}_2$ satisfy $\mathbf{y}_1 = \bar{y}$, $\mathbf{y}_2 = y$; for every $\vartheta \in (0, 1)$ we set $y_\vartheta = \mathbf{y}((1 - \vartheta), \vartheta)$. Since $g(y_\vartheta) = \sup_x B(y_\vartheta, x)$, and $x \mapsto -B(y, x) = B(x, y)$ is ξ -convexlike with respect to $\mathcal{B}_2$, there exists a unique point x_ϑ such that $g(y_\vartheta) = B(y_\vartheta, x_\vartheta)$. We have

$$g(\bar{y}) \leq g(y_\vartheta) = B(y_\vartheta, x_\vartheta) \stackrel{(1)}{\leq} (1 - \vartheta)B(\bar{y}, x_\vartheta) + \vartheta B(y, x_\vartheta) \stackrel{(2)}{\leq} (1 - \vartheta)g(\bar{y}) + \vartheta B(y, x_\vartheta)$$

where we used the convexity inequality in (1) and the minimality property of g in (2) ($g(\bar{y}) \geq b(\bar{y}, x)$ for every $x \in D$). We deduce that

$$g(\bar{y}) \leq B(y, x_\vartheta) \quad \text{for every } \vartheta \in (0, 1). \quad (4.41)$$

As $\vartheta \downarrow 0$ y_ϑ is a minimizing sequence for g since $g(y_\vartheta) \leq (1 - \vartheta)g(\bar{y}) + \vartheta g(y)$ so that $\liminf_{\vartheta \downarrow 0} g(y_\vartheta) = g(\bar{y})$. We deduce that $y_\vartheta \rightarrow \bar{y}$ as $\vartheta \downarrow 0$.

We claim that x_ϑ converges to the unique maximizer $\bar{x}$ of $x \mapsto B(\bar{y}, x)$ as $\vartheta \downarrow 0$. In fact

$$\begin{aligned} (1 - \vartheta)B(\bar{y}, x_\vartheta) + \vartheta g(y) &\geq (1 - \vartheta)B(\bar{y}, x_\vartheta) + \vartheta B(y, x_\vartheta) \geq B(y_\vartheta, x_\vartheta) = g(y_\vartheta) \\ &\geq B(y_\vartheta, x) \quad \text{for every } x \in D, \end{aligned}$$

so that

$$\limsup_{\vartheta \downarrow 0} B(\bar{y}, x_\vartheta) \geq \limsup_{\vartheta \downarrow 0} B(y_\vartheta, x) \geq B(\bar{y}, x) \quad \text{for every } x \in D,$$

and therefore

$$\limsup_{\vartheta \downarrow 0} B(\bar{y}, x_\vartheta) \geq \sup_{x \in D} B(\bar{y}, x) = g(\bar{y})$$

showing that x_ϑ is a maximizing sequence for the ξ -concavelike function $x \mapsto B(\bar{y}, x)$. Proposition 4.9 then shows that $x_\vartheta \rightarrow \bar{x}$. We can eventually pass to the limit in (4.41) as $\vartheta \downarrow 0$ using the lower semicontinuity of B w.r.t. x obtaining

$$g(\bar{y}) = B(\bar{y}, \bar{x}) \leq B(\bar{x}, y) \quad \text{for every } y \in D,$$

so that $g(\bar{y}) \leq f(\bar{x}) \leq 0$.

To show uniqueness, if (y', y') is another saddle point for B , then

$$g(y') = B(y', y') = 0$$

yields $g(y') = g(\bar{y}) = 0$ showing that $\bar{y} = y'$ thanks to the ξ -convexity of g . $\square$

4.3.2 Examples

In some settings, showing the primal problem properties is easy due to compactness.

Example 4.24 (Reinforcement Learning and Primal Problem). *In reinforcement learning, the cost function often depends on a finite set of state-action pairs. Consider $X = \Delta_N$ with metric $\mathcal{W}_2$, where N is the total number of state-action pairs. A commonly-analyzed cost function $C : X \rightarrow \mathbb{R}$ [161] is*

$$C(\rho) = c^\top \rho + \sum_{i=1}^M H(a_i^\top \rho), \quad H(b) = b \log b,$$

where the vectors a_i index all state-action pairs associated with action i . The summation represents an incentive to play different actions (exploration). A gradient flow on this cost function will have

$$b(y, x) = C(y) - C(x).$$

For this particular choice of B , we can show that the assumptions in Theorem 4.23 for the primal problem 4.15 hold. To show that $y \mapsto b(y, x)$ is quasi-convexlike with respect to $\mathcal{B}$, which is defined by $\mathbf{y}(\alpha) := \sum_{i=1}^J \alpha_i \mathbf{y}_i$,

$$\begin{aligned} b(\mathbf{y}(\alpha), x) &= c^\top \sum_{i=1}^J \alpha_i \mathbf{y}_i + \sum_{k=1}^M \left(a_k^\top \sum_{j=1}^J \alpha_j \mathbf{y}_j \right) \log \left(a_k^\top \sum_{j=1}^J \alpha_j \mathbf{y}_j \right) - C(x) \\ &\leq \sum_{j=1}^J \alpha_j \left(c^\top \mathbf{y}_j + \sum_{k=1}^M a_k^\top \mathbf{y}_j \log a_k^\top \mathbf{y}_j \right) - C(x) \leq \max_j b(\mathbf{y}_j, x), \end{aligned}$$

where the first inequality holds by Jensen's inequality. Finally, all superlevels $\sigma^+(y) = \{x \in D : b(y, x) \geq 0\}$ are closed in D and compact in $\mathcal{W}_2$ because in finite dimensions, closedness and boundedness implies compactness.

4.3.2.1 Coupled Gradient Flow Setting: Monotonicity and Dissipativity

In the setting where b specifies a coupled n -species gradient flow system, where each species aims to minimize $F^{(i)} : X \rightarrow \mathbb{R}$, b takes the form

$$b(y, x) = \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}). \quad (4.42)$$

In the game theory literature, strong monotonicity of a game is a sufficient condition for existence of a Nash equilibrium. In Euclidean space, a game $(F^{(i)})$ λ -monotone if

$$\sum_{i=1}^N \langle \nabla_{x^{(i)}} F^{(i)}(x) - \nabla_{x^{(i)}} F^{(i)}(y), x - y \rangle \geq \lambda \|x - y\|^2 \quad \text{for all } x, y \in \mathbb{R}^d. \quad (4.43)$$

When $\lambda > 0$, it can be shown that coupled gradient descent converges to the unique Nash equilibrium of the game ($F^{(i)}$); however, this definition relies on a notion of gradient of $F^{(i)}$. The η interaction dissipativity and barycentrically κ -convexlike conditions for the dual equilibrium problem point toward a zeroth-order condition for monotonicity. We propose the following zeroth order definition of monotonicity for coupled games in metric spaces.

Definition 4.25 (Game-Theoretic Monotonicity). *A set of energy functionals ($F^{(i)}$) are $-\lambda$ -monotone, where $\lambda = \eta - \kappa$, if*

- κ convexity: $y \mapsto B(y, x)$ is barycentrically κ -convexlike with respect to a system $\mathcal{B}$.
- η dissipativity: $b(y, x) + b(x, y) \leq \eta d_X^2(x, y)$.

Equivalently, each $F^{(i)}(\cdot, \cdot) - F^{(i)}(x^{(i)}, \cdot)$ is barycentrically κ -convexlike and for every $x, y \in D$,

$$\sum_{i=1}^N F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}) + F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) \leq \eta d_X^2(x, y).$$

In Lemma 4.55, we show that Definition 4.25 implies the standard Euclidean game theory monotonicity inequality (4.43) when $\mathcal{B}$ includes linear interpolations, and in Lemma 4.57 we show Definition 4.25 implies the λ -monotonicity as defined for coupled Wasserstein-2 gradient flows in [64] when $\mathcal{B}$ includes geodesics for $J = 2$. In the Wasserstein case, the set of curves needed for first-order monotonicity are the geodesics, $\mathcal{G}_2$, which are a 2-point, rather than J -point, interpolation. For many common functionals such as potential, self-interacting, and entropy functionals, geodesic convexity implies convexity along generalized geodesics [9, Section 4.2.2], which can then be used to show barycentric convexity.

The reverse direction in both Euclidean and Wasserstein-2 spaces does not necessarily hold. This is because the zeroth order condition requires an upper-bound on $F^{(i)}$, whereas the first order requires a lower bound on $\nabla F^{(i)}$, and so the value of λ obtained from the first order condition can be tighter. See Example 4.26 for an example of this gap.

Example 4.26 (Zeroth- and first-order monotonicity discrepancy). *This example illustrates the gap between first order and zeroth order monotonicity. Consider the two-player game in $X^{(1)} = \mathbb{R}^{d_1}$, $X^{(2)} = \mathbb{R}^{d_2}$ with energies*

$$F^{(1)}(z) = \frac{1}{2} z^{(1)\top} Q_1 z^{(1)} + z^{(1)\top} B z^{(2)}, \quad F^{(2)}(z) = \frac{1}{2} z^{(2)\top} Q_2 z^{(2)}$$

The convexity parameter is $\kappa = \lambda_{\underline{Q}}$, where $\lambda_{\underline{Q}}$ is the smallest eigenvalue among the eigenvalues of Q_1 and Q_2 . The first-order monotonicity condition depends on the smallest eigenvalue of M ,

denoted $\lambda_{\underline{M}}$, where

$$\sum_{i=1}^2 \langle \nabla_{z^{(i)}} F^{(i)}(z) - \nabla_{z^{(i)}} F^{(i)}(y), z^{(i)} - y^{(i)} \rangle = (z - y)^\top M (z - y) \geq \lambda_{\underline{M}} \|z - y\|^2,$$

$$M = \begin{bmatrix} Q_1 & \frac{1}{2}B \\ \frac{1}{2}B^\top & Q_2 \end{bmatrix}.$$

The zeroth order condition is

$$\begin{aligned} & \sum_{i=1}^2 F^{(i)}(z^{(i)}, z^{(-i)}) - F^{(i)}(y^{(i)}, z^{(-i)}) + F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(z^{(i)}, y^{(-i)}) \\ &= (z^{(1)} - y^{(1)})^\top B (z^{(2)} - y^{(2)}) \\ &= (z - y)^\top M (z - y) - (z^{(1)} - y^{(1)})^\top Q_1 (z^{(1)} - y^{(1)}) - (z^{(2)} - y^{(2)})^\top Q_2 (z^{(2)} - y^{(2)}) \\ &\leq \lambda_{\overline{M}} \|z - y\|^2 - (z^{(1)} - y^{(1)})^\top Q_1 (z^{(1)} - y^{(1)}) - (z^{(2)} - y^{(2)})^\top Q_2 (z^{(2)} - y^{(2)}) \\ &\leq (\lambda_{\overline{M}} - \lambda_{\underline{Q}}) \|z - y\|^2, \end{aligned}$$

where $\lambda_{\overline{M}}$ is the largest eigenvalue of M . The upper bound η is given by $\eta = \lambda_{\overline{M}} - \lambda_{\underline{Q}}$. According to the first-order monotonicity definition, $\lambda = -\lambda_{\underline{M}}$, whereas according to the zeroth order definition, $\lambda = \eta - \kappa = \lambda_{\overline{M}} - \lambda_{\underline{Q}} - \lambda_{\underline{Q}} = \lambda_{\overline{M}} - 2\lambda_{\underline{Q}}$. If $\lambda_{\overline{M}} - 2\lambda_{\underline{Q}} > -\lambda_{\underline{M}}$, then the zeroth order and first order definitions are not equivalent, with the first-order condition providing a tighter value of η .

For a game with all players optimizing over probability measures, another definition of monotonicity is presented in [206], which the authors call "linear monotonicity," and they call the definition in [64] "displacement monotonicity." Linear monotonicity (with $\lambda = 0$) is defined as

$$\sum_{i=1}^N \int \left(\delta_{\rho^{(i)}} F^{(i)}[\rho](x) - \delta_{\rho^{(i)}} F^{(i)}[\tilde{\rho}](x) \right) d(\rho^{(i)} - \tilde{\rho}^{(i)})(x) \geq 0 \quad \text{for all } \rho^{(i)}, \tilde{\rho}^{(i)} \in \mathcal{P}_2(\mathbb{R}^{d_i}).$$

The notion of convexity of $F^{(i)}(\cdot, \rho^{(i)})$ is linear convexity, such that

$$F^{(i)}(\alpha \rho^{(i)} + (1 - \alpha) \tilde{\rho}^{(i)}, \rho^{(-i)}) \leq \alpha F^{(i)}(\rho^{(i)}, \rho^{(-i)}) + (1 - \alpha) F^{(i)}(\tilde{\rho}^{(i)}, \rho^{(-i)}).$$

For functionals that satisfy linear monotonicity, we can use linear interpolations rather than geodesic interpolations to show existence of a Nash equilibrium.

4.3.2.2 Interaction Dissipativity for Coupled Gradient Flow System

As discussed in Section 4.3.2.1, κ is the convexity parameter for all $F^{(i)}(\cdot, y^{(-i)})$, and η is the dissipativity lost (or gained, if $\eta < 0$) from the coupling structure. While our results are formulated

for $\eta \in \mathbb{R}$ for the case of general b , we can show that in the coupled gradient flow setting, $\eta \geq 0$, that is, the coupling structure cannot improve system dissipativity in comparison with each species evolving separately.

Proposition 4.27. *Let b take the coupled gradient flow form (4.42) and suppose $D := \prod_{i=1}^n D^{(i)}$. Then $\eta \geq 0$.*

Proof. We wish to show that, given the condition

$$b(x, y) + b(y, x) \leq \eta d_X^2(x, y), \quad \forall x, y \in D,$$

it is not possible for $\eta < 0$ when b takes the form (4.42). Let the swap operator $S^{(i)} : D \times D \rightarrow D \times D$, which switches the i^{th} entry in the two input elements, be defined as

$$S^{(i)}(x, y) = (\tilde{x}, \tilde{y}), \quad \tilde{x}^{(j)} = x^{(j)}, \tilde{y}^{(j)} = y^{(j)} \quad \forall j \neq i, \quad \text{and } \tilde{x}^{(i)} = y^{(i)}, \tilde{y}^{(i)} = x^{(i)}.$$

We rewrite $b(y, x) + b(x, y)$ as

$$b(y, x) + b(x, y) = \sum_{i=1}^N G_i(x, y),$$

$$G_i(x, y) := F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) + F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}).$$

Importantly, using (4.42), $G_i(S^{(i)}(x, y)) = -G_i(x, y)$. For the sake of contradiction, suppose $\eta < 0$ so that, for any x, y such that $d_X(x, y) > 0$,

$$\sum_{i=1}^N G_i(x, y) < 0. \tag{4.44}$$

Since (4.44) holds for all $x \neq y$, the point $S^{(i)}(x, y) \in D \times D$ must also satisfy

$$\sum_{i=1}^N G_i(S^{(1)}(x, y)) < 0,$$

which, in combination with (4.44), implies that

$$\sum_{i=2}^N G_i(S^{(1)}(x, y)) + G_i(x, y) < 0. \tag{4.45}$$

Now can we plug in $S^{(2)}(x, y)$ to (4.45) and add that to (4.45), resulting in

$$\sum_{i=3}^N G_i(S^{(1)} \circ S^{(2)}(x, y)) + G_i(S^{(2)}(x, y)) + G_i(S^{(1)}(x, y)) + G_i(x, y) < 0.$$

Recursively applying $S^{(i+1)}$, we obtain

$$\sum_{l=1}^{N-1} \sum_{w \in W_l} G_n(w(x, y)) < 0, \quad W_l := \left\{ S^{(i_1)} \circ \dots \circ S^{(i_l)} \mid i_j \in [1, \dots, N-1], i_j \neq i_k \forall j \neq k \right\}.$$

Finally, applying the last $S^{(N)}$ and summing results in complete cancellation on the left-hand side and a contradiction. Therefore $\eta \geq 0$. $\square$

4.3.2.3 Example (I): Bifunctions associated with monotone operators in Hilbert spaces

For the Hilbert space example where $B(y, x) := \langle A(x), y - x \rangle + \varphi(y) - \varphi(x)$ with $x, y \in D$, the primal equilibrium problem is a solution x^* of the system of variational inequalities:

$$\langle A(x^*), y - x^* \rangle + \varphi(y) \geq \varphi(x^*) \quad \text{for every } y \in D,$$

whereas the dual problem aims for $y_* \in D$:

$$\langle A(x), y_* - x \rangle + \varphi(y_*) \leq \varphi(x) \quad \text{for every } x \in D.$$

We are in the classic framework of [14]. If the sublevels $\{x \in D \mid \langle A(x), x - y \rangle + \varphi(x) \leq c\}$ are compact for every $y \in D, c \in \mathbb{R}$ then the primal problem has a solution (cf. [14, Theorem 10.3]).

Turning to the dual problem, an example of natural conditions involves monotonicity of A , more precisely for Theorem 4.22

1. A is monotone:

$$\langle A(x) - A(y), x - y \rangle \geq -\eta \|x - y\|^2 \quad \text{for every } x, y \in D.$$

2. D is closed, φ is lower semicontinuous and ξ -uniformly convex for some $\xi > 0$,
3. A is hemicontinuous, i.e. for all $x_0, x_1 \in D, v \in H$ the map $t \mapsto \langle A((1-t)x_0 + tx_1), v \rangle$ is continuous in $[0, 1]$.

Concerning the ξ -uniform convexity of φ , it is worth noticing that for the application to evolution problems the function φ is typically perturbed by the squared norm: for $\tau > 0$ the variational movement scheme involves $\varphi_\tau(x) := \frac{1}{2\tau} \|x - \bar{x}\|^2 + \varphi(x)$, which is $1/\tau$ -uniformly convex whenever φ is convex. So our metric setting includes the case of hemicontinuous monotone operators defined on closed convex subsets of a Hilbert space.

4.3.2.4 Example (II): Bifunctions and equilibrium solutions for multispecies systems

In the coupled Wasserstein-2 gradient flow example (Example (II)), we consider the conditions under which the primal equilibrium problem has a solution x^* :

$$\sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)*}, y^{(-i)}) \geq 0 \quad \text{for every } y \in D,$$

and conditions under which the dual problem has a solution y_* :

$$\sum_{i=1}^N F^{(i)}(y_*^{(i)}, y_*^{(-i)}) - F^{(i)}(x^{(i)}, y_*^{(-i)}) \leq 0, \quad \text{for every } x \in D.$$

For the primal problem, Theorem 4.16 requires $y \mapsto \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)})$ to be barycentrically quasi-convexlike with respect to a system of continuous barycentric maps $\mathcal{B}$ in D . This condition is possible to satisfy, with difficulty depending on the particular choice of cross terms in $F^{(i)}$, that is, terms in $F^{(i)}$ that depend on $y^{(j)}$ for $j \neq i$. The second condition is easier; the superlevels

$$\sigma^+(y) = \left\{ x \in D : \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) \geq 0 \right\}$$

need to be closed in D and at least one compact; sufficient smoothness and convexity of each $F^{(i)}(\cdot, y^{(i)})$ results in compactness.

To satisfy the requirements of Theorem 4.22, the first two conditions hold if $(F^{(i)})$ are $\lambda > 0$ -monotone according to Definition 4.25. Next, a sufficient condition for the third criterium is the existence of some x_0 such that each $F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x_0^{(i)}, y^{(-i)})$ has a minimizer over y , which again may be challenging or easy depending on the cross terms in $F^{(i)}$. Finally, we require that $x \mapsto \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)})$ is upper-hemicontinuous with respect to $\mathcal{B}_2$, which is equivalent to lower-semicontinuity of $F^{(i)}(\cdot, y^{(-i)})$.

4.4 Metric λ -dissipative evolutions

In this section we establish some general properties for dissipative evolution in metric spaces, governed by a suitable class of evolution variational inequalities.

4.4.1 Driving bifunctions and EVI flows.

Let (X, d_X) be a metric space, let $D \subset X$ and let $\eta, \kappa \in \mathbb{R}$ with $\lambda := \eta - \kappa$, and let $b : D \times D \rightarrow \mathbb{R}$. We set

$$b^\kappa(y, x) := b(y, x) - \frac{\kappa}{2} d_X^2(x, y), \quad x, y \in D \times D, \quad (4.46)$$

and we introduce a suitable notion of slope related to b .

Definition 4.28 (Global slope of b). For every $x \in D$ we set

$$\Delta^\kappa(x) := \sup_{y \in D} \frac{b^\kappa(y, x)_-}{d_X(x, y)}, \quad \Delta(x) := \inf_{\kappa \in \mathbb{R}} \Delta^\kappa(x); \quad \Delta^\kappa(x), \Delta(x) \in [0, +\infty]. \quad (4.47)$$

We set $D^\kappa(\Delta) := \{x \in D : \Delta^\kappa(x) < \infty\}$ and $D(\Delta) := \{x \in D : \Delta(x) < \infty\}$.

Remark 4.29. In order to better understand the role of $\Delta^\kappa(x)$ it could be useful to observe that in the simple case $b(y, x) = \varphi(y) - \varphi(x)$ for a λ -geodesically convex function $\varphi : D \rightarrow \mathbb{R}$, $\Delta^\kappa(x)$ coincides with the metric slope of φ [9]

$$\Delta^\kappa(x) = |\partial\varphi|(x) = \limsup_{y \rightarrow x} \frac{(\varphi(y) - \varphi(x))_-}{d_X(y, x)} \quad \text{for every } \kappa \leq \lambda.$$

In the operator-Hilbertian case $b(y, x) = \langle A(x), y - x \rangle$ where A is globally defined, $\Delta^\kappa(x) = \|A(x)\|$ for every $\kappa \leq 0$. In both cases $D(\Delta)$ is dense in D .

Definition 4.30. Let $\kappa, \eta \in \mathbb{R}$ be given real parameters, with $\lambda := \eta - \kappa$. We say that $b : D \times D \rightarrow \mathbb{R}$ is an admissible (κ, η) -interaction-dissipative bifunction if

1. for every $x \in D$, $y \mapsto b(y, x)$ has complete sublevels (thus it is also lower semicontinuous);
2. there exist $x_0 \in D$, $a, b \geq 0$ such that

$$b(y, x_0) \geq -a - b d_X^2(x_0, y) \quad \text{for every } y \in D; \quad (4.48)$$

3. it satisfies

$$b(x, x) = 0, \quad b(x, y) + b(y, x) \leq \eta d_X^2(x, y) \quad \text{for every } x, y \in D. \quad (4.49)$$

4. For some $\kappa \in \mathbb{R}$, $D^\kappa(\Delta)$ is dense in D and the following conditional upper semicontinuity holds

$$y, x, x_n \in D, \quad x_n \rightarrow x, \quad \sup_n \Delta^\kappa(x_n) < +\infty \quad \Rightarrow \quad \limsup_{n \rightarrow \infty} b(y, x_n) \leq b(y, x). \quad (4.50)$$

Remark 4.31. The simplest situation is when b is skew-symmetric, i.e.

$$b(x, y) + b(y, x) = 0, \quad \text{i.e.} \quad b(y, x) = -b(x, y), \quad \text{for every } x, y \in D. \quad (4.51)$$

In this case, $x \mapsto b(y, x)$ is upper semicontinuous and therefore condition 4. above reduces to the density of $D^\kappa(\Delta)$ in D .

It is worth noticing that (4.50) implies that $\Delta^\kappa(\cdot)$ is lower semicontinuous. In fact, if $(x_n)_{n \in \mathbb{N}}$ is a sequence in D converging to $x \in D$ such that $\Delta^\kappa(x_n) \leq L$, then we have

$$\frac{\kappa}{2} d_X^2(y, x_n) \leq L d_X(y, x_n) + b(y, x_n)$$

and we can pass to the limit in the inequality as $n \rightarrow \infty$ thanks to (4.50), obtaining $\Delta^\kappa(x) \leq L$ as well.

We extend b to $\bar{D} \times D$ by setting

$$b(y, x) := \inf \left\{ \liminf_{n \rightarrow \infty} b(y_n, x) : y_n \in D, y_n \rightarrow y \right\} \quad \text{for every } y \in \bar{D}, x \in D. \quad (4.52)$$

Notice that the notation (4.52) is coherent when $y \in D$ thanks to the lower semicontinuity of b with respect to y .

We want to study continuous solutions $x : [0, +\infty) \rightarrow D$ of the b -Evolution Variational Inequality

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), y) \leq \bar{b}^\kappa(y, x(t)) \quad t \in (0, +\infty), \quad \text{for every } y \in D; \quad (\text{EVI}_{2,\kappa})$$

Observing that

$$\bar{b}^\kappa(y, x) \leq (\eta - \frac{\kappa}{2}) d_X^2(y, x) - b(x, y) = -b^{2\eta-\kappa}(x, y) \quad (4.53)$$

we will also consider the dual relaxed formulation

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), y) \leq -b^{2\eta-\kappa}(x(t), y) \quad t \in (0, +\infty), \quad \text{for every } y \in D^\kappa(\Delta). \quad (\text{d-EVI}_{2,2\eta-\kappa})$$

for a continuous curve x with values in $\bar{D}$. Thanks to (4.53), clearly any solution to $(\text{EVI}_{2,\kappa})$ also solves $(\text{d-EVI}_{2,2\eta-\kappa})$.

Example 4.32. Choosing $b(y, x) := \varphi(y) - \varphi(x)$ ($\text{EVI}_{2,\kappa}$) corresponds to the EVI_κ formulation of metric gradient flows [9, 171].

Example 4.33. If X is a Hilbert space with scalar product $\langle \cdot, \cdot \rangle$ and $A : D \rightarrow X$ is monotone operator, we can choose

$$b(y, x) := \langle Ax, y - x \rangle, \quad \kappa = 0; \quad (4.54)$$

$(\text{EVI}_{2,\kappa})$ is a metric formulation of the differential equation

$$\frac{d}{dt} x(t) + Ax(t) = 0. \quad (4.55)$$

Definition 4.34. A strong solution $x : [0, +\infty) \rightarrow D^\kappa(\Delta)$ to $(\text{EVI}_{2,\kappa})$ is a locally Lipschitz curve satisfying $(\text{EVI}_{2,\kappa})$ pointwise a.e. in $(0, \infty)$.

A mild solution $\mathbf{x}$ is a continuous curve with values in $\overline{D}$ which is the locally uniform limit of a sequence of strong solutions.

A weak solution is a continuous curve $\mathbf{x} : [0, +\infty) \rightarrow \overline{D}$ solution to $(d\text{-EVI}_{2,2\eta-\kappa})$ according to one of the equivalent form of Lemma 4.12.

A s-EVI flow in $D \subset X$ is a family of continuous maps $S_t : \overline{D} \rightarrow \overline{D}$ such that for every $x_0 \in D$

$$S_{t+h}(x_0) = S_h(S_t(x_0)), \quad \lim_{h \downarrow 0} S_h(x_0) = x_0, \quad t \mapsto S_t(x_0) \text{ is a solution to } (\text{EVI}_{2,\kappa}) \quad (4.56)$$

and whose restriction to $D^\kappa(\Delta)$ is a strong solution to $(\text{EVI}_{2,\kappa})$.

Using (4.53), it is clear that every strong solution is also a weak solution. According to Lemma 4.12, weak solutions can be characterized by the integral inequalities

$$\frac{1}{2}d_X^2(y, \mathbf{x}(t)) - \frac{1}{2}d_X^2(y, \mathbf{x}(s)) \leq \int_s^t b^{2\eta-\kappa}(\mathbf{x}(r), y) dr \quad \text{for every } y \in D^\kappa(\Delta), \quad 0 < s < t. \quad (4.57)$$

Remark 4.35. If $\mathbf{x}$ is a strong solution and L is its Lipschitz constant in $[0, T]$ we easily get

$$-b(y, \mathbf{x}(t)) \leq L d_X(y, \mathbf{x}(t)) \quad \text{for every } y \in D, \quad t \in [0, T],$$

so that $\sup_{t \in [0, T]} \Delta^\kappa(\mathbf{x}(t)) \leq L$. It follows that the map $t \mapsto b(y, \mathbf{x}(t))$ is upper semicontinuous and in particular locally bounded from above for every $y \in D$.

If $\mathbf{x}$ is a weak solution, the map $t \mapsto -b^{2\eta-\kappa}(\mathbf{x}(t), y)$ is upper semicontinuous, so that Lemma 4.12 can be applied as well.

4.4.2 Equivalent formulations

We introduce the function $b_1 : D \times D \rightarrow \mathbb{R}$ defined for every $x, y \in D$ by

$$b_1(y, x) := \frac{b(y, x)}{d_X(x, y)} \quad \text{if } x \neq y, \quad b_1(x, x) := 0, \quad (4.58)$$

and thus satisfying

$$b_1(y, x) + b_1(x, y) \leq \eta d_X(x, y). \quad (4.59)$$

As in (4.46) we consider the perturbations

$$b_1^\kappa(y, x) := b_1(y, x) - \frac{\kappa}{2} d_X(x, y), \quad x, y \in D, \quad (4.60)$$

resulting in

$$b_1^\kappa(y, x) + b_1^\kappa(x, y) \leq \lambda d_X(x, y). \quad (4.61)$$

Lemma 4.36. *A locally Lipschitz curve $x : [0, +\infty) \rightarrow D^\kappa(\Delta)$ is a strong solution to $(\text{EVI}_{2,\kappa})$ if and only if it is a solution to*

$$\frac{d}{dt} d_X(x(t), y) \leq b_1^\kappa(y, x(t)) \quad \text{a.e. in } (0, \infty), \text{ for every } y \in D. \quad (\text{EVI}_{1,\kappa})$$

(or, equivalently, according to one of the forms of Lemma 4.12).

A continuous curve $x : [0, +\infty) \rightarrow \bar{D}$ is a weak solution to $(\text{d-EVI}_{2,2\eta-\kappa})$ if and only if it is a solution to

$$\frac{d}{dt} d_X(x(t), y) \leq -b_1^{2\eta-\kappa}(x(t), y) \quad \text{for every } y \in D^\kappa(\Delta). \quad (\text{d-EVI}_{1,2\eta-\kappa})$$

according to one of the equivalent forms of Lemma 4.12.

Proof. For strong solutions the equivalence is easy since the locally Lipschitz condition implies

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), y) = d_X(x(t), y) \cdot \frac{d}{dt} d_X(x(t), y), \quad \frac{d}{dt} d_X(x(t), y) = 0 \text{ a.e. where } x(t) = y.$$

In the case of weak solutions, we apply Lemma 4.12 and we prove the equivalence of $(\text{d-EVI}_{1,2\eta-\kappa})$ and $(\text{d-EVI}_{2,2\eta-\kappa})$ using the lower left Dini derivative.

The implication $(\text{d-EVI}_{1,2\eta-\kappa}) \Rightarrow (\text{d-EVI}_{2,2\eta-\kappa})$ is an immediate consequence of the formula

$$\frac{1}{2} \frac{d}{dt_-} \zeta^2(t) = \liminf_{h \downarrow 0} \frac{\zeta^2(t) - \zeta^2(t-h)}{2h} = \zeta(t) \liminf_{h \downarrow 0} \frac{\zeta(t) - \zeta(t-h)}{h} = \zeta(t) \frac{d}{dt_-} \zeta(t) \quad (4.62)$$

which holds for every nonnegative continuous function ζ .

Let us now prove the converse implication: we suppose that $(\text{d-EVI}_{2,2\eta-\kappa})$ holds, we fix $y \in D^\kappa(\Delta)$, and we want to prove that

$$\frac{d}{dt_-} d_X(x(t), y) \leq -b_1^{2\eta-\kappa}(x(t), y) \quad \text{for every } t \in (0, \infty) \text{ such that } x(t) \neq y, \quad (4.63)$$

and that the right-hand side of $(\text{d-EVI}_{1,2\eta-\kappa})$ is locally bounded from above. (4.63) can be obtained by applying (4.62).

We also have

$$-b_1^{2\eta-\kappa}(x(t), y) = -b_1^\kappa(x(t), y) + (\eta - \kappa) d_X(x(t), y) \leq \left[\Delta^\kappa(y) + \eta - \kappa \right] d_X(x(t), y),$$

which shows the required local boundedness. $\square$

Theorem 4.37. *Let x be a strong solution to $(\text{EVI}_{2,\kappa})$ and let y be a weak solution to $(\text{d-EVI}_{2,2\eta-\kappa})$. Then*

$$d_X(y(t), x(t)) \leq e^{\lambda(t-s)} d_X(y(s), x(s)) \quad \text{whenever } 0 \leq s \leq t. \quad (4.64)$$

If for every $x_0 \in D^k(\Delta)$ there is a (strong) solution to $(\text{EVI}_{2,\kappa})$ starting from x_0 , then it is possible to define a unique s-EVI flow $S_t : \bar{D} \rightarrow \bar{D}$ of Lipschitz maps. In particular there is at most one weak solution to $(\text{EVI}_{2,\kappa})$ starting from $x_0 \in \bar{D}$.

If $x_0 \in D$ then any continuous solution to (4.67) is locally Lipschitz. Any locally absolutely continuous solution with values in D is locally Lipschitz and satisfies

$$|\dot{x}^+|(t) \leq e^{\lambda(t-s)} |\dot{x}^+|(s) \quad \text{whenever} \quad 0 \leq s \leq t. \quad (4.65)$$

Proof. Let us first consider the case in which also y is locally absolutely continuous. Then using [9, Lemma 4.3.4], setting $g(t, s) = d_X(y(t), x(s))$ and using the equivalence of right- and left-hand derivatives,

$$\begin{aligned} \frac{d}{dt} g(t, t) &\leq \limsup_{h \downarrow 0} \frac{d_X(x(t), y(t)) - d_X(x(t-h), y(t))}{h} \\ &\quad + \limsup_{h \downarrow 0} \frac{d_X(x(t), y(t+h)) - d_X(x(t), y(t))}{h} \\ &= \frac{d}{ds} [d_X(x(s), y(t))] \Big|_{s=t} + \frac{d}{ds} [d_X(x(t), y(s))] \Big|_{s=t} \end{aligned}$$

we get by Lemma 4.36

$$\frac{d}{dt} d_X(y(t), x(t)) \leq b_1^\kappa(y(t), x(t)) + b_1^{\kappa+2\lambda}(x(t), y(t)) = \lambda d_X(y(t), x(t)) \quad \text{a.e. in } (0, +\infty) \quad (4.66)$$

which immediately yields (4.64). Applying the same estimate to the solutions $x(t)$ and $x_h(t) := x(t+h)$ we get

$$d_X(x(t+h), x(t)) \leq e^{\lambda(t-s)} d_X(x(s+h), x(s)) \quad \text{if } 0 < s < t;$$

dividing by h and passing to the limit as $h \downarrow 0$ we get (4.65), since the left and right limits coincide for absolutely continuous curves. The fact that x is locally Lipschitz in $(0, +\infty)$ follows since $|\dot{x}^+|(s) < \infty$ for $\mathcal{L}^1$ -a.e. $s > 0$. Finally, strong (Lipschitz) solutions can be extended to the closure of D by completeness of the underlying metric space. The resulting curve is continuous but not necessarily locally Lipschitz.

In the case when y is just a continuous weak solution, we argue by a standard doubling variable technique (see e.g. [175] and the appendix). We extend x in the interval $(-\infty, 0)$ by setting $x(t) := x(0)$ for $t < 0$. Let us set $\zeta(s, t) := e^{-\frac{1}{2}\lambda(s+t)} d_X(x(s), y(t))$; we easily see that

$$\frac{\partial}{\partial s} \zeta(s, t) + \frac{\partial}{\partial t} \zeta(s, t) \leq f(s, t) \quad \text{in } \mathcal{D}'(\mathbb{R} \times (0, +\infty))$$

where

$$f(s, t) := \begin{cases} 0 & \text{if } s > 0, \\ -e^{-\frac{\lambda}{2}(s+t)} b_1^{\kappa+\lambda}(y(t), x(0)) & \text{if } s \leq 0. \end{cases}$$

For every $T > 0$ we thus get

$$e^{-\lambda T} \int_{T-\varepsilon}^T d_X(y(T), x(t)) dt \leq \varepsilon d_X(y(0), x(0)) + \iint_{Q(\varepsilon, T)} f(s, t) ds dt$$

where $Q(\varepsilon) := \{(s, t) : -\varepsilon < s < 0, 0 < t < s + \varepsilon\}$. Dividing by $\varepsilon > 0$ and passing to the limit as $\varepsilon \downarrow 0$, since the integrand is uniformly bounded from above being $x(0) \in D^\kappa(\Delta)$, we obtain (4.64). $\square$

4.4.3 Discrete b-EVI

Definition 4.38. Let $\tau > 0$. A sequence $(X_\tau^n)_{n \in \mathbb{N}}$ is a solution of the discrete Evolution Variational Inequality generated by b with step size τ if

$$\frac{1}{2\tau} \left(d_X^2(X_\tau^n, Y) - d_X^2(X_\tau^{n-1}, Y) \right) + \frac{1}{2\tau} d_X^2(X_\tau^n, X_\tau^{n-1}) \leq b^\kappa(Y, X_\tau^n). \quad (4.67)$$

Lemma 4.39. If a sequence $(X_\tau^n)_{n \in \mathbb{N}}$ is a solution of (4.67) then it also solves

$$\frac{1}{\tau} \left(d_X(X_\tau^n, Y) - d_X(X_\tau^{n-1}, Y) \right) \leq b_1^\kappa(Y, X_\tau^n). \quad (4.68)$$

Proof. Let us first suppose $Y \neq X_\tau^n$. Setting $a := d_X(X_\tau^n, Y)$, $b := d_X(X_\tau^{n-1}, Y)$, $c = b_1^\kappa(Y, X_\tau^n)$, since $|a - b| \leq d_X(X_\tau^n, X_\tau^{n-1})$, (4.67) yields

$$\frac{1}{2\tau} (a^2 - b^2 + (a - b)^2) \leq ca,$$

which, after developing the squares, yields

$$\frac{1}{\tau} (a^2 - ab) \leq ca;$$

dividing by $a > 0$ we get (4.68). Taking now the infimum of all the limits along sequences Y_n in D converging to X_τ^n we get (4.68) also for $Y = X_\tau^n$. $\square$

Lemma 4.40. Define $c(\tau, T) := (1 - \lambda\tau)^{-(T/\tau+1)}$ and $\tilde{\tau} := \left(\frac{1}{\tau} - \lambda\right)^{-1}$. Every discrete solution (X_τ^n) with $X_\tau^0 \in D^\kappa(\Delta)$ satisfies

$$\tau \Delta^\kappa(X_\tau^n) \leq d_X(X_\tau^n, X_\tau^{n-1}) \leq \tilde{\tau} \Delta^\kappa(X_\tau^{n-1}).$$

This results in $(X_\tau^n)_{n \geq 0} \in D^\kappa(\Delta)$ and

$$\Delta^\kappa(X_\tau^n) \leq c(\tau, T) \Delta^\kappa(X_\tau^0) \leq c(\tau_0, T) \Delta^\kappa(X_\tau^0), \quad \tau \in [0, \tau_0], \quad n \in [0, \dots, \lfloor T/\tau \rfloor]. \quad (4.69)$$

Proof. Choosing $Y := X_\tau^{n-1}$ we get

$$\frac{d_X(X_\tau^n, X_\tau^{n-1})}{\tau} \stackrel{(4.68)}{\leq} b_1^\kappa(X_\tau^{n-1}, X_\tau^n) \stackrel{(4.61)}{\leq} -b_1^\kappa(X_\tau^n, X_\tau^{n-1}) + \lambda d_X(X_\tau^n, X_\tau^{n-1}), \quad (4.70)$$

which results in

$$\left(\frac{1}{\tau} - \lambda\right) d_X(X_\tau^n, X_\tau^{n-1}) \leq -b_1^\kappa(X_\tau^n, X_\tau^{n-1}) \leq \Delta^\kappa(X_\tau^{n-1}).$$

On the other hand, changing sign we obtain (second inequality is triangle inequality)

$$-b_1^\kappa(Y, X_\tau^n) \stackrel{(4.68)}{\leq} \frac{1}{\tau} \left(d_X(X_\tau^{n-1}, Y) - d_X(X_\tau^n, Y) \right) \leq \frac{d_X(X_\tau^n, X_\tau^{n-1})}{\tau},$$

and taking the supremum over Y , we obtain

$$\Delta^\kappa(X_\tau^n) \leq \frac{1}{\tau} d_X(X_\tau^n, X_\tau^{n-1}), \quad \left(\frac{1}{\tau} - \lambda\right) d_X(X_\tau^n, X_\tau^{n-1}) \leq \Delta^\kappa(X_\tau^{n-1}).$$

□

Combining these estimates,

$$\tau \Delta^\kappa(X_\tau^n) \leq d_X(X_\tau^n, X_\tau^{n+1}) \leq \left(\frac{1}{\tau} - \lambda\right)^{-1} \Delta^\kappa(X_\tau^{n-1}),$$

which means that

$$\Delta^\kappa(X_\tau^n) \leq (1 - \tau\lambda)^{-1} \Delta^\kappa(X_\tau^{n-1}).$$

Applying this inequality recursively, we have

$$\Delta^\kappa(X_\tau^n) \leq (1 - \tau\lambda)^{-n} \Delta^\kappa(X_\tau^0).$$

In the case where $\lambda \leq 0$, $\Delta^\kappa(X_\tau^n) \leq \Delta^\kappa(X_\tau^0)$. Otherwise, since $n \leq \frac{T}{\tau} + 1$,

$$(1 - \tau\lambda)^{-n} \leq (1 - \tau\lambda)^{-(T/\tau+1)} =: c(\tau, T),$$

where $c(\tau, T)$ satisfies $c(\tau, T) \leq c(\tau_0, T)$ for all $\tau \in [0, \tau_0]$. The final result is

$$\Delta^\kappa(X_\tau^n) \leq c(\tau_0, T) \Delta^\kappa(X_\tau^0) \quad \forall \tau \in [0, \tau_0], \quad n \in [0, \dots, \lfloor T/\tau \rfloor].$$

We call $\bar{X}_\tau$ the piecewise constant function taking the value X_τ^n in $((n-1)\tau, n\tau]$. $\underline{X}_\tau$ takes the value X_τ^{n-1} in the same interval.

Theorem 4.41. *If X_τ^0 belongs to $D^\kappa(\Delta)$, the family $(\bar{X}_\tau)_\tau$ constructed from the discrete solution to (4.67) satisfies the Cauchy condition*

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq (1 + \lambda T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) \\ &\quad + \left[\hat{c}(\tau_0, T)(1 + \lambda T) + \frac{3}{2^{5/4}} \sqrt{\hat{c}(\tau_0, T) h_{\tau_1}(T) + h_{\tau_2}(T)} \right] (\Delta^\kappa(X_{\tau_1}^0) + \Delta^\kappa(Y_{\tau_2}^0)) \end{aligned}$$

where $\hat{c}(\tau, T) := \frac{T+\tau}{1-\lambda\tau}$, and $h_\tau(T) = T + \frac{\tau}{2}\ell_\tau(T)(1 - \ell_\tau(T))$.

The proof follows the approach in [175] to obtain the error estimate and is in the appendix.

Theorem 4.42. *Let X be complete and $(\bar{X}_\tau)_\tau$ be a solution to (4.67) with $X_\tau^0 \in D^\kappa(\Delta)$. Then there exists $x(t)$ which is the pointwise limit of $\bar{X}_\tau(t)$ as $\tau \rightarrow 0$. The limit $x(t)$ satisfies $(\text{EVI}_{2,\kappa})$.*

Proof. We construct $\bar{X}_\tau$ as the piecewise constant function taking the value X_τ^n in $((n-1)\tau, n\tau]$. We will show there exists $x(t)$ which is the pointwise limit of $\bar{X}_\tau(t)$ as $\tau \rightarrow 0$. From Theorem 4.41,

$$d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) \leq (1 + \tilde{\eta}T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) + \sqrt{2(\tilde{\tau}_1 + \tilde{\tau}_2) h_{\tau_1, \tau_2}(T)} (\Delta^\kappa(X_{\tau_1}^0) + \Delta^\kappa(Y_{\tau_2}^0)),$$

Taking a limit as $\tau_1, \tau_2 \rightarrow 0$, we have $d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) \rightarrow 0$ and therefore the existence of a limit $x(T)$ which satisfies $d_X(\bar{X}_\tau(T), x(T)) \rightarrow 0$ as $\tau \rightarrow 0$. Now we will show that x satisfies $(\text{EVI}_{2,\kappa})$. By definition of X_τ^n and using $\frac{1}{2} \frac{d}{dt} \zeta_\tau(t, Y) = \frac{1}{2\tau} (d_X(X_\tau^n, Y) - d_X(X_\tau^{n-1}, Y))$ for $t \in ((n-1)\tau, n\tau]$,

$$\begin{aligned} \frac{1}{2\tau} (d_X(X_\tau^n, Y) - d_X(X_\tau^{n-1}, Y)) - \frac{\lambda}{2} d_X^2(X_\tau^n, Y) &= \frac{1}{2} \frac{d}{dt} \zeta_\tau(t, Y) - \frac{\lambda}{2} d_X^2(\bar{X}_\tau(t), Y), \\ b(Y, X_\tau^n) - \frac{1}{2\tau} d_X^2(X_\tau^n, X_\tau^{n-1}) &= b(Y, \bar{X}_\tau(t)) - d_X^2(\bar{X}_\tau(t), \underline{X}_\tau(t)), \quad t \in ((n-1)\tau, n\tau]. \end{aligned}$$

Plugging these into (4.67) for all $t \geq 0$,

$$\frac{1}{2} \frac{d}{dt} \zeta_\tau(t, Y) - \frac{\lambda}{2} d_X^2(\bar{X}_\tau(t), Y) \leq b(Y, \bar{X}_\tau(t)) - d_X^2(\bar{X}_\tau(t), \underline{X}_\tau(t)) \quad t \geq 0, \quad (4.71)$$

Integrating both sides of (4.71) from S to T results in

$$\frac{1}{2} (\zeta_\tau(T, Y) - \zeta_\tau(S, Y)) + \int_S^T \left(-\frac{\lambda}{2} d_X^2(\bar{X}_\tau(t), Y) + d_X^2(\bar{X}_\tau(t), \underline{X}_\tau(t)) - b(Y, \bar{X}_\tau(t)) \right) dt \leq 0$$

We would like to take the limit $\tau \rightarrow 0$; in order to estimate the limit for the term $b(Y, \bar{X}_\tau(t))$, we will maximize over all possible sequences $\tau_n \rightarrow 0$:

$$\begin{aligned} \lim_{\tau \rightarrow 0} b(Y, \bar{X}_\tau(t)) &\leq \sup_{n \rightarrow \infty} \{ \limsup b(Y, \bar{X}_{\tau_n}(t)) : \tau_n \rightarrow 0 \} \\ &\leq \sup_{n \rightarrow \infty} \{ \limsup b(Y, X_n) : X_n \rightarrow x(t) \} = \bar{b}(Y, x(t)). \end{aligned}$$

Now taking the limit as $\tau \rightarrow 0$,

$$\frac{1}{2} \left(d_X^2(x(T), Y) - d_X^2(x(S), Y) \right) + \int_S^T \left(-\frac{\lambda}{2} d_X^2(x(t), Y) - \bar{b}(Y, x(t)) \right) dt \leq 0.$$

Finally, we can divide by $T - S$ and take the limit as $|T - S| \rightarrow 0$ to obtain (EVI_{2,κ}),

$$\frac{1}{2} \frac{d}{dt} d_X^2(x(t), Y) \leq \bar{b}^\kappa(Y, x(t)).$$

□

4.5 Solution to the Discrete b-EVI

Let $b : D \times D \rightarrow \mathbb{R}$ be an admissible (κ, η) -interaction-dissipative bifunction according to Definition 4.30. For every $\bar{x} \in X$, $\tau > 0$, we introduce the generating functions $B(\tau, \bar{x}; \cdot, \cdot) : D \times D \rightarrow \mathbb{R}$ defined by

$$B(\tau, \bar{x}; y, x) := \frac{1}{2\tau} d_X^2(\bar{x}, y) - \frac{1}{2\tau} d_X^2(\bar{x}, x) + b(y, x). \quad (4.72)$$

Theorem 4.43 (Resolvent equation). *Let $b : D \times D \rightarrow \mathbb{R}$ be an admissible (κ, η) -interaction-dissipative bifunction, let B be defined as in (4.72), and let $\tau_o > 0$ be such that $\tau_o^{-1} + \kappa > \eta_+$. For every $\bar{x} \in D$ and for every $\tau \in (0, \tau_o)$ we suppose that*

the family of functions $y \mapsto B(\tau, \bar{x}; y, x)$, $x \in D$, are jointly barycentrically $\left(\frac{1}{\tau} + \kappa\right)$ -convexlike (with respect to some system $\mathcal{B}$ possibly depending on τ and $\bar{x}$) and for every $y \in D$ the maps $x \mapsto b(y, x)$ are $\mathcal{B}_2$ -upper hemicontinuous.

Then for every $\bar{x} \in D$ there exists a unique $x_\tau \in D$ solving the equilibrium problems

$$B(\tau, \bar{x}; y, x_\tau) \geq 0 \quad \text{for every } y \in D, \quad (4.73)$$

$$B(\tau, \bar{x}; x_\tau, x) \leq 0 \quad \text{for every } x \in D. \quad (4.74)$$

Moreover, every solution to (4.73) also solves

$$\frac{1}{2\tau} \left(d_X^2(x_\tau, y) - d_X^2(\bar{x}, y) + d_X^2(\bar{x}, x_\tau) \right) \leq \bar{b}^\kappa(y, x_\tau). \quad (4.75)$$

Proof. Let us fix $\tau, \bar{x}$ as in the statement of the Theorem and let us define $B(y, x) := B(\tau, \bar{x}; y, x)$. We have to check that B satisfies the properties of Theorem 4.22 with $\xi := \frac{1}{\tau} + \kappa$.

(4.37) is immediate since

$$B(\tau, \bar{x}; y, x) + B(\tau, \bar{x}; x, y) = b(y, x) + b(x, y) \leq \eta d_X^2(y, x).$$

The second condition is satisfied as well, so that we derive a solution x_τ of the dual formulation (4.74). In order to obtain (4.73) it is sufficient to check that the map $x \mapsto B(\tau, \bar{x}; y, x)$ satisfies the conditions of the metric Minty's Lemma, which are guaranteed by the upper hemicontinuity of $x \mapsto b(y, x)$.

In order to prove (4.75), we fix $y \in D$ and we select a path $\mathbf{x} \in \mathcal{B}_2$, with $\mathbf{x}_1 := x_\tau$ and $\mathbf{x}_2 := y$; we call $y_\vartheta := \mathbf{x}((1 - \vartheta), \vartheta)$, $\vartheta \in (0, 1)$. We get

$$\begin{aligned} 0 &\leq B(\tau, \bar{x}; y_\vartheta, x_\tau) \\ &\leq (1 - \vartheta)B(\tau, \bar{x}; x_\tau, x_\tau) + \vartheta B(\tau, \bar{x}; y, x_\tau) - \frac{1}{2}(\kappa + 1/\tau)\vartheta(1 - \vartheta)d_X^2(y, x_\tau) \end{aligned}$$

so that using $B(\tau, \bar{x}; x_\tau, x_\tau) = 0$, rearranging the terms, dividing by $\vartheta > 0$ and passing to the limit as $\vartheta \downarrow 0$ we get

$$B(\tau, \bar{x}; y, x_\tau) \geq \frac{1}{2}(\kappa + 1/\tau)d_X^2(y, x_\tau)$$

which yields (4.75). $\square$

Remark 4.44 (Metric Space Curvature). *Implicit in point 1 of Theorem 4.43 is the curvature of (X, d_X) . In particular, $\frac{1}{\tau} + \kappa$ barycentric convexity in condition 1 in Theorem 4.43 implies that*

$$\frac{1}{2\tau}d_X^2(\bar{x}, \mathbf{x}(\alpha)) + b(\bar{x}, \mathbf{x}(\alpha)) \leq \sum_{j=1}^J \alpha_j \left[\frac{1}{2\tau}d_X^2(\mathbf{x}_j, \bar{x}) + b(\bar{x}, \mathbf{x}_j) \right] - \frac{\tau^{-1} + \kappa}{4} \sum_{j,k=1}^J \alpha_j \alpha_k d_X^2(\mathbf{x}_j, \mathbf{x}_k).$$

Since this needs to hold for all $\tau \rightarrow 0$, it must hold for all the τ -dependent terms:

$$\frac{1}{2\tau}d_X^2(\bar{x}, \mathbf{x}(\alpha)) \leq \frac{1}{2\tau} \sum_{j=1}^J \alpha_j d_X^2(\mathbf{x}_j, \bar{x}) - \frac{1}{4\tau} \sum_{i,j=1}^J \alpha_i \alpha_j d_X^2(\mathbf{x}_i, \mathbf{x}_j). \quad (4.76)$$

For $J = 2$ and a geodesic interpolant, this inequality does not hold for positively curved spaces. Compare this with the CAT(0) / nonpositively curved metric space property, where x is a geodesic path from $x_0 \in X$ to $x_1 \in X$,

$$d_X^2(y, \mathbf{x}(t)) \leq t d_X^2(y, x_0) + (1 - t) d_X^2(y, x_1) - t(1 - t) d_X^2(x_0, x_1).$$

In Hilbert spaces (and all flat spaces), the above inequality holds with equality for $\mathbf{x}$ a geodesic interpolant. By allowing $\mathbf{x}$ to have a non-optimal interpolant (like generalized geodesics), even spaces with positive curvature can potentially satisfy (4.76).

Remark 4.45 (Choice of Discrete Scheme). *From a discrete-time gradient flow perspective, consider the following two natural options for how each species i can update to move in the direction of steepest descent of $F^{(i)}$. One option is minimizing $F^{(i)}(\cdot, (X_\tau^{n-1})^{(-i)})$ with a distance penalty,*

$$X_\tau^{(i),n} = \arg \min_{x^{(i)} \in X^{(i)}} \left\{ F^{(i)}(x^{(i)}, X_\tau^{(-i),n-1}) + \frac{1}{2\tau} d_{X^{(i)}}^2(x^{(i)}, X_\tau^{(i),n-1}) \right\}, \quad \text{for every } i \in [1, \dots, N].$$

This update is implicit in $x^{(i)}$ but explicit in $x^{(-i)}$ because of the dependence on $X_\tau^{(-i),n-1}$. The second option is to solve

$$X_\tau^n \text{ is a Nash equilibrium of } \left\{ x \mapsto F^{(i)}(x^{(i)}, x^{(-i)}) + \frac{1}{2\tau} d_{X^{(i)}}^2(x^{(i)}, X_\tau^{(i),n-1}) \right\}_i,$$

which is a fully implicit update. Proving the existence of minimizers for the partially-explicit update is easier than proving existence of a Nash equilibrium for a set of energies. On the other hand, we expect better stability properties from the fully implicit method. In [112], the implicit-explicit update scheme was used along with Lipschitz bounds on the energy functionals to prove existence of solutions for a two-species PDE system with nonlocal interaction potentials and linear coupling potentials. To illustrate the lack of stability of the first method, consider the min-max game where $X^{(1)} = X^{(2)} = \mathbb{R}$, and $F^{(1)}(x^{(1)}, x^{(2)}) = -F^{(2)}(x^{(2)}, x^{(1)}) = \frac{b}{2}(x^{(1)})^2 + ax^{(1)}x^{(2)} - \frac{b}{2}(x^{(2)})^2$ with $b > 0$ and $a \in \mathbb{R}$. The corresponding continuous-time gradient flow solution converges exponentially to the Nash equilibrium $x = 0$ with rate b . The schemes via the implicit-explicit and the Nash update are

$$\begin{aligned} \text{(implicit-explicit)} \quad X_\tau^{n+1} &= (1 + \tau b)^{-1} \begin{bmatrix} 1 & -\tau a \\ \tau a & 1 \end{bmatrix} X_\tau^n, \\ \text{(fully implicit)} \quad X_\tau^{n+1} &= ((1 + \tau b)^2 + \tau^2 a^2)^{-1} \begin{bmatrix} 1 + \tau b & -\tau a \\ \tau a & 1 + \tau b \end{bmatrix} X_\tau^n. \end{aligned}$$

The maximum eigenvalue magnitudes for each system are $\sqrt{a^2\tau^2 + 1}(b\tau + 1)^{-1}$ and $[\tau^2(a^2 + b^2) + 2b\tau + 1]^{-1/2}$, with exponential stability if the magnitudes are less than one. If the coupling term a is larger than the convexity parameter b , the first system is exponentially stable only for a small enough τ , dependent on a and b , whereas the second system is exponentially stable for all $\tau > 0$, $a \in \mathbb{R}$, and $b > 0$. This mirrors the need for an upper-Lipschitz bound on the gradient of the energy for explicit Euler for the stability of gradient descent, along with a small enough time step relative to the Lipschitz constant. For the remainder of this work, we therefore define the update scheme for multispecies coupled gradient flows via the fully implicit Nash equilibrium scheme. This motivated our choice of saddle point scheme (VMS) for general bifunctions b .

Definition 4.46 (Variational Movement Scheme). Let $\tau > 0$ be a given step size and let $X_\tau^0 \in D$. A sequence $(X_\tau^n)_{n \in \mathbb{N}}$ in D is a solution of the recursive Variational Movement Scheme if for every $n > 0$

$$(X_\tau^n, X_\tau^n) \text{ is a saddle point of } (y, x) \mapsto B(\tau, X_\tau^{n-1}; y, x) \quad (\text{VMS})$$

i.e. for every $(x, y) \in D \times D$

$$B(\tau, X_\tau^{n-1}; X_\tau^n, x) \leq 0 = B(\tau, X_\tau^{n-1}; X_\tau^n, X_\tau^n) \leq B(\tau, X_\tau^{n-1}; y, X_\tau^n). \quad (4.77)$$

In this case we define the piecewise constant interpolations $\bar{X}_\tau$ defined by

$$\bar{X}_\tau(0) := X_\tau^0, \quad \bar{X}^\tau(t) := X_\tau^n \quad \text{if } t \in ((n-1)\tau, n\tau]. \quad (4.78)$$

Theorem 4.47. *Under the same assumptions of Theorem 4.43, any VMS solution $(X_\tau^n)_{n \in \mathbb{N}}$ with initial value $X_\tau^0 \in D^\kappa(\Delta)$ also satisfies (4.67).*

Proof. Let us fix $Y \in D$ and let us select a path $\mathbf{x} \in \mathcal{B}_{X_\tau^{n-1}, 2}$, with $\mathbf{x}_1 := X_\tau^n$ and $\mathbf{x}_2 := Y$; we call $X_\vartheta := \mathbf{x}((1-\vartheta), \vartheta)$, $\vartheta \in (0, 1)$. We get

$$\begin{aligned} 0 &\leq B(\tau, X_\tau^{n-1}; X_\vartheta, X_\tau^n) \\ &\leq (1-\vartheta)B(\tau, X_\tau^{n-1}; X_\tau^n, X_\tau^n) + \vartheta B(\tau, X_\tau^{n-1}; Y, X_\tau^n) - \frac{1}{2}(\kappa + 1/\tau)\vartheta(1-\vartheta)d_X^2(Y, X_\tau^n) \end{aligned}$$

so that using $B(\tau, X_\tau^{n-1}; X_\tau^n, X_\tau^n) = 0$, rearranging the terms, dividing by $\vartheta > 0$ and passing to the limit as $\vartheta \downarrow 0$ we get

$$B(\tau, X_\tau^{n-1}; Y, X_\tau^n) \geq \frac{1}{2}(\kappa + 1/\tau)d_X^2(Y, X_\tau^n)$$

which yields

$$\frac{1}{2\tau} \left(d_X^2(Y, X_\tau^n) - d_X^2(Y, X_\tau^{n-1}) + d_X^2(X_\tau^n, X_\tau^{n-1}) \right) \leq b(Y, X_\tau^n) - \frac{\kappa}{2} d_X^2(Y, X_\tau^n). \quad (4.79)$$

□

Lemma 4.48. *If $\Delta^{\kappa_0}(x) < \infty$ for some $\kappa_0 \in \mathbb{R}$ then $\Delta^\kappa(x) = \Delta^{\kappa_0}(x)$ for every $\kappa < \kappa_0$. In particular*

$$D(\Delta) = \left\{ x \in D : \Delta^\kappa(x) < \infty \text{ for some } \kappa \in \mathbb{R}. \right\} \quad (4.80)$$

Proof. Since the map $k \mapsto b^k(y, x)_-$ is increasing, it is easy to check that $\kappa \mapsto \Delta^\kappa(x)$ is increasing as well. Therefore, it is sufficient to check that

$$\kappa < \kappa_0, \quad 0 < \Delta^{\kappa_0}(x) < \infty \quad \Rightarrow \quad \Delta^\kappa(x) \geq \Delta^{\kappa_0}(x).$$

Let $0 < S < \Delta^{\kappa_0}(x)$. We can find $y \in D$ such that

$$\frac{\kappa_0}{2} d_X^2(y, x) - b(y, x) > S$$

□

Theorem 4.49 (Existence Result). *Suppose $\kappa > -\infty$ and $\eta < +\infty$ such that $b : D \times D \rightarrow \mathbb{R}$ is an admissible (κ, η) -interaction-dissipative bifunction. If for every $\bar{x} \in D$ there exists a system $\mathcal{B}_{\bar{x}}$ of barycentric maps such that*

1. the function $d_X^2(\cdot, \bar{x})$ is 1-barycentrically convexlike with respect to $\mathcal{B}_{\bar{x}}$;
2. for every $y \in D$ the functions $y \mapsto b(y, x)$ are barycentrically κ -convexlike with respect to $\mathcal{B}_{\bar{x}}$;
3. for every $y \in D$ the maps $x \mapsto b(y, x)$ are $\mathcal{B}_{\bar{x}, 2}$ -upper hemicontinuous,

then for every $x_0 \in D^k(\Delta)$ there exists a unique continuous curve x satisfying $(\text{EVI}_{2, \kappa})$ with $x(0) = x_0$ and any two solutions x, y satisfy

$$d_X(x(t), y(t)) \leq e^{\lambda(t-s)} d_X(x(s), y(s)), \quad \text{where } \lambda = \eta - \kappa.$$

Proof. Since $\kappa > -\infty$ and $\eta < +\infty$, there exists τ_0 such that for all $\tau \in (0, \tau_0)$, there exists a solution (X_τ^n) to the VMS by Theorem 4.43, which in turn solves discrete EVI (4.67) by Theorem 4.47. From Theorem 4.42, we have the existence of $x(t)$ which satisfies $(\text{EVI}_{2, \kappa})$. We now show that the limiting curve x is continuous. From (4.67), selecting $Y = X_\tau^{k-1}$

$$\frac{2 + \kappa\tau}{2\tau} d_X^2(X_\tau^k, X_\tau^{k-1}) \leq b(X_\tau^{k-1}, X_\tau^k).$$

Using the interpolation $\bar{X}_\tau(t)$, it holds for $s < t$ where $s \in ((m-1)\tau, m\tau]$, $t \in ((n-1)\tau, n\tau]$ that $d_X(\bar{X}_\tau(t), \bar{X}_\tau(s)) \leq \sqrt{n-m} \left(\sum_{k=m+1}^n d_X^2(X_\tau^k, X_\tau^{k-1}) \right)^{1/2}$. Restricting t such that $n \in \{m, m+1\}$, and using that $(n-m)\tau \leq t-s+\tau$, we have

$$d_X^2(\bar{X}_\tau(t), \bar{X}_\tau(s)) \leq \frac{2(t-s-\tau)}{2 + \kappa\tau} b(\bar{X}_\tau(t), \bar{X}_\tau(s)).$$

Taking $\tau \rightarrow 0$, we have

$$d_X^2(x(t), x(s)) \leq (t-s)b(x(t), x(s)).$$

Using the upper-hemicontinuity of $y \mapsto b(y, x)$ and taking $t \rightarrow s$, we have that x is continuous in time. From Theorem 4.37, we obtain the contraction estimate. $\square$

4.5.1 Examples

We discuss examples illustrating the interplay between the convexity of $y \mapsto b(y, x)$ and $d_X^2(\cdot, \bar{x})$.

Example 4.50 (Convexity of Distance Fails for Spherical Hellinger). *Consider a system of coupled gradient flows with respect to the spherical Hellinger metric: for all $i \in [1, \dots, N]$,*

$$\partial_t \rho^{(i)}(x) = -\nabla_{\mathcal{H}, \rho^{(i)}} F^{(i)}[\rho^{(1)}, \dots, \rho^{(n)}] = -\delta_{\rho_i} F^{(i)}[\rho](x) \rho^{(i)}(x) + \int \delta_{\rho_i} F^{(i)}[\rho](x) \rho^{(i)}(x) dx.$$

We would like to prove well-posedness of the solutions to the above system of evolutions. While results in [46] show that a class of f -divergences are geodesically convex in the Hellinger metric, existence of solutions was left as an open problem. In [143], the authors propose a min-max flow system for use in machine learning applications, and study the long time behavior, also leaving well-posedness for future work. Unfortunately, our results do not apply to this setting because there is no known set of curves along which $d_X^2(\cdot, \rho)$ is 1-convex. Therefore, we cannot apply Theorem 4.43 due to poor convexity properties of the spherical Hellinger metric.

Example 4.51 (Nonconvexity and Existence of Solutions). *For the following choice of parameters, the system in [64, Example 4.5], where $\rho^{(1)}, \rho^{(2)} \in \mathcal{P}_2(\mathbb{R})$ and $\kappa > 0$,*

$$\begin{aligned} F^{(1)}[\rho^{(1)}, \rho^{(2)}] &= \frac{\kappa}{2} \int z_1^2 d\rho^{(1)}(z_1) + 2\kappa \iint z_1 z_2 d\rho^{(1)}(z_1) d\rho^{(2)}(z_2), \\ F^{(2)}[\rho^{(1)}, \rho^{(2)}] &= \frac{\kappa}{2} \int z_2^2 d\rho^{(2)}(z_2) + 2\kappa \iint z_1 z_2 d\rho^{(1)}(z_1) d\rho^{(2)}(z_2), \end{aligned}$$

is non-monotone, even though F_1 and F_2 are strongly displacement in ρ_1 and ρ_2 respectively. The coupled Wasserstein-2 gradient flow formulation can be written as a joint gradient flow of an energy functional that has no lower bound. However, Theorem 4.43 can still be applied to provide existence of solutions. Computing the dissipativity, we have, with $\gamma^{(1)} \in \Gamma^*(\rho^{(1)}, \mu^{(1)})$ and $\gamma^{(2)} \in \Gamma^*(\rho^{(2)}, \mu^{(2)})$,

$$\begin{aligned} b(\rho, \mu) + b(\mu, \rho) &= 2\kappa \iiint (z_1 z'_2 + z'_1 z_2 - z_1 z_2 - z'_1 z'_2) d\gamma^{(1)}(z_1, z'_1) d\gamma^{(2)}(z_2, z'_2) \\ &= 2\kappa \iiint (z_1 - z'_1)(z'_2 - z_2) d\gamma^{(1)}(z_1, z'_1) d\gamma^{(2)}(z_2, z'_2) \leq \kappa d_X^2(\rho, \mu), \end{aligned}$$

where the inequality follows from Young's inequality. Therefore $\eta = \kappa$; selecting $\tau_0 = (2\kappa)^{-1}$, the requirements for Theorem 4.43 are satisfied even without strong monotonicity.

Example 4.52 (Single Species). *In the single species gradient flow setting, $X = X^{(1)}$ and energy $F : X \rightarrow \mathbb{R}$. Solving the equilibrium problem parameterized by $B(\tau, X^{n-1}; y, x)$ reduces to solving*

$$\begin{aligned} \frac{1}{2\tau} d_X^2(X^{n-1}, y) - \frac{1}{2\tau} d_X^2(X^{n-1}, X^n) + F(y) - F(X^n) &\geq 0, \quad \forall y \in X, \\ \Rightarrow \frac{1}{2\tau} d_X^2(X^{n-1}, y) + F(y) &\geq \frac{1}{2\tau} d_X^2(X^{n-1}, X^n) + F(X^n), \quad \forall y \in X, \end{aligned}$$

which is the same minimizing movement scheme in [9].

Example 4.53 (Min-Max Evolution). *Let $F : X^{(1)} \times X^{(2)} \rightarrow \mathbb{R}$ be a Lagrangian function satisfying the usual lower and upper semicontinuity condition:*

$$\begin{aligned} \forall x^{(1)} \in X^{(1)} : \quad x^{(2)} &\mapsto F(x^{(1)}, x^{(2)}) \quad \text{is upper semicontinuous in } X^{(2)}, \\ \forall x^{(2)} \in X^{(2)} : \quad x^{(1)} &\mapsto F(x^{(1)}, x^{(2)}) \quad \text{is lower semicontinuous in } X^{(1)}. \end{aligned} \tag{4.81}$$

The bifunction is given by $b(y, x) = F(y^{(1)}, x^{(2)}) - F(x^{(1)}, y^{(2)})$ (due to the cancellation of all other terms), and the generating bifunction is

$$B(\tau, X_\tau^{n-1}; x, y) := \frac{1}{2\tau} d_X^2(x, X_\tau^{n-1}) - \frac{1}{2\tau} d_X^2(y, X_\tau^{n-1}) + F(x^{(1)}, y^{(2)}) - F(y^{(1)}, x^{(2)}). \quad (4.82)$$

The condition (4.77) states that X_τ^n must satisfy

$$\frac{1}{2\tau} d_X^2(X_\tau^n, X_\tau^{n-1}) - \frac{1}{2\tau} d_X^2(y, X_\tau^{n-1}) + F(X_\tau^{(1),n}, y^{(2)}) - F(y^{(1)}, X_\tau^{(2),n}) \geq 0 \quad \forall y \in X,$$

from which, rearranging to separate by species,

$$\begin{aligned} & \frac{1}{2\tau} d_X^2(X_\tau^{(1),n}, X_\tau^{(1),n-1}) - \frac{1}{2\tau} d_X^2(y^{(1)}, X_\tau^{(1),n-1}) + F(X_\tau^{(1),n}, y^{(2)}) \\ & \geq \frac{1}{2\tau} d_X^2(y^{(2)}, X_\tau^{(2),n-1}) - \frac{1}{2\tau} d_X^2(X_\tau^{(2),n}, X_\tau^{(2),n-1}) + F(y^{(1)}, X_\tau^{(2),n}) \quad \forall y \in X, \end{aligned}$$

we observe the left-hand side is convex in $X_\tau^{(1),n}$ and the right-hand side is concave in $X_\tau^{(2),n}$ along generalized geodesics when F is convex-concave along generalized geodesics.

4.5.1.1 Steady State is Nash Equilibrium for Coupled Multispecies Systems

Corollary 4.54 (Steady State is Nash Equilibrium). *Let b correspond to a multispecies gradient flow system with a solution such that the semigroup is continuous. If (D, d) is complete and $\lambda < 0$, then there exists a unique steady state $x_\infty \in D$, which is also the unique Nash equilibrium of $(F^{(i)})$.*

Proof. Using the completeness and the contraction from $\lambda < 0$, by [52, Lemma 7.3] there exists a unique stationary state $x_\infty \in D$. Then x_∞ satisfies $(\text{EVI}_{2,\kappa})$

$$\frac{\kappa}{2} d_{X^{(i)}}(x_\infty, y) \leq \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x_\infty^{(i)}, y^{(-i)}) \quad \text{for all } y^{(i)} \in D.$$

Fixing a particular $i \in [1, \dots, n]$ and selecting $y^{(-i)} = x_\infty^{(-i)}$,

$$\frac{\kappa}{2} d_{X^{(i)}}(x_\infty^{(i)}, y^{(i)}) \leq F^{(i)}(y^{(i)}, x_\infty^{(-i)}) - F^{(i)}(x_\infty^{(i)}, x_\infty^{(-i)}) \quad \text{for all } y^{(i)} \text{ such that } (y^{(i)}, x_\infty^{(-i)}) \in D.$$

Since $\kappa > 0$ (because $\lambda < 0$ and $\eta \geq 0$ by Proposition 4.27),

$$F^{(i)}(x_\infty^{(i)}, x_\infty^{(-i)}) \leq F^{(i)}(y^{(i)}, x_\infty^{(-i)}) \quad \text{for all } y^{(i)} \text{ such that } (y^{(i)}, x_\infty^{(-i)}) \in D.$$

Because this can be shown for all $i \in [1, \dots, n]$, x_∞ is a Nash equilibrium of $(F^{(i)})$. To show that the Nash equilibrium is unique, for the sake of contradiction consider two Nash equilibria $x_* \in D$

and $y_* \in D$ with $d_X(x_*, y_*) > 0$. Since $y \mapsto b(y, x)$ is κ -barycentrically convex, comparing the points x_* and y_* in $b(\cdot, x_*)$,

$$\sum_{i=1}^N F^{(i)}(x_*^{(i)}, x_*^{(-i)}) - F^{(i)}(x_*^{(i)}, x_*^{(-i)}) \leq \sum_{i=1}^N F^{(i)}(y_*^{(i)}, y_*^{(-i)}) - F^{(i)}(x_*^{(i)}, y_*^{(-i)}) - \frac{\kappa}{2} d_X^2(x_*, y_*).$$

This implies that

$$\sum_{i=1}^N F^{(i)}(y_*^{(i)}, y_*^{(-i)}) - F^{(i)}(x_*^{(i)}, y_*^{(-i)}) \geq \frac{\kappa}{2} d_X^2(x_*, y_*).$$

We can similarly compare the points x_* and y_* in $b(\cdot, y_*)$. Combining this with the η -interaction dissipativity inequality,

$$\begin{aligned} \kappa d_X^2(x_*, y_*) &\leq \sum_{i=1}^N F^{(i)}(y_*^{(i)}, y_*^{(-i)}) - F^{(i)}(x_*^{(i)}, y_*^{(-i)}) + F^{(i)}(x_*^{(i)}, x_*^{(-i)}) - F^{(i)}(y_*^{(i)}, x_*^{(-i)}) \\ &\leq \eta d_X^2(x_*, y_*), \end{aligned}$$

we have a contradiction since $\kappa > \eta$. □

The assumption that D is complete is quite strong; for example, if an energy functional on probability measures has an entropy term, the domain is restricted to $L^1 \cap \mathcal{P}_2$, which is not complete. One may require additional regularity in the definition of a steady state, and then it must be shown explicitly that x_∞ has such regularity.

4.5.1.2 Example (I): Existence for Hilbert space evolutions

The main difference in applying Theorem 4.43 versus Theorem 4.22 is the presence of $\frac{1}{2\tau} d^2(x, \bar{x})$ in the bifunction. In order for $b(y, x) = \langle A(x), y - x \rangle$ to satisfy the criteria for Theorem 4.43, we can simply apply the same requirements as discussed in Example 4.3.2.3, because norms on Hilbert spaces satisfy the desired 1-convexity property and continuity properties.

4.5.1.3 Example (II): Convexity along generalized geodesics

Let $(F^{(i)})$ be a coupled gradient flow with each species in the Wasserstein-2 metric space. To show existence of a global Nash equilibrium, we can select geodesics $\mathcal{G}_2$ to show that the maps $y \mapsto b(y, x)$ are κ -displacement convex. In combination with $\eta < \kappa$ dissipativity and sufficient continuity, as discussed in Section 4.3.2.4, we obtain existence of a Nash equilibrium from Theorem 4.22. However, to show well-posedness of the corresponding coupled gradient flow dynamics, we can no longer use $\mathcal{G}_2$ to show that the family $y \mapsto B(\tau, \bar{x}; y, x)$ are barycentrically $\left(\frac{1}{\tau} + \kappa\right)$ -convexlike, because $\mathcal{W}_2$ is not 1-convex along geodesics, and this term appears in the definition of B . This convexity can be shown using generalized geodesics instead, as discussed in point 6 in Example 4.13.

4.6 Multispecies Systems

In this section, we provide additional results and examples for multispecies coupled gradient flows in Euclidean space and the Wasserstein-2 metric. In both settings, we show that the zeroth-order notion of monotonicity that we propose in Definition 4.25 implies the respective first order conditions. In the Euclidean setting, we show how coupling worsens overall dissipativity, and in the Wasserstein-2 setting we provide an example of energy functionals with a commonly-analyzed nonlocal interaction term, and we show how to obtain estimates for dissipativity of the system.

4.6.1 Euclidean Multispecies Gradient Flows

First, we show that monotonicity according to Definition 4.25 implies the classical game-theoretic notion of monotonicity for games on Euclidean space. For an example showing how the first-order condition does not imply the zeroth order condition, see Example 4.26.

Lemma 4.55 (Monotone Functions in $\mathbb{R}^d$). *Let $X^{(i)} = \mathbb{R}^{d_i}$, $X = \mathbb{R}^d$ and $d_{X^{(i)}}$ be the Euclidean norm. If $(F^{(i)}) \in C^1(\mathbb{R}^d; \mathbb{R})$ are $-\lambda$ -monotone with respect to linear interpolations according to Definition 4.25, then the following classical definition of monotonicity holds:*

$$\sum_{i=1}^N \langle \nabla_{x^{(i)}} F^{(i)}(x^{(i)}, x^{(-i)}) - \nabla_{y^{(i)}} F^{(i)}(y^{(i)}, y^{(-i)}), x^{(i)} - y^{(i)} \rangle \geq -\lambda \|x - y\|^2.$$

Proof. To show this, we first use that $y \mapsto b(y, x)$ is κ -barycentrically convex, selecting the family linear interpolations for $\mathcal{B}$. The first-order κ convexity inequality, comparing (y, x) and (x, x) , gives

$$\begin{aligned} \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) &\geq \sum_{i=1}^N \left(F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(x^{(i)}, x^{(-i)}) \right. \\ &\quad \left. + \sum_{j=1}^N \nabla_{y^{(j)}} (F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)})) \Big|_{y=x} (y^{(j)} - x^{(j)}) \right) + \frac{\kappa}{2} \|x - y\|^2. \end{aligned}$$

Since the gradients with respect to $x^{(j)}$ where $j \neq i$ cancel, the inequality reduces to

$$\sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) \geq \sum_{i=1}^N \nabla_{x^{(i)}} F^{(i)}(x^{(i)}, x^{(-i)})(y^{(i)} - x^{(i)}) + \frac{\kappa}{2} \|x - y\|^2.$$

Repeating the calculation with test points (x, y) and (y, y) and summing, we have

$$\begin{aligned} \sum_{i=1}^N F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) &+ F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}) \\ &\geq - \sum_{i=1}^N \langle \nabla_{x^{(i)}} F^{(i)}(x^{(i)}, x^{(-i)}) - \nabla_{y^{(i)}} F^{(i)}(y^{(i)}, y^{(-i)}), y^{(i)} - x^{(i)} \rangle + \kappa \|x - y\|^2. \end{aligned}$$

Now we apply the η interaction dissipativity bound to the left-hand side and rearrange to obtain

$$\sum_{i=1}^N \langle \nabla_{x^{(i)}} F^{(i)}(x^{(i)}, x^{(-i)}) - \nabla_{y^{(i)}} F^{(i)}(y^{(i)}, y^{(-i)}), x^{(i)} - y^{(i)} \rangle \geq (\kappa - \eta) \|x - y\|^2,$$

which is exactly the first-order game theoretic monotonicity condition with $\lambda = \eta - \kappa$. $\square$

The following example illustrates the role of η interaction dissipativity; here it quantifies how much dissipativity is lost due to the coupling between the two players.

Example 4.56 (Non-symmetric two-player game). *This example shows the setting when $\eta > 0$, that is, the coupling term causes contractivity to slow down compared to the convexity of each term independently. Let $X^{(1)} = X^{(2)} = \mathbb{R}$, and*

$$F^{(1)}(x^{(1)}, x^{(2)}) = \frac{a}{2}x^{(1)2} + x^{(1)}x^{(2)}, \quad F^{(2)}(x^{(1)}, x^{(2)}) = \frac{a}{2}(x^{(2)} - b)^2.$$

The second species is independent of the first species, but the first species evolves based on the second. The classical definition of monotonicity is

$$\begin{aligned} & \left\langle \begin{bmatrix} \nabla_{x^{(1)}} F^{(1)}(x^{(1)}, x^{(2)}) - \nabla_{x^{(1)}} F^{(1)}(y^{(1)}, y^{(2)}) \\ \nabla_{x^{(2)}} F^{(2)}(x^{(1)}, x^{(2)}) - \nabla_{x^{(2)}} F^{(2)}(y^{(1)}, y^{(2)}) \end{bmatrix}, \begin{bmatrix} x^{(1)} - y^{(1)} \\ x^{(2)} - y^{(2)} \end{bmatrix} \right\rangle \\ &= a(x^{(1)} - y^{(1)})^2 + (x^{(2)} - y^{(2)})(x^{(1)} - y^{(1)}) + a(x^{(2)} - y^{(2)})^2 \\ &= \left(a - \frac{1}{2}\right) \|x - y\|^2 + \frac{1}{2}(x^{(2)} - y^{(2)} + x^{(1)} - y^{(1)})^2 \geq \left(a - \frac{1}{2}\right) \|x - y\|^2, \end{aligned}$$

and therefore $\lambda = -(a - 1/2)$ and since $\kappa = a$, we have $\eta = 1/2$. Now checking monotonicity according to the zeroth-order definition, we have

$$\begin{aligned} & \sum_{i=1,2} F^{(i)}(y^{(i)}, y^{(-i)}) - F^{(i)}(x^{(i)}, y^{(-i)}) + F^{(i)}(x^{(i)}, x^{(-i)}) - F^{(i)}(y^{(i)}, x^{(-i)}) \\ &= \frac{1}{2} \|x - y\|^2 - \frac{1}{2}(x^{(1)} - y^{(1)} + x^{(2)} - y^{(2)})^2 \leq \frac{1}{2} \|x - y\|^2, \end{aligned}$$

which satisfies (4.37) with $\eta = 1/2$. Note that $\kappa = a$ must be large enough for $\lambda < 0$ to obtain contraction; this reflects that the strength of the convexity must be large enough to overcome the effects of the coupling. Additionally, any value of $\lambda \geq -a + 1/2$ maintains the first-order monotonicity inequality, in the same way that a κ -convex function is also $\kappa - \varepsilon$ convex, for any $\varepsilon \geq 0$.

4.6.2 Multispecies Coupled Wasserstein-2 Gradient Flows

In this section, we illustrate how the zeroth-order definition of monotonicity implies that proposed in [64], and provide a specific choice of energy functionals to which the results can be applied.

Lemma 4.57 (Wasserstein-2 Monotonicity). *Let $X^{(i)} = \mathcal{P}_2(\mathbb{R}^{d_i})$, $d_{X^{(i)}} = \mathcal{W}_2$, and let each $F^{(i)}$ be smooth enough so that $\nabla_{x^{(i)}} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](x^{(i)})$ is weakly well-posed. If $(F^{(i)})$ are $-\lambda$ -monotone with respect to geodesics $\mathcal{G}_2$ according to Definition 4.25, then the monotonicity condition in [64] holds,*

$$\begin{aligned} \sum_{i=1}^N \int \langle \nabla_{z_i} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](z_i) - \nabla_{v_i} \delta_{\mu^{(i)}} F^{(i)}[\mu^{(i)}, \mu^{(-i)}](v_i), z_i - v_i \rangle d\gamma^{(i)}(z_i, v_i) \\ \geq -\lambda \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}), \end{aligned}$$

where $\gamma^{(i)} \in \Gamma^*(\rho^{(i)}, \mu^{(i)})$ are Wasserstein-2 optimal plans with marginals $\rho^{(i)}$ and $\mu^{(i)}$.

Proof. Using the κ -convexity of $y \mapsto b(y, x)$ along geodesics $\mathcal{G}_2$, using (μ, ρ) and (ρ, ρ) as comparison points in the convexity inequality,

$$\begin{aligned} \sum_{i=1}^N F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}, \mu^{(-i)}) &\geq \sum_{i=1}^N \left(F^{(i)}(\rho^{(i)}, \rho^{(-i)}) - F^{(i)}(\rho^{(i)}, \rho^{(-i)}) \right. \\ &\quad \left. + \sum_{j=1}^N \int \delta_{\mu^{(j)}} [F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}, \mu^{(-i)})] \Big|_{\mu=\rho} (\mu^{(j)} - \rho^{(j)}) + \frac{\kappa}{2} \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}) \right) \end{aligned}$$

which can be rewritten, with enough smoothness in $(F^{(i)})$, using the optimal maps $\gamma^{(j)} \in \Gamma^*(\rho^{(j)}, \mu^{(j)})$, as

$$\begin{aligned} \sum_{i=1}^N F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}, \mu^{(-i)}) &\geq \sum_{i=1}^N \left(\frac{\kappa}{2} \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}) \right. \\ &\quad \left. + \sum_{j=1}^N \int \langle \nabla_{z_j} \delta_{\mu^{(j)}} (F^{(i)}[\mu^{(i)}, \mu^{(-i)}] - F^{(i)}[\rho^{(i)}, \mu^{(-i)}]) (z_j) \Big|_{\mu=\rho}, v_j - z_j \rangle d\gamma^{(j)}(z_j, v_j) \right) \end{aligned}$$

The right hand side, using that the j^{th} derivatives cancel for all $j \neq i$, reduces to give

$$\sum_{i=1}^N F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}, \mu^{(-i)}) \geq \sum_{i=1}^N \left(\frac{\kappa}{2} \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}) \right. \quad (4.83)$$

$$\left. - \int \langle \nabla_{z_i} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](z_i), z_i - v_i \rangle d\gamma^{(i)}(z_i, v_i) \right) \quad (4.84)$$

Now using the convexity inequality again with test points (ρ, μ) and (μ, μ) and summing,

$$\begin{aligned} \sum_{i=1}^N F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}, \mu^{(-i)}) + F^{(i)}(\rho^{(i)}, \rho^{(-i)}) - F^{(i)}(\mu^{(i)}, \rho^{(-i)}) &\geq \sum_{i=1}^N \left(\kappa \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}) \right. \\ &\quad \left. - \int \langle \nabla_{z_i} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](z_i) - \nabla_{v_i} \delta_{\mu^{(i)}} F^{(i)}[\mu^{(i)}, \mu^{(-i)}](v_i), z_i - v_i \rangle d\gamma^{(i)}(z_i, v_i) \right). \end{aligned}$$

From the η -dissipativity of $b(y, x)$, we upper-bound the left-hand side of the inequality by $\eta \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)})$ to obtain

$$\begin{aligned} & (\eta - \kappa) \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}) \\ & \geq \sum_{i=1}^N \left(- \int \langle \nabla_{z_i} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](z_i) - \nabla_{v_i} \delta_{\mu^{(i)}} F^{(i)}[\mu^{(i)}, \mu^{(-i)}](v_i), z_i - v_i \rangle d\gamma^{(i)}(z_i, v_i) \right), \end{aligned}$$

and inverting the inequality results in exactly $-\lambda = \kappa - \eta$ monotonicity of $(F^{(i)})$ from [64],

$$\begin{aligned} & \sum_{i=1}^N \int \langle \nabla_{z_i} \delta_{\rho^{(i)}} F^{(i)}[\rho^{(i)}, \rho^{(-i)}](z_i) - \nabla_{v_i} \delta_{\mu^{(i)}} F^{(i)}[\mu^{(i)}, \mu^{(-i)}](v_i), z_i - v_i \rangle d\gamma^{(i)}(z_i, v_i) \\ & \geq -\lambda \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}, \mu^{(i)}). \end{aligned}$$

□

In the next result, we show formally that multispecies gradient flow dynamics of the form $\frac{d}{dt}\rho^{(i)} = \operatorname{div} \left(\rho^{(i)} \nabla_{x^{(i)}} \delta_{\rho^{(i)}} F^{(i)}[\rho] \right)$ for all $i \in [1, \dots, N]$ satisfy $(\operatorname{EVI}_{2,\kappa})$.

Lemma 4.58 (Multispecies Wasserstein-2 Gradient Flow Satisfies EVI (Formal)). *Consider dynamics given by*

$$\frac{d}{dt}\rho^{(i)} = \operatorname{div} \left(\rho^{(i)} \nabla_{x^{(i)}} \delta_{\rho^{(i)}} F^{(i)}(\rho) \right) \quad \text{for all } i \in [1, \dots, N]. \quad (4.85)$$

Let $(F^{(i)})$ be $-\lambda$ -monotone according to Definition 4.25 with respect to geodesics. Any (weak) solution to (4.85) satisfies the EVI, for all $\mu \in \mathbf{D}$, and all $t \geq 0$,

$$\frac{1}{2} \frac{d}{dt} \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}(t), \mu^{(i)}) + \frac{\kappa}{2} \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}(t), \mu^{(i)}) \leq \sum_{i=1}^N F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}(t), \mu^{(-i)}),$$

where $\mathbf{D}$ is the joint domain of $(F^{(i)})$.

Proof. Since $\mu \mapsto b(\mu, \rho(t))$ is κ -convex with respect to geodesics we can use again (4.83), for all $\mu \in \mathbf{D}$,

$$\begin{aligned} & \frac{\kappa}{2} \sum_{i=1}^N \mathcal{W}_2^2(\rho^{(i)}(t), \mu^{(i)}) + \sum_{i=1}^N \int \langle \nabla_{x^{(i)}} \delta_{\rho^{(i)}} F^{(i)}(v^{(i)}), z^{(i)} - v^{(i)} \rangle d\gamma^{(i)}(t, z^{(i)}, v^{(i)}) \\ & \leq \sum_{i=1}^N F^{(i)}(\rho^{(i)}(t), \rho^{(-i)}(t)) - F^{(i)}(\rho^{(i)}(t), \rho^{(-i)}(t)) - F^{(i)}(\mu^{(i)}, \mu^{(-i)}) - F^{(i)}(\rho^{(i)}(t), \mu^{(-i)}), \end{aligned}$$

where $\gamma^{(i)}(t) \in \Gamma^*(\rho^{(i)}(t), \mu^{(i)})$. Identifying

$$\frac{1}{2} \frac{d}{dt} \mathcal{W}_2^2(\rho^{(i)}(t), \mu^{(i)}) = \langle \nabla_{x^{(i)}} \delta_{\rho^{(i)}} F^{(i)}(v^{(i)}), z^{(i)} - v^{(i)} \rangle d\gamma^{(i)}(t, z^{(i)}, v^{(i)}),$$

using [9, Lemma 8.4.7] allows us to conclude the EVI. $\square$

The following example illustrates how Definition 4.25 can be applied to a specific choice of coupling between two species in the Wasserstein-2 space.

Example 4.59 (Attractive-Repulsive Kernels). *Let $\hat{F}^{(1)} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ and $\hat{F}^{(2)} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ be $\kappa_1 > 0$ and $\kappa_2 > 0$ displacement convex. Consider the setting with a nonlocal interaction kernel $W \in C^2(\mathbb{R}^d, \mathbb{R})$ that is a symmetric function $W(x) = W(-x)$, and energy functionals*

$$F^{(1)}(\rho^{(1)}, \rho^{(2)}) = \hat{F}^{(1)}(\rho^{(1)}) + \iint W(x_1 - x_2) d\rho^{(1)}(x_1) d\rho^{(2)}(x_2), \quad F^{(2)}(\rho) = \hat{F}^{(2)}(\rho^{(2)}).$$

First we check the η interaction dissipativity inequality from Definition 4.25. Only the cross-interaction term coupling the two species appears:

$$\begin{aligned} & \iint [W(y_1 - y_2) - W(x_1 - y_2) \\ & \quad + W(x_1 - x_2) - W(y_1 - x_2)] d\rho^{(1)}(x_1) d\mu^{(1)}(y_1) d\rho^{(2)}(x_2) d\mu^{(2)}(y_2) \\ & = \iint \int_0^1 \int_0^1 \langle x_1 - y_1, \nabla^2 W(I(s, \tau, x, y)) \cdot (x_2 - y_2) \rangle ds d\tau d\gamma^{(1)}(x_1, y_1) d\gamma^{(2)}(x_2, y_2) \end{aligned}$$

where $I(s, \tau, x, y) := (1-s)x_1 + sy_1 - (1-\tau)x_2 - \tau y_2$ and $\gamma^{(1)} \in \Gamma^*(\rho^{(1)}, \mu^{(1)})$, $\gamma^{(2)} \in \Gamma^*(\rho^{(2)}, \mu^{(2)})$. The η interaction dissipativity inequality holds if

$$\int_0^1 \int_0^1 \langle x_1 - y_1, \nabla^2 W(I(s, \tau, x, y)) \cdot (x_2 - y_2) \rangle ds d\tau \leq \eta \|x - y\|^2.$$

For example, consider the case where $W = W^{(m)}$ is an attractive-repulsive kernel, with $l_a > l_r > 0$, $C_r > C_a > 0$:

$$W^{(m)}(x) = C_r \exp\left(-\frac{\|x\|^2}{l_r}\right) - C_a \exp\left(-\frac{\|x\|^2}{l_a}\right).$$

This kernel has been studied in a number of applications; see [15, 67] and references therein. The Hessian of $W^{(m)}$ is uniformly bounded in x because

$$\begin{aligned} \nabla^2 W^{(m)} &= 2I_{d^{(i)}} \left(-\frac{C_r}{l_r} \exp(-\|x\|^2/l_r) + \frac{C_a}{l_a} \exp(-\|x\|^2/l_a) \right) \\ &\quad + 4xx^\top \left(\frac{C_r}{l_r^2} \exp(-\|x\|^2/l_r) - \frac{C_a}{l_a^2} \exp(-\|x\|^2/l_a) \right). \end{aligned}$$

If κ_1, κ_2 are large enough, then $\lambda = \eta - \kappa = \eta - \min\{\kappa_1, \kappa_2\}$ can be less than zero.

4.A Proof of Cauchy Estimate

In this section we present the proof of the Cauchy estimate, Theorem 4.41.

Proof of Theorem 4.41. We now introduce the functions

$$\ell(t) := t - n \quad \text{if } n < t \leq n + 1, \quad \ell_\tau(t) := \ell(t/\tau), \quad \ell_\tau(0) := 0,$$

and for a discrete solution $X_{\tau_1}^n$

$$\zeta_{\tau_1}(s, Y) := (1 - \ell_{\tau_1}(s))d_X(\underline{X}_{\tau_1}(s), Y) + \ell_{\tau_1}(s)d_X(\bar{X}_{\tau_1}(s), Y)$$

Similarly, for a discrete solution $Y_{\tau_2}^n$ we set

$$\omega_{\tau_2}(X, t) := (1 - \ell_{\tau_2}(t))d_X(\underline{Y}_{\tau_2}(t), X) + \ell_{\tau_2}(t)d_X(\bar{Y}_{\tau_2}(t), X).$$

We immediately check from Lemma 4.39 that

$$\partial_s \zeta_{\tau_1}(s, Y) \leq b_1^K(Y, \bar{X}_{\tau_1}(s)), \quad \partial_t \omega_{\tau_2}(X, t) \leq b_1^K(X, \bar{Y}_{\tau_2}(t)), \quad \text{for a.e. } t, s > 0.$$

Choosing $Y := \bar{Y}_{\tau_2}(t)$ in the first inequality and $X := \bar{X}_{\tau_1}(s)$ in the second one, and extending $\bar{X}_{\tau_1}(s) \equiv X_{\tau_1}^0$ for $s < 0$ we get

$$\partial_s \zeta_{\tau_1}(s, \bar{Y}_{\tau_2}(t)) + \partial_t \omega_{\tau_2}(\bar{X}_{\tau_1}(s), t) \leq f(s, t), \quad (4.86)$$

where

$$f(s, t) = \begin{cases} b_1^K(X_{\tau_1}^0, \bar{Y}_{\tau_2}(t)) & \text{if } s < 0, \\ \lambda d_X(\bar{X}_{\tau_1}(s), \bar{Y}_{\tau_2}(t)) & \text{otherwise} \end{cases}, \quad t \geq 0. \quad (4.87)$$

We will integrate over the strip $Q_{0,T}^\varepsilon$ as in [175]:


$$Q_{0,T}^\varepsilon := \{(s, t) \in \mathbb{R}^2 : 0 \leq t \leq T, \quad t - \varepsilon \leq s \leq t\}.$$

The integral of (4.86) over $Q_{0,T}^\varepsilon$,

$$\iint_{Q_{0,T}^\varepsilon} \operatorname{div}_{s,t} \left(\begin{bmatrix} \zeta_{\tau_1}(s, \bar{Y}_{\tau_2}(t)) \\ \omega_{\tau_2}(\bar{X}_{\tau_1}(s), t) \end{bmatrix} \right) ds dt \leq \iint_{Q_{0,T}^\varepsilon} f(s, t) ds dt := a_0, \quad (4.88)$$

will require five estimates; one from the right hand side, and four from the divergence theorem. In Figure 4.3, the edges of $Q_{0,T}^\varepsilon$ are labeled with the corresponding boundary integrals. Applying the divergence theorem to the lower bound, let

$$\begin{aligned} \iint_{Q_{0,T}^\varepsilon} \operatorname{div}_{s,t} \left(\begin{bmatrix} \zeta_{\tau_1}(s, \bar{Y}_{\tau_2}(t)) \\ \omega_{\tau_2}(\bar{X}_{\tau_1}(s), t) \end{bmatrix} \right) ds dt &= \sum_{i=1}^4 a_i, \quad a_1 := \int_{T-\varepsilon}^T \omega_{\tau_2}(\bar{X}_{\tau_1}(s), T) ds, \\ a_2 &:= - \int_{-\varepsilon}^0 \omega_{\tau_2}(\bar{X}_{\tau_1}(s), 0) ds, \quad a_3 := - \frac{1}{\sqrt{2}} \int_0^T \left(\omega_{\tau_2}(\bar{X}_{\tau_1}(t), t) - \zeta_{\tau_1}(t, \bar{Y}_{\tau_2}(t)) \right) dt, \\ a_4 &:= - \frac{1}{\sqrt{2}} \int_0^T \left(\zeta_{\tau_1}(t - \varepsilon, \bar{Y}_{\tau_2}(t)) - \omega_{\tau_2}(\bar{X}_{\tau_1}(t - \varepsilon), t) \right) dt. \end{aligned}$$

Figure 4.3: The strip $Q_{0,T}^\epsilon$.

Therefore (4.88) can be written as

$$\sum_{i=1}^4 a_i \leq a_0. \quad (4.89)$$

We compute estimates for $a_0, a_1, \dots, a_4$ to obtain the Cauchy estimate. We will begin with a_0 . Under Definition 4.28 with $x \in D^\kappa(\Delta)$, we have

$$\Delta^\kappa(x) := \sup_{y' \in D} \frac{b^\kappa(y', x)_-}{d_X(x, y')} \geq -\frac{b(y, x)}{d_X(x, y)} + \frac{\kappa}{2} d_X(x, y) = -b_1^\kappa(y, x) \quad \text{for all } y \in D. \quad (4.90)$$

Combining this estimate with (4.61), we construct the following bound on $b_1^\kappa(X_{\tau_1}^0, \bar{Y}_{\tau_2}(t))$

$$b_1^\kappa(X_{\tau_1}^0, \bar{Y}_{\tau_2}(t)) \leq \Delta^\kappa(X_{\tau_1}^0) + \lambda d_X(X_{\tau_1}^0, \bar{Y}_{\tau_2}(t)).$$

Applying this to the definition of a_0 ,

$$\begin{aligned} a_0 &\leq \int_0^\epsilon \int_{t-\epsilon}^t \Delta^\kappa(X_{\tau_1}^0) ds dt + \tilde{\eta} \iint_{Q_{0,T}^\epsilon} d_X(\bar{X}_{\tau_1}(s), \bar{Y}_{\tau_2}(t)) ds dt \\ &= \frac{\epsilon^2}{2} \Delta^\kappa(X_{\tau_1}^0) + \lambda \iint_{Q_{0,T}^\epsilon} d_X(\bar{X}_{\tau_1}(s), \bar{Y}_{\tau_2}(t)) ds dt. \end{aligned}$$

The integral over the distance term has the upper bound

$$\begin{aligned} \iint_Q d_X(\bar{X}_{\tau_1}(t), \bar{Y}_{\tau_2}(s)) ds dt &\leq \varepsilon T d_X(\bar{X}_{\tau_1}(0), \bar{Y}_{\tau_2}(0)) + \iint_{Q_{0,T}^\varepsilon} d_X(\bar{X}_{\tau_1}(0), \bar{X}_{\tau_1}(t)) ds dt \\ &\quad + \iint_{Q_{0,T}^\varepsilon} d_X(\bar{Y}_{\tau_2}(0), \bar{Y}_{\tau_2}(s)) ds dt \\ &\leq \varepsilon T d_X(\bar{X}_{\tau_1}(0), \bar{Y}_{\tau_2}(0)) + \varepsilon T \sum_{i=0}^{\lfloor T/\tau_1 \rfloor} d_X(X_{\tau_1}^i, X_{\tau_1}^{i+1}) + \varepsilon T \sum_{i=0}^{\lfloor T/\tau_2 \rfloor} d_X(Y_{\tau_2}^i, Y_{\tau_2}^{i+1}). \end{aligned}$$

The summations are upper-bounded, using (4.69),

$$\begin{aligned} \sum_{i=0}^{\lfloor T/\tau_1 \rfloor} d_X(X_{\tau_1}^i, X_{\tau_1}^{i+1}) &\leq \sum_{i=0}^{\lfloor T/\tau_1 \rfloor} \tilde{\tau}_1 \Delta^\kappa(X_{\tau_1}^i) \leq (\lfloor T/\tau_1 \rfloor + 1) \tilde{\tau}_1 c(\tau_1, T) \Delta^\kappa(X_{\tau_1}^0) \leq \frac{T + \tau_0}{1 - \lambda \tau_0} c(\tau_1, T) \Delta^\kappa(X_{\tau_1}^0) \\ \sum_{i=0}^{\lfloor T/\tau_2 \rfloor} d_X(Y_{\tau_2}^i, Y_{\tau_2}^{i+1}) &\leq \frac{T + \tau_0}{1 - \lambda \tau_0} c(\tau_2, T) \Delta^\kappa(Y_{\tau_2}^0). \end{aligned}$$

Therefore, defining $\hat{c}(\tau_0, T) := \frac{T + \tau_0}{1 - \lambda \tau_0} c(\tau_0, T)$,

$$\iint_{Q_{0,T}^\varepsilon} d_X(\bar{X}_{\tau_1}(t), \bar{Y}_{\tau_2}(s)) ds dt \leq \varepsilon T \left(d_X(\bar{X}_{\tau_1}(0), \bar{Y}_{\tau_2}(0)) + \hat{c}(\tau_1, T) \Delta^\kappa(X_{\tau_1}^0) + \hat{c}(\tau_2, T) \Delta^\kappa(Y_{\tau_2}^0) \right).$$

and

$$a_0 \leq \frac{\varepsilon^2}{2} \Delta^\kappa((X_{\tau_1}^0) + \lambda \varepsilon T \left(d_X(\bar{X}_{\tau_1}(0), \bar{Y}_{\tau_2}(0)) + \hat{c}(\tau_1, T) \Delta^\kappa(X_{\tau_1}^0) + \hat{c}(\tau_2, T) \Delta^\kappa(Y_{\tau_2}^0) \right) \quad (4.91)$$

Next, a_1 can be lower-bounded using the triangle inequality to obtain

$$\begin{aligned} \omega_{\tau_2}(\bar{X}_{\tau_1}(s), T) &\leq (1 - \ell_{\tau_2}(T)) \left[d_X(\underline{Y}_{\tau_2}(T), \bar{X}_{\tau_1}(s)) + d_X(\bar{X}_{\tau_1}(s), \bar{X}_{\tau_1}(T)) \right] \\ &\quad + \ell_{\tau_2}(T) \left[d_X(\bar{Y}_{\tau_2}(T), \bar{X}_{\tau_1}(s)) + d_X(\bar{X}_{\tau_1}(s), \bar{X}_{\tau_1}(T)) \right] \\ &= \omega_{\tau_2}(\bar{X}_{\tau_1}(s), T) + d_X(\bar{X}_{\tau_1}(s), \bar{X}_{\tau_1}(T)). \end{aligned}$$

Again using the triangle inequality twice, we have

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq d_X(\bar{X}_{\tau_1}(s), \bar{Y}_{\tau_2}(T)) + d_X(\bar{X}_{\tau_1}(T), \bar{X}_{\tau_1}(s)) \\ &\leq \ell_{\tau_2}(T) d_X(\bar{X}_{\tau_1}(s), \bar{Y}_{\tau_2}(T)) + (1 - \ell_{\tau_2}(T)) \left(d_X(\bar{X}_{\tau_1}(s), \underline{Y}_{\tau_2}(T)) \right. \\ &\quad \left. + d_X(\underline{Y}_{\tau_2}(T), \bar{Y}_{\tau_2}(T)) \right) + d_X(\bar{X}_{\tau_1}(T), \bar{X}_{\tau_1}(s)) \\ &\leq \omega_{\tau_2}(\bar{X}_{\tau_1}(s), T) + d_X(\bar{X}_{\tau_1}(T), \bar{X}_{\tau_1}(s)) + d_X(\bar{Y}_{\tau_2}(T), \underline{Y}_{\tau_2}(T)). \end{aligned}$$

We will integrate both sides of the above inequality from $T - \varepsilon$ to ε . We will use that

$$\int_{T-\varepsilon}^T d_X(\bar{X}_{\tau_1}(T), \bar{X}_{\tau_1}(s)) ds \leq \varepsilon \sup_{i: \tau_1 i \in [T-\varepsilon, T]} d_X(X_{\tau_1}^i, X_{\tau_1}^{i-1}) \leq \varepsilon c(\tau_1, T) \tilde{\tau}_1 \Delta^\kappa(X_{\tau_1}^0),$$

in combination with $\tilde{\tau}_i c(\tau_i, T) \leq \hat{c}(\tau_i, T)$, to obtain

$$\varepsilon(d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) - \hat{c}(\tau_1, T)\Delta^\kappa(X_{\tau_1}^0) - \hat{c}(\tau_2, T)\Delta^\kappa(Y_{\tau_2}^0)) \leq \int_{T-\varepsilon}^T \omega_{\tau_2}(\bar{X}_{\tau_1}(s), T) ds =: a_1.$$

Next, for a_2

$$a_2 = - \int_{-\varepsilon}^0 \omega_{\tau_2}(\bar{X}_{\tau_1}(s), 0) ds = - \int_{-\varepsilon}^0 d_X(\bar{Y}_{\tau_2}(0), \bar{X}_{\tau_2}(0)) ds = -\varepsilon d_X(X_{\tau_1}^0, Y_{\tau_2}^0). \quad (4.92)$$

Next we will bound a_3 . We set

$$\delta_{X,\tau_1}(s) := \sup_{Y \in D} \left| \zeta_{\tau_1}(s, Y) - d_X(\bar{X}_{\tau_1}(s), Y) \right|, \quad \delta_{Y,\tau_2}(t) := \sup_{X \in D} \left| \omega_{\tau_2}(X, t) - d_X(\bar{Y}_{\tau_2}(t), X) \right| \quad (4.93)$$

observing that

$$\begin{aligned} \delta_{Y,\tau_2}(t) &= \sup_Y \left| (1 - \ell_{\tau_2}(t)) d_X(\underline{Y}_{\tau_2}(t), X) + \ell_{\tau_2}(t) d_X(\bar{Y}_{\tau_2}(t), X) - d_X(\bar{Y}_{\tau_2}(t), X) \right| \\ &= \sup_Y \left| (1 - \ell_{\tau_2}(t)) [d_X(\underline{Y}_{\tau_2}(t), X) - d_X(\bar{Y}_{\tau_2}(t), X)] + \ell_{\tau_2}(t) [d_X(\bar{Y}_{\tau_2}(t), X) - d_X(\bar{Y}_{\tau_2}(t), X)] \right| \\ &= (1 - \ell_{\tau_2}(t)) \sup_Y \left| d_X(\underline{Y}_{\tau_2}(t), X) - d_X(\bar{Y}_{\tau_2}(t), X) \right| \leq (1 - \ell_{\tau_2}(t)) d_X(\underline{Y}_{\tau_2}(t), \bar{Y}_{\tau_2}(t)), \end{aligned}$$

and from (4.69)

$$\begin{aligned} \delta_{X,\tau_1}(s) &\leq c(\tau_1, T) \left(\frac{1}{\tau_1} - \lambda \right)^{-1} (1 - \ell_{\tau_1}(s)) \Delta^\kappa(X_{\tau_1}^0), \\ \delta_{Y,\tau_2}(t) &\leq c(\tau_2, T) \left(\frac{1}{\tau_2} - \lambda \right)^{-1} (1 - \ell_{\tau_2}(t)) \Delta^\kappa(Y_{\tau_2}^0). \end{aligned} \quad (4.94)$$

Recall $\tilde{\tau}_i := \left(\frac{1}{\tau_i} - \lambda \right)^{-1}$ and let

$$I_{X,\tau_1} := \int_0^T \delta_{X,\tau_1}(s) ds \leq c(\tau_1, T) \frac{\tilde{\tau}_1}{2} h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0), \quad h_\tau(T) = T + \frac{\tau}{2} \ell_\tau(T) (1 - \ell_\tau(T)), \quad (4.95)$$

$$I_{Y,\tau_2} := \int_0^T \delta_{Y,\tau_2}(s) ds \leq c(\tau_2, T) \frac{\tilde{\tau}_2}{2} h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0) \quad (4.96)$$

Therefore

$$\begin{aligned} a_3 &= -\frac{1}{\sqrt{2}} \int_0^T \left(\omega_{\tau_2}(\bar{X}_{\tau_1}(t), t) - d_X(\bar{X}_{\tau_1}(t), \bar{Y}_{\tau_2}(t)) + d_X(\bar{X}_{\tau_1}(t), \bar{Y}_{\tau_2}(t)) - \zeta_{\tau_1}(t, \bar{Y}_{\tau_2}(t)) \right) dt \\ &\geq -\frac{1}{\sqrt{2}} \int_0^T \left(\delta_{X,\tau_1}(t) + \delta_{Y,\tau_2}(t) \right) dt = -\frac{1}{\sqrt{2}} (I_{X,\tau_1} + I_{Y,\tau_2}). \end{aligned}$$

A similar calculation for a_4 gives

$$\begin{aligned} a_4 &= -\frac{1}{\sqrt{2}} \int_0^T \left(\zeta_{\tau_1}(t - \varepsilon, \bar{Y}_{\tau_2}(t)) - d_X(\bar{X}_{\tau_1}(t - \varepsilon), \bar{Y}_{\tau_2}(t)) + d_X(\bar{X}_{\tau_1}(t - \varepsilon), \bar{Y}_{\tau_2}(t)) \right. \\ &\quad \left. - \omega_{\tau_2}(\bar{X}_{\tau_1}(t - \varepsilon), t) \right) dt \geq -\frac{1}{\sqrt{2}} \int_0^T \left(\delta_{X,\tau_1}(t - \varepsilon) + \delta_{Y,\tau_2}(t) \right) dt = -\frac{1}{\sqrt{2}} (I_{X,\tau_1} + I_{Y,\tau_2}). \end{aligned}$$

Combining the estimates for the five terms results in

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq (1 + \lambda T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) + \frac{\varepsilon}{2} \Delta^\kappa(X_{\tau_1}^0) + \hat{c}(\tau_1, T)(1 + \lambda T) \Delta^\kappa(X_{\tau_1}^0) \\ &\quad + \hat{c}(\tau_2, T)(1 + \lambda T) \Delta^\kappa(Y_{\tau_2}^0) + \hat{c}(\tau_1, T) \frac{1}{\sqrt{2}\varepsilon} h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0) + \hat{c}(\tau_2, T) \frac{1}{\sqrt{2}\varepsilon} h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0), \end{aligned}$$

and selecting $\tau_0 = \max\{\tau_1, \tau_2\}$,

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq (1 + \lambda T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) + \frac{\varepsilon}{2} \Delta^\kappa(X_{\tau_1}^0) + \hat{c}(\tau_0, T)(1 + \lambda T)(\Delta^\kappa(X_{\tau_1}^0) + \Delta^\kappa(Y_{\tau_2}^0)) \\ &\quad + \frac{\hat{c}(\tau_0, T)}{\sqrt{2}\varepsilon} \left(h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0) + h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0) \right), \end{aligned}$$

and optimizing w.r.t. ε ,

$$\varepsilon = 2^{3/4} \left(\hat{c}(\tau_0, T) \left(h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0) + h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0) \right) \right)^{1/2} \left(\frac{1}{\Delta^\kappa(X_{\tau_1}^0)} \right)^{1/2}.$$

With this value of ε ,

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq (1 + \lambda T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) \\ &\quad + 2^{-1/4} \left(\Delta^\kappa(X_{\tau_1}^0) \hat{c}(\tau_0, T) \left(h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0) + h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0) \right) \right)^{1/2} \\ &\quad + \hat{c}(\tau_0, T)(1 + \lambda T)(\Delta^\kappa(X_{\tau_1}^0) + \Delta^\kappa(Y_{\tau_2}^0)) \\ &\quad + \frac{2^{-3/4}}{\sqrt{2}} \left(\hat{c}(\tau_0, T) \left(h_{\tau_1}(T) \Delta^\kappa(X_{\tau_1}^0) + h_{\tau_2}(T) \Delta^\kappa(Y_{\tau_2}^0) \right) \right)^{1/2} \left(\Delta^\kappa(X_{\tau_1}^0) \right)^{1/2}. \end{aligned}$$

Adding positive terms dependent on $\Delta^\kappa(Y_{\tau_u}^0)$ inside the two square roots,

$$\begin{aligned} d_X(\bar{X}_{\tau_1}(T), \bar{Y}_{\tau_2}(T)) &\leq (1 + \lambda T) d_X(X_{\tau_1}^0, Y_{\tau_2}^0) \\ &\quad + \left[\hat{c}(\tau_0, T)(1 + \lambda T) + \frac{3}{2^{5/4}} \sqrt{\hat{c}(\tau_0, T) h_{\tau_1}(T) + h_{\tau_2}(T)} \right] (\Delta^\kappa(X_{\tau_1}^0) + \Delta^\kappa(Y_{\tau_2}^0)) \end{aligned}$$

□

*Chapter 5***STRATEGIC DISTRIBUTION SHIFT OF INTERACTING AGENTS VIA
COUPLED GRADIENT FLOWS**

L. Conger, F. Hoffmann, E. Mazumdar, and L. Ratliff. Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows. en. *Advances in Neural Information Processing Systems*, 36:45971–46006, Dec. 2023

ABSTRACT

We propose a novel framework for analyzing the dynamics of distribution shift in real-world systems that captures the feedback loop between learning algorithms and the distributions on which they are deployed. Prior work largely models feedback-induced distribution shift as adversarial or via an overly simplistic distribution-shift structure. In contrast, we propose a coupled partial differential equation model that captures fine-grained changes in the distribution over time by accounting for complex dynamics that arise due to strategic responses to algorithmic decision-making, non-local endogenous population interactions, and other exogenous sources of distribution shift. We consider two common settings in machine learning: cooperative settings with information asymmetries, and competitive settings where a learner faces strategic users. For both of these settings, when the algorithm retrain via gradient descent, we prove asymptotic convergence of the retraining procedure to a steady-state, both in finite and in infinite dimensions, obtaining explicit rates in terms of the model parameters. To do so we derive new results on the convergence of coupled PDEs that extends what is known on multi-species systems. Empirically, we show that our approach captures well-documented forms of distribution shifts like polarization and disparate impacts that simpler models cannot capture.

5.1 Introduction

In many machine learning tasks, there are commonly sources of exogenous and endogenous distribution shift, necessitating that the algorithm be retrained repeatedly over time. Some of these shifts occur without the influence of an algorithm; for example, individuals influence each other to become more or less similar in their attributes, or benign forms of distributional shift occur [185]. Other shifts, however, are in response to algorithmic decision-making. Indeed, the very use of a decision-making algorithm can incentivize individuals to change or mis-report their data to achieve desired outcomes—a phenomenon known in economics as Goodhart’s law. Such phenomena have been empirically observed, a well-known example being in [42], where researchers observed a population in Columbia strategically mis-reporting data to game a poverty index score used for distributing government assistance. Works such as [169, 209], which investigate the effects of distribution shift over time on a machine learning algorithm, point toward the need for evaluating the robustness of algorithms to distribution shifts.

Many existing approaches for modeling distribution shift focus on simple metrics like optimizing over moments or covariates [83, 146, 25]. Other methods consider worst-case scenarios, as in distributionally robust optimization [1, 151, 93, 139]. However, when humans respond to algorithms, these techniques may not be sufficient to holistically capture the impact an algorithm has on a population. For example, an algorithm that takes into account shifts in a distribution’s mean might inadvertently drive polarization, rendering a portion of the population disadvantaged.

Motivated by the need for a more descriptive model, we present an alternative perspective which allows us to fully capture complex dynamics that might drive distribution shifts in real-world systems. Our approach is general enough to capture various sources of exogenous and endogenous distribution shift including the feedback loop between algorithms and data distributions studied in the literature on performative prediction [182, 128, 187, 172, 170], the strategic interactions studied in strategic classification [124, 90], and also endogenous factors like intra-population dynamics and distributional shifts. Indeed, while previous works have studied these phenomena in isolation, our method allows us to capture all of them as well as their interactions. For example, in [215], the authors investigate the effects of dynamics in strategic classification problems—but the model they analyze does not capture individual interactions in the population. In [128], the authors model the interaction between a population that repeatedly responds to algorithmic decision-making by shifting its mean. Additionally, [187] study settings in which the population has both exogenous and endogenous distribution shifts due to feedback, but much like the other cited work, the focus remains on average performance. Each of these works fails to account for diffusion or intra-population interactions that can result in important qualitative changes to the distribution.

Contributions. Our approach to this problem relies on a detailed non-local PDE model of the data

distribution which captures each of these factors. One term driving the evolution of the distribution over time captures the response of the population to the deployed algorithm, another draws on models used in the PDE literature for describing non-local effects and consensus in biological systems to model intra-population dynamics, and the last captures a background source of distribution shift. This is coupled with an ODE, lifted to a PDE, which describes the training of a machine learning algorithm results in a coupled PDE system which we analyze to better understand the behaviors that can arise among these interactions.

In one subcase, our model exhibits a joint gradient flow structure, where both PDEs can be written as gradients flows with respect to the same joint energy, but considering infinite dimensional gradients with respect to the different arguments. This mathematical structure provides powerful tools for analysis and has been an emerging area of study with a relatively small body of prior work, none of which related to distribution shifts in societal systems, and a general theory for multi-species gradient flows is still lacking. We give a brief overview of the models that are known to exhibit this joint gradient flow structure: in [81] the authors consider a two-species tumor model with coupling through Brinkman’s Law. A number of works consider coupling via convolution kernels [112, 118, 131, 49, 94, 91] and cross-diffusion [147, 7, 159], with applications in chemotaxis among other areas. In the models we consider here, the way the interaction between the two populations manifests is neither via cross-diffusion, nor via the non-local self-interaction term. A related type of coupling has recently appeared in [125], however in the setting of graphs. Recent work [89] provides particle-based methods to approximately compute the solution to a minimax problem where the optimization space is over measures; following that work, [207] provides another particle-based method using mirror descent-ascent to solve a similar problem. Other recent work [154] proves that a mean-field gradient ascent-descent scheme with an entropy annealing schedule converges to the solution of a minimax optimization problem with a timescale separation parameter that is also time-varying; in contrast, our work considers fixed timescale separation setting. [116] show that the mean-field description of a particle method for solving minimax problems has proveable convergence guarantees in the Wasserstein-Fisher-Rao metric. Each of these references considers an energy functional that is linear in the distribution of each species respectively; our energy includes nonlinearities in the distributions via a self-interaction term as well as diffusion for the population. Moreover, the above works introduce a gradient flow dynamic as a tool for obtaining and characterizing the corresponding steady states, whereas in our setting we seek to capture the time-varying behavior that models distributions shifts. In the other subcase, we prove exponential convergence in two competitive, timescale separated settings where the algorithm and strategic population have conflicting objectives.

We show numerically that retraining in a competitive setting leads to polarization in the population,

illustrating the importance of fine-grained modeling.

5.2 Problem Formulation

Machine learning algorithms that are deployed into the real world for decision-making often become part of complex feedback loops with the data distributions and data sources with which they interact. In an effort to model these interactions, consider a machine learning algorithm that has loss given by $L(z, x)$ where $x \in \mathbb{R}^d$ are the algorithm parameters and $z \in \mathbb{R}^d$ are the population attributes, and the goal is to solve

$$\operatorname{argmin}_{x \in \mathcal{X}} \mathbb{E}_{z \sim \rho} L(z, x),$$

where $\mathcal{X}$ is the class of model parameters and $\rho(z)$ is the population distribution. Individuals have an objective given by $J(z, x)$ in response to a model parameterized by x , and they seek to solve

$$\operatorname{argmin}_{z \in \mathbb{R}^d} J(z, x).$$

When individuals in the population and the algorithm have access to gradients, we model the optimization process as a gradient-descent-type process. Realistically, individuals in the population will have nonlocal information and influences, as well as external perturbations, the effects of which we seek to capture in addition to just minimization. To address this, we propose a partial differential equation (PDE) model for the population, that is able to capture nonlocal interactions between individuals on the level of a collective population. To analyse how the population evolves over time, a notion of derivative in infinite dimensions is needed. A natural, and in this context physically meaningful, way of measuring the dissipation mechanism for probability distributions is the Wasserstein-2 metric (see Definition 5.15). The following expression appears when computing the gradient of an energy functional with respect to the Wasserstein-2 topology.

Definition 5.1. *[First Variation] For a map $G : \mathcal{P}(\mathbb{R}^d) \mapsto \mathbb{R}$ and fixed probability distribution $\rho \in \mathcal{P}(\mathbb{R}^d)$, the first variation of G at the point ρ is denoted by $\delta_\rho G[\rho] : \mathbb{R}^d \rightarrow \mathbb{R}$, and is defined via the relation*

$$\int \delta_\rho G[\rho](z) \psi(z) dz = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (G(\rho + \epsilon \psi) - G(\rho))$$

for all $\psi \in C_c^\infty(\mathbb{R}^d)$ such that $\int d\psi = 0$, assuming that G is regular enough for all quantities to exist.

Here, $\mathcal{P}(\mathbb{R}^d)$ denotes the space of probability measures on the Borel sigma algebra. Using the first variation, we can express the gradient in Wasserstein-2 space, see for example [204, Exercise 8.8].

Lemma 5.2. *The gradient of an energy $G : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ in the Wasserstein-2 space is given by*

$$\nabla_{W_2} G(\rho) = -\operatorname{div} (\rho \nabla \delta_\rho G[\rho]) .$$

Here, $\mathcal{P}_2(\mathbb{R}^d)$ denotes the set of probability measures with bounded second moments, also see Appendix 5.A.2. As a consequence, the infinite dimensional steepest descent in Wasserstein-2 space can be expressed as the PDE

$$\partial_t \rho = -\nabla_{W_2} G(\rho) = \operatorname{div} (\rho \nabla \delta_\rho G[\rho]) . \quad (5.1)$$

All the coupled gradient flows considered in this work have this Wasserstein-2 structure. In particular, when considering that individuals minimize their own loss, we can capture these dynamics via a gradient flow in the Wasserstein-2 metric on the level of the distribution of the population. Then for given algorithm parameters $x \in \mathbb{R}^d$, the evolution for this strategic population is given by

$$\partial_t \rho = \operatorname{div} \left(\rho \nabla \delta_\rho \left[\mathbb{E}_{z \sim \rho} J(z, x) + E(\rho) \right] \right) , \quad (5.2)$$

where $E(\rho)$ is a functional including terms for internal influences and external perturbations. In real-world deployment of algorithms, decision makers update their algorithm over time, leading to an interaction between the two processes. We also consider the algorithm dynamics over time, which we model as

$$\dot{x} = -\nabla_x \left[\mathbb{E}_{z \sim \rho} L(z, x) \right] . \quad (5.3)$$

In this work, we analyze the behavior of the dynamics under the following model. The algorithm suffers a cost $f_1(z, x)$ for a data point z under model parameters x in the strategic population, and a cost $f_2(z, x)$ for a data point in a fixed, non-strategic population. The strategic population is denoted by $\rho \in \mathcal{P}$, and the non-strategic population by $\bar{\rho} \in \mathcal{P}$. The algorithm aims to minimize

$$\mathbb{E}_{z \sim \rho} L(z, x) = \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 ,$$

where the norm is the vector inner product $\|x\|^2 = \langle x, x \rangle$ and $\beta > 0$ weights the cost of updating the model parameters from its initial condition.

We consider two settings: (i) aligned objectives, and (ii) competing objectives. Case (i) captures the setting in which the strategic population minimization improves the performance of the algorithm, subject to a cost for deviating from a reference distribution $\bar{\rho} \in \mathcal{P}$. This cost stems from effort required to manipulate features, such as a loan applicant adding or closing credit cards. On the other hand, Case (ii) captures the setting in which the strategic population minimization worsens the performance of the algorithm, again incurring cost from distributional changes.

5.2.1 Case (i): Aligned Objectives

In this setting, we consider the case where the strategic population and the algorithm have aligned objectives. This occurs in examples such as recommendation systems, where users and algorithm designers both seek to develop accurate recommendations for the users. This corresponds to the population cost

$$\mathbb{E}_{z \sim \rho, x \sim \mu} J(z, x) = \iint f_1(z, x) d\rho(z) d\mu(x) + \alpha KL(\rho \mid \tilde{\rho}),$$

where $KL(\cdot \mid \cdot)$ denotes the Kullback-Leibler divergence. Note that the KL divergence introduces diffusion to the dynamics for ρ . The weight $\alpha > 0$ parameterizes the cost of distribution shift to the population. To account for nonlocal information and influence among members of the population, we include a kernel term $E(\rho) = \frac{1}{2} \int \rho W * \rho dz$, where $(W * \rho)(z) = \int W(z - \bar{z}) d\rho(\bar{z})$ is a convolution integral and W is a suitable interaction potential.

5.2.2 Case (ii): Competing Objectives

In settings such as online internet forums, where algorithms and users have used manipulative strategies for marketing [84], the strategic population may be incentivized to modify or mis-report their attributes. The algorithm has a competitive objective, in that it aims to maintain performance against a population whose dynamics cause the algorithm performance to suffer. When the strategic population seeks an outcome contrary to the algorithm, we model strategic population cost as

$$\mathbb{E}_{z \sim \rho, x \sim \mu} J(z, x) = - \iint f_1(z, x) d\rho(z) d\mu(x) + \alpha KL(\rho \mid \tilde{\rho}).$$

A significant factor in the dynamics for the strategic population is the timescale separation between the two "species"—i.e., the population and the algorithm. In our analysis, we will consider two cases: one, where the population responds much faster than the algorithm, and two, where the algorithm responds much faster than the population. We illustrate the intermediate case in a simulation example.

5.3 Results

We are interested in characterizing the long-time asymptotic behavior of the population distribution, as it depends on the decision-makers action over time. The structure of the population distribution gives us insights about how the decision-makers actions influences the entire population of users. For instance, as noted in the preceding sections, different behaviors such as bimodal distributions or large tails or variance might emerge, and such effects are not captured in simply looking at average performance. To understand this intricate interplay, one would like to characterize the behavior of both the population and the algorithm over large times. Our main contribution towards this goal is a novel analytical framework as well as analysis of the long-time asymptotics.

A key observation is that the dynamics in (5.2) and (5.3) can be re-formulated as a gradient flow; we lift x to a probability distribution μ by representing it as a Dirac delta μ sitting at the point x . As a result, the evolution of μ will be governed by a PDE, and combined with the PDE for the population, we obtain a system of coupled PDEs,

$$\begin{aligned}\partial_t \rho &= \operatorname{div} \left(\rho \nabla_z \delta_\rho \left[\mathbb{E}_{z \sim \rho, x \sim \mu} J(z, x) + E(\rho) \right] \right) \\ \partial_t \mu &= \operatorname{div} \left(\mu \nabla_x \delta_\mu \left[\mathbb{E}_{z \sim \rho, x \sim \mu} L(z, x) \right] \right),\end{aligned}$$

where δ_ρ and δ_μ are first variations with respect to ρ and μ according to Definition 5.1. The natural candidates for the asymptotic profiles of this coupled system are its steady states, which - thanks to the gradient flow structure - can be characterized as ground states of the corresponding energy functionals. In this work, we show existence and uniqueness of minimizers (maximizers) for the functionals under suitable conditions on the dynamics. We also provide criteria for convergence and explicit convergence rates. We begin with the case where the interests of the population and algorithm are aligned, and follow with analogous results in the competitive setting. We show convergence in energy, which in turn ensures convergence in a product Wasserstein metric. For convergence in energy, we use the notion of relative energy and prove that the relative energy converges to zero as time increases.

Definition 5.3 (Relative Energy). *The relative energy of a functional G is given by $G(\gamma|\gamma_\infty) = G(\gamma) - G(\gamma_\infty)$, where $G(\gamma_\infty)$ is the energy at the steady state.*

Since we consider the joint evolution of two probability distributions, we define a distance metric $\overline{W}$ on the product space of probability measures with bounded second moment.

Definition 5.4 (Joint Wasserstein Metric). *The metric over $\mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d)$ is called $\overline{W}$ and is given by*

$$\overline{W}((\rho, \mu), (\tilde{\rho}, \tilde{\mu}))^2 = W_2(\rho, \tilde{\rho})^2 + W_2(\mu, \tilde{\mu})^2$$

for all pairs $(\rho, \mu), (\tilde{\rho}, \tilde{\mu}) \in \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d)$, and where W_2 denotes the Wasserstein-2 metric (see Definition 5.15). We denote by $\overline{\mathcal{W}}(\mathbb{R}^d) := (\mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d), \overline{W})$ the corresponding metric space.

5.3.1 Gradient Flow Structure

In the case where the objectives of the algorithm and population are *aligned*, we can write the dynamics as a gradient flow by using the same energy functional for both species. Let

$G_a(\rho, \mu) : \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \mapsto [0, \infty]$ be the energy functional given by

$$G_a(\rho, \mu) = \iint f_1(z, x) d\rho(z) d\mu(x) + \iint f_2(z, x) d\bar{\rho}(z) d\mu(x) + \alpha KL(\rho|\bar{\rho}) + \frac{1}{2} \int \rho W * \rho + \frac{\beta}{2} \int \|x - x_0\|^2 d\mu(x).$$

This expression is well-defined as the relative entropy $KL(\rho|\bar{\rho})$ can be extended to the full set $\mathcal{P}(\mathbb{R}^d)$ by setting $G_a(\rho, \mu) = +\infty$ in case ρ is not absolutely continuous with respect to $\bar{\rho}$.

In the *competitive* case we define $G_c(\rho, x) : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \mapsto [-\infty, \infty]$ by

$$G_c(\rho, x) = \int f_1(z, x) d\rho(z) + \int f_2(x, z') d\bar{\rho}(z') - \alpha KL(\rho|\bar{\rho}) - \frac{1}{2} \int \rho W * \rho + \frac{\beta}{2} \|x - x_0\|^2.$$

In settings like recommender systems, the population and algorithm have aligned objectives; they seek to minimize the same cost but are subject to different dynamic constraints and influences, modeled by the regularizer and convolution terms. In the case where the objectives are aligned, the dynamics are given by

$$\begin{aligned} \partial_t \rho &= \operatorname{div} (\rho \nabla_z \delta_\rho G_a[\rho, \mu]) \\ \partial_t \mu &= \operatorname{div} (\mu \nabla_x \delta_\mu G_a[\rho, \mu]) . \end{aligned} \tag{5.4}$$

Note that (5.4) is a joint gradient flow, because the dynamics can be written in the form

$$\partial_t \gamma = \operatorname{div} (\gamma \nabla \delta_\gamma G_a(\gamma)) ,$$

where $\gamma = (\rho, \mu)$ and where the gradient and divergence are taken in both variables (z, x) . We discuss the structure of the dynamics (5.4) as well as the meaning of the different terms appearing in the energy functional G_a in Appendix 5.A.1.

In other settings, such as credit score reporting, the objectives of the population are competitive with respect to the algorithm. Here we consider two scenarios; one, where the algorithm responds quickly relative to the population, and two, where the population responds quickly relative to the algorithm. In the case where the algorithm can immediately adjust optimally (best-respond) to the distribution, the dynamics are given by

$$\begin{aligned} \partial_t \rho &= -\operatorname{div} (\rho (\nabla_z \delta_\rho G_c[\rho, x])|_{x=b(\rho)}) , \\ b(\rho) &:= \operatorname{argmin}_{\bar{x}} G_c(\rho, \bar{x}) . \end{aligned} \tag{5.5}$$

Next we can consider the population immediately responding to the algorithm, which has dynamics

$$\begin{aligned} \frac{d}{dt} x &= -\nabla_x G_c(\rho, x)|_{\rho=r(x)} , \\ r(x) &:= \operatorname{argmin}_{\hat{\rho} \in \mathcal{P}} -G_c(\hat{\rho}, x) . \end{aligned} \tag{5.6}$$

In this time-scale separated setting, model (5.5) represents a dynamic maximization of G_c with respect to ρ in Wasserstein-2 space, and an instantaneous minimization of G_c with respect to the algorithm parameters x . Model (5.6) represents an instantaneous maximization of G_c with respect to ρ and a dynamic minimization of G_c with respect to the algorithm parameters x . The key results on existence and uniqueness of a ground state as well as the convergence behavior of solutions depend on convexity (concavity) of G_a and G_c . The notion of convexity that we will employ for energy functionals in the Wasserstein-2 geometry is *(uniform) displacement convexity*, which is analogous to (strong) convexity in Euclidean spaces. One can think of displacement convexity for an energy functional defined on $\mathcal{P}_2$ as convexity along the shortest path in the Wasserstein-2 metric (linear interpolation in the Wasserstein-2 space) between any two given probability distributions. For a detailed definition of (uniform) displacement convexity and concavity, see Section 5.A.2. In fact, suitable convexity properties of the input functions f_1, f_2, W and $\tilde{\rho}$ will ensure (uniform) displacement convexity of the resulting energy functionals appearing in the gradient flow structure, see for instance [204, Chapter 5.2].

We make the following assumptions in both the competitive case and aligned interest cases. Here, I_d denotes the $d \times d$ identity matrix, $\text{Hess}(f)$ denotes the Hessian of f in all variables, while $\nabla_x^2 f$ denotes the Hessian of f in the variable x only.

Assumption 5.5 (Convexity of f_1 and f_2). *The functions $f_1, f_2 \in C^2(\mathbb{R}^d \times \mathbb{R}^d; [0, \infty))$ satisfy for all $(z, x) \in \mathbb{R}^d \times \mathbb{R}^d$ the following:*

- *There exists constants $\lambda_1, \lambda_2 \geq 0$ such that $\text{Hess}(f_1) \succeq \lambda_1 I_{2d}$ and $\nabla_x^2 f_2 \succeq \lambda_2 I_d$;*
- *There exist constants $a_i > 0$ such that $x \cdot \nabla_x f_i(z, x) \geq -a_i$ for $i = 1, 2$;*

Assumption 5.6 (Reference Distribution Shape). *The reference distribution $\tilde{\rho} \in \mathcal{P}(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ satisfies $\log \tilde{\rho} \in C^2(\mathbb{R}^d)$ and $\nabla_z^2 \log \tilde{\rho}(z) \preceq -\tilde{\lambda} I_d$ for some $\tilde{\lambda} > 0$.*

Assumption 5.7 (Convex Interaction Kernel). *The interaction kernel $W \in C^2(\mathbb{R}^d; [0, \infty))$ is convex, symmetric $W(-z) = W(z)$, and for some $D > 0$ satisfies*

$$z \cdot \nabla_z W(z) \geq -D, \quad |\nabla_z W(z)| \leq D(1 + |z|) \quad \forall z \in \mathbb{R}^d.$$

We make the following observations regarding the assumptions above:

- The convexity in Assumption 5.7 can be relaxed and without affecting the results outlined below by following a more detailed analysis analogous to the approach in [50].
- If f_1 and f_2 are strongly convex, the proveable convergence rate increases, but without strict or strong convexity of f_1 and f_2 , the regularizers $KL(\rho|\tilde{\rho})$ and $\int \|x - x_0\|_2^2 dx$ provide the convexity guarantees necessary for convergence.

For concreteness, one can consider the following classical choices of input functions to the evolution:

- Using the log-loss function for f_1 and f_2 satisfies Assumption 5.5.
- Taking the reference measure $\tilde{\rho}$ to be the normal distribution satisfies Assumption 5.6, which ensures the distribution is not too flat.
- Taking quadratic interactions $W(z) = \frac{1}{2}|z|^2$ satisfies Assumption 5.7.

Remark 5.8 (Cauchy-Problem). *To complete the arguments on convergence to equilibrium, we require sufficient regularity of solutions to the PDEs under consideration. In fact, it is sufficient if we can show that equations (5.4), (5.5), and (5.6) can be approximated by equations with smooth solutions. Albeit tedious, these are standard techniques in the regularity theory for partial differential equations, see for example [50, Proposition 2.1 and Appendix A], [176], [204, Chapter 9], and the references therein. Similar arguments as in [85] are expected to apply to the coupled gradient flows considered here, guaranteeing existence of smooth solutions with fast enough decay at infinity, and we leave a detailed proof for future work.*

5.3.2 Analysis of Case (i): Aligned Objectives

The primary technical contribution of this setting consists of lifting the algorithm dynamics from an ODE to a PDE, which allows us to model the system as a joint gradient flow on the product space of probability measures. The coupling occurs in the potential function, rather than as cross-diffusion or non-local interaction as more commonly seen in the literature for multi-species systems.

Theorem 5.9. *Suppose that Assumptions 5.5-5.7 are satisfied and let $\lambda_a := \lambda_1 + \min(\lambda_2 + \beta, \alpha\tilde{\lambda}) > 0$. Consider solutions $\gamma_t := (\rho_t, \mu_t)$ to the dynamics (5.4) with initial conditions satisfying $\gamma_0 \in \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^d)$ and $G_a(\gamma_0) < \infty$. Then the following hold:*

- There exists a unique minimizer $\gamma_\infty = (\rho_\infty, \mu_\infty)$ of G_a , which is also a steady state for equation (5.4). Moreover, $\rho_\infty \in L^1(\mathbb{R}^d)$, has the same support as $\tilde{\rho}$, and its density is continuous.*
- The solution γ_t converges exponentially fast in $G_a(\cdot | \gamma_\infty)$ and $\bar{W}$,*

$$G_a(\gamma_t | \gamma_\infty) \leq e^{-2\lambda_a t} G_a(\gamma_0 | \gamma_\infty) \quad \text{and} \quad \bar{W}(\gamma_t, \gamma_\infty) \leq c e^{-\lambda_a t} \quad \text{for all } t \geq 0,$$

where $c > 0$ is a constant only depending on γ_0 , γ_∞ and the parameter λ_a .

Proof. (Sketch) For existence and uniqueness, we leverage classical techniques in the calculus of variations. To obtain convergence to equilibrium in energy, our key result is a new HWI-type inequality, providing as a consequence generalizations of the log-Sobolev inequality and the Talagrand inequality. Together, these inequalities relate the energy (classically denoted by H in the

case of the Boltzmann entropy), the metric (classically denoted by W in the case of the Wasserstein-2 metric) and the energy dissipation (classically denoted by I in the case of the Fisher information)¹. Combining these inequalities with Gronwall's inequality allows us to deduce convergence both in energy and in the metric $\overline{W}$. $\square$

5.3.3 Analysis of Case (ii): Competing Objectives

In this setting, we consider the case where the algorithm and the strategic population have goals in opposition to each other; specifically, the population benefits from being classified incorrectly. First, we will show that when the algorithm instantly best-responds to the population, then the distribution of the population converges exponentially in energy and in W_2 . Then we will show a similar result for the case where the population instantly best-responds to the algorithm.

In both cases, we begin by proving two Danskin-type results (see [74, 24]) which will be used for the main convergence theorem, including convexity (concavity) results. To this end, we make the following assumption ensuring that the regularizing component in the evolution of ρ is able to control the concavity introduced by f_1 and f_2 .

Assumption 5.10 (Upper bounds for f_1 and f_2). *There exists a constant $\Lambda_1 > 0$ such that*

$$\nabla_z^2 f_1(z, x) \preceq \Lambda_1 I_d \quad \text{for all } (z, x) \in \mathbb{R}^d \times \mathbb{R}^d,$$

and for any $R > 0$ there exists a constant $c_2 = c_2(R) \in \mathbb{R}$ such that

$$\sup_{x \in B_R(0)} \int f_2(z, x) d\bar{\rho}(z) < c_2.$$

Equipped with Assumption 5.10, we state the result for a best-responding algorithm.

Theorem 5.11. *Suppose Assumptions 5.5-5.10 are satisfied with $\alpha\tilde{\lambda} > \Lambda_1$. Let $\lambda_b := \alpha\tilde{\lambda} - \Lambda_1$. Define $G_b(\rho) := G_c(\rho, b(\rho))$. Consider a solution ρ_t to the dynamics (5.5) with initial condition $\rho_0 \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho_0) < \infty$. Then the following hold:*

- (a) *There exists a unique maximizer ρ_∞ of $G_b(\rho)$, which is also a steady state for equation (5.5). Moreover, $\rho_\infty \in L^1(\mathbb{R}^d)$, has the same support as $\tilde{\rho}$, and its density is continuous.*
- (b) *The solution ρ_t converges exponentially fast to ρ_∞ with rate λ_b in $G_b(\cdot | \rho_\infty)$ and W_2 ,*

$$G_b(\rho_t | \rho_\infty) \leq e^{-2\lambda_b t} G_a(\rho_0 | \rho_\infty) \quad \text{and} \quad W_2(\rho_t, \rho_\infty) \leq c e^{-\lambda_b t} \quad \text{for all } t \geq 0,$$

where $c > 0$ is a constant only depending on ρ_0 , ρ_∞ and the parameter λ_b .

¹Hence the name HWI inequalities.

Proof. (Sketch) The key addition in this setting as compared with Theorem 5.9 is proving that $G_b(\rho)$ is bounded below, uniformly displacement concave and guaranteeing its smoothness via Berge's Maximum Theorem. This is non-trivial as it uses the properties of the best response $b(\rho)$. A central observation for our arguments to work is that $\delta_\rho G_b[\rho] = (\delta_\rho G_c[\rho, x])|_{x=b(\rho)}$. We can then conclude using the direct method in the calculus of variations and the HWI method. $\square$

Here, the condition that $\alpha\tilde{\lambda}$ must be large enough corresponds to the statement that the system must be subjected to a strong enough regularizing effect.

In the opposite case, where ρ instantly best-responds to the algorithm, we show Danskin-like results for derivatives through the best response function and convexity of the resulting energy in x which allows to deduce convergence.

Theorem 5.12. *Suppose Assumptions 5.5-5.10 are satisfied with $\alpha\tilde{\lambda} > \Lambda_1$, and that $r(x)$ is differentiable (as shown by example conditions in Lemmas 5.52 and 5.53). Define $G_d(x) := G_c(r(x), x)$. Then it holds:*

- (a) *There exists a unique minimizer x_∞ of $G_d(x)$ which is also a steady state for (5.6).*
- (b) *The vector $x(t)$ solving the dynamics (5.6) with initial condition $x(0) \in \mathbb{R}^d$ converges exponentially fast to x_∞ with rate $\lambda_d := \lambda_1 + \lambda_2 + \beta > 0$ in G_d and in the Euclidean norm:*

$$\begin{aligned} \|x(t) - x_\infty\| &\leq e^{-\lambda_d t} \|x(0) - x_\infty\|, \\ G_d(x(t)) - G_d(x_\infty) &\leq e^{-2\lambda_d t} (G_d(x(0)) - G_d(x_\infty)) \end{aligned}$$

for all $t \geq 0$.

These two theorems illustrate that, under sufficient convexity conditions on the cost functions, we expect the distribution ρ and the algorithm x to converge to a steady state. In practice, when the distributions are close enough to the steady state there is no need to retrain the algorithm.

While we have proven results for the extreme timescale cases, we anticipate convergence to the same equilibrium in the intermediate cases. Indeed, it is well known [27] (especially for systems in Euclidean space) that for two-timescale stochastic approximations of dynamical systems, with appropriate stepsize choices, converge asymptotically, and finite-time high probability concentration bounds can also be obtained. These results have been leveraged in strategic classification [215] and Stackelberg games [108, 110, 109]. We leave this intricate analysis to future work.

In the following section we show numerical results in the case of a best-responding x , best-responding ρ , and in between where x and ρ evolve on a similar timescale. Note that in these settings, the dynamics do not have a gradient flow structure due to a sign difference in the energies, requiring conditions to ensure that one species does not dominate the other.

5.4 Numerical Examples

We illustrate numerical results for the case of a classifier, which are used in scenarios such as loan or government aid applications [42], school admissions [180], residency match [188], and recommendation algorithms [141], all of which have some population which is incentivized to submit data that will result in a desirable classification. For all examples, we select classifiers of the form $x \in \mathbb{R}$, so that a data point $z \in \mathbb{R}$ is assigned a label of 1 with probability $q(z, x) = (1 + \exp(-b^\top z + x))^{-1}$ where $b > 0$. Let f_1 and f_2 be given by

$$f_1(z, x) = -\log(1 - q(z, x)), \quad f_2(z, x) = -\log q(z, x).$$

Note that $\text{Hess}(f_1) \succeq 0$ and $\nabla_x^2 f_2 \succeq 0$, so $\lambda_1 = \lambda_2 = 0$. Here, the strictness of the convexity of the functional is coming from the regularizers, not the cost functions, with $\tilde{\rho}$ a scaled normal distribution. We show numerical results for two scenarios with additional settings in the appendix. First we illustrate competitive interests under three different timescale settings. Then we simulate the classifier taking an even more naïve strategy than gradient descent and discuss the results. The PDEs were implemented based on the finite volume method from [47].

5.4.1 Competitive Objectives

In the setting with competitive objectives, we utilize $G_c(\rho, x)$ with $W = 0$, f_1 and f_2 as defined above with $b = 3$ fixed as it only changes the steepness of the classifier for $d = 1$, and $\alpha = 0.1$ and $\beta = 0.05$. In Figure 5.1, we simulate two extremes of the timescale setting; first when ρ is nearly best-responding and then when x is best-responding. The simulations have the same initial conditions and end with the same distribution shape; however, the behavior of the strategic population differs in the intermediate stages. When ρ is nearly best-responding, we see that the distribution quickly shifts




Figure 5.1: When x versus ρ best-responds, we observe the same final state but different intermediate states. Modes appear in the strategic population which simpler models cannot capture.

mass over the classifier threshold. Then the classifier shifts right, correcting for the shift in ρ , which then incentivizes ρ to shift more mass back to the original mode. In contrast, when x best-responds, the right-hand mode slowly increases in size until the system converges.

Figure 5.2 shows simulation results from the setting where ρ and x evolve on the same timescale. We observe that the distribution shift in ρ appears to fall between the two extreme timescale cases, which we expect. We highlight two important observations for the competitive case. One, a single-mode




Figure 5.2: In this experiment the population and classifier have similar rates of change, and the distribution change for ρ exhibits behaviors from both the fast ρ and fast x simulations; the right-hand mode does not peak as high as the fast ρ case but does exceed its final height and return to the equilibrium.


distribution becomes bimodal, which would not be captured using simplistic metrics such as the mean and variance. This split can be seen as polarization in the population, a phenomenon that a mean-based strategic classification model would not capture. Two, the timescale on which the classifier updates significantly impacts the intermediate behavior of the distribution. In our example, when x updated slowly relative to the strategic population, the shifts in the population were greater than in the other two cases. This suggests that understanding the effects of timescale separation are important for minimizing volatility of the coupled dynamics.

5.4.2 Naïve Behavior


In this example, we explore the results of the classifier adopting a non-gradient-flow strategy, where the classifier chooses an initially-suboptimal value for x and does not move, allowing the strategic population to respond. All functions and parameters are the same as in the previous example. When comparing with the gradient descent strategy, we observe that while the initial loss for the classifier is worse for the naive strategy, the final cost is better. While this results is not surprising, because one can view this as a general-sum game where the best response to a fixed decision may be better than the equilibrium, it illustrates how our method provides a framework for evaluating how different training strategies perform in the long run against a strategic population.

5.5 Future Directions, Limitations, and Broader Impact

Our work presents a method for evaluating the robustness of an algorithm to a strategic population, and investigating a variety of robustness using our techniques opens a range of future research directions. Our application suggests many questions relevant to the PDE literature, such as: (1) Does convergence still hold with the gradient replaced by an estimated gradient? (2) Can we prove convergence in between the two timescale extremes? (3) How do multiple dynamic populations respond to an algorithm, or multiple algorithms? In the realm of learning algorithms, our framework



(a) Both species minimize their respective losses; when the classifier uses a naïve strategy, the final performance is better for the classifier and uniformly worse for the population.



(b) The classifier selects a suboptimal initial condition $x = 2.2$, instead of $x = 1.5$ which minimizes the initial loss, and then does not move in response to the population.

Figure 5.3: Although the classifier starts with a larger cost by taking the naïve strategy, the final loss is better. This illustrates how our model can be used to compare robustness of different strategies against a strategic population.

can be extended to other learning update strategies and presents a way to model how we can design these update strategies to induce desired behaviors in the population.

A challenge in our method is that numerically solving high-dimensional PDEs is computationally expensive and possibly unfeasible. Here we note that in many applications, agents in the population do not alter more than a few features due to the cost of manipulation. We are encouraged by the recent progress using deep learning to solve PDEs, which could be used in our application.

Broader Impacts Modeling the full population distribution rather than simple metrics of the distribution is important because not all individuals are affected by the algorithm in the same way. For example, if there are tails of the distribution that have poor performance even if on average the model is good, we need to know how that group is advantaged or disadvantaged relative to the rest of the population. Additionally, understanding how people respond to algorithms offers an opportunity to incentivise people to move in a direction that increases social welfare.

LC is supported by an NDSEG fellowship from the Air Force Office of Scientific Research. FH is supported by start-up funds at the California Institute of Technology. LR is supported by ONR YIP N00014-20-1-2571 P00003 and NSF Awards CAREER 1844729, and CPS 1931718. EM acknowledges support from NSF Award 2240110. We are grateful for helpful discussions with José A. Carrillo.

5.A General structure and preliminaries

In this section, we give more details on the models discussed in the main article, and introduce definitions and notation that are needed for the subsequent proofs.

5.A.1 Structure of the dynamics

For the case of aligned objectives, the full coupled system of PDEs (5.4) can be written as

$$\partial_t \rho = \alpha \Delta \rho + \operatorname{div} \left(\rho \nabla_z \left(\int f_1 d\mu - \alpha \log \tilde{\rho} + W * \rho \right) \right), \quad (5.7a)$$

$$\partial_t \mu = \operatorname{div} \left(\mu \nabla_x \left(\int f_1 d\rho + \int f_2 d\bar{\rho} + \frac{\beta}{2} \|x - x_0\|^2 \right) \right). \quad (5.7b)$$

In other words, the population ρ in (5.7a) is subject to an isotropic diffusive force with diffusion coefficient $\alpha > 0$, a drift force due to the time-varying confining potential $\int f_1 d\mu(t) - \alpha \log \tilde{\rho}$, and a self-interaction force via the interaction potential W . If we consider the measure μ to be given and fixed in time, this corresponds exactly to the type of parabolic equation studied in [50]. Here however the dynamics are more complex due to the coupling of the confining potential with the dynamics (5.7b) for $\mu(t)$ via the coupling potential f_1 . Before presenting the analysis of this model, let us give a bit more intuition on the meaning and the structure of these dynamics.

In the setting where μ represents a binary classifier, we can think of the distribution $\bar{\rho}$ as modelling all those individuals carrying the true label 1, say, and the distribution $\rho(t)$ as modelling all those individuals carrying a true label 0, say, where 0 and 1 denote the labels of two classes of interest. The term $\int f_1(z, x) \mu(t, dx)$ represents a penalty for incorrectly classifying an individual at z with true label 0 when using the classifier $\mu(t, x)$. In other words, $\int f_1(z, x) \mu(t, dx) \in [0, \infty)$ is increasingly large the more z digresses from the correct classification 0. Similarly, $\int f_1(z, x) \rho(t, dz) \in [0, \infty)$ is increasingly large if the population ρ shifts mass to locations in z where the classification is incorrect. The terminology *aligned objectives* refers to the fact that in (5.7) both the population and the classifier are trying to evolve in a way as to maximize correct classification. Analogously, the term $\int f_2(z, x) \bar{\rho}(dz)$ is large if x would incorrectly classify the population $\bar{\rho}$ that carries the label 1. A natural extension of the model (5.7) would be a setting where also the population carrying labels 1 evolves over time, which is simulated in Section 5.E.2. Most elements of the framework presented here would likely carry over the setting of three coupled PDEs: one for the evolution of $\rho(t)$, one for the evolution of $\bar{\rho}(t)$ and one for the classifier $\mu(t)$.

The term

$$\alpha \Delta \rho - \alpha \operatorname{div} (\rho \nabla \log \tilde{\rho}) = \alpha \operatorname{div} (\rho \nabla \delta_\rho KL(\rho | \tilde{\rho}))$$

forces the evolution of $\rho(t)$ to approach $\tilde{\rho}$. In other words, it penalizes (in energy) deviations from a given reference measure $\tilde{\rho}$. In the context of the application at hand, we take $\tilde{\rho}$ to be the initial

distribution $\rho(t = 0)$. The solution $\rho(t)$ then evolves away from $\tilde{\rho}$ over time due to the other forces that are present. Therefore, the term $KL(\rho \mid \tilde{\rho})$ in the energy both provides smoothing of the flow and a penalization for deviations away from the reference measure $\tilde{\rho}$.

The self-interaction term $W * \rho$ introduces non-locality into the dynamics, as the decision for any given individual to move in a certain direction is influenced by the behavior of all other individuals in the population. The choice of W is application dependent. Very often, the interaction between two individuals only depends on the distance between them. This suggests a choice of W as a radial function, i.e. $W(z) = \omega(|z|)$. A choice of $\omega : \mathbb{R} \rightarrow \mathbb{R}$ such that $\omega'(r) > 0$ corresponds to an *attractive* force between individuals, whereas $\omega'(r) < 0$ corresponds to a *repulsive* force. The statement $|z|\omega'(|z|) = z \cdot \nabla_z W(z) \geq -D$ in Assumption 5.7 therefore corresponds to a requirement that the self-interaction force is not too repulsive. Neglecting all other forces in (5.7a), we obtain the non-local interaction equation

$$\partial_t \rho = \operatorname{div}(\rho \nabla W * \rho)$$

which appears in many instances in mathematical biology, mathematical physics, and material science, and it is an equation that has been extensively studied over the past few decades, see for example [45, 15, 53, 22, 23, 51, 44] and references therein.

Using the results from these works, our assumptions on the interaction potential W can be relaxed in many ways, for example by allowing discontinuous derivatives at zero for W , or by allowing W to be negative.

The dynamics (5.7b) for the algorithm μ is a non-autonomous transport equation,

$$\partial_t \mu = \operatorname{div}(\mu v) ,$$

where the time-dependence in the velocity field

$$v(t, x) := \nabla_x \left(\int f_1(z, x) d\rho(t, z) + \int f_2(z, x) d\tilde{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 \right) ,$$

comes through the evolving population $\rho(t)$. This structure allows to obtain an explicit solution for $\mu(t)$ in terms of the initial condition μ_0 and the solution $\rho(t)$ to (5.7a) using the method of characteristics.

Proposition 5.13. *Assume that there exists a constant $c > 0$ such that*

$$\left\| \int \nabla_x f_1(z, x) d\rho(z) + \int \nabla_x f_2(z, x) d\tilde{\rho}(z) \right\| \leq c(1 + \|x\|) \quad \forall \rho \in \mathcal{P}_2(\mathbb{R}^d) \text{ and } \forall x \in \mathbb{R}^d . \quad (5.8)$$

Then the unique distributional solution $\mu(t)$ to (5.7b) is given by

$$\mu(t) = \Phi(t, 0, \cdot)_{\#} \mu_0 , \quad (5.9)$$

where $\Phi(t, s, x)$ solves the characteristic equation

$$\partial_s \Phi(s, t, x) + v(s, \Phi(s, t, x)) = 0, \quad \Phi(t, t, x) = x. \quad (5.10)$$

Proof. Thanks to Assumption 5.5, we have that $v \in C^1(\mathbb{R} \times \mathbb{R}^d; \mathbb{R}^d)$, and by (5.8), we have

$$\|v(t, x)\| \leq c(1 + \|x\|) \quad \text{for all } t \geq 0, x \in \mathbb{R}^d.$$

By classical Cauchy-Lipschitz theory for ODEs, this guarantees the existence of a unique global solution $\Phi(t, s, x)$ solving (5.10). Then it can be checked directly that $\mu(t)$ as defined in (5.9) is a distributional solution to (5.7b). $\square$

In the characteristic equation (5.10), $\Phi(s, t, x)$ is a parametrization of all trajectories: if a particle was at location x at time t , then it is at location $\Phi(s, t, x)$ at time s . Our assumptions on f_1, f_2 and $\bar{\rho}$ also ensure that $\Phi(s, t, \cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a C^1 -diffeomorphism for all $s, t \in \mathbb{R}$. For more details on transport equations, see for example [77].

Remark 5.14. Consider the special case where $\mu_0 = \delta_{x_0}$ for some initial position $x_0 \in \mathbb{R}^d$. Then by Proposition 5.13, the solution to (5.7b) is given by $\mu(t) = \delta_{x(t)}$, where $x(t) := \Phi(t, 0, x_0)$ solves the ODE

$$\dot{x}(t) = -v(t, x(t)), \quad x(0) = x_0,$$

which is precisely of type (5.3).

For the case of competing objectives, the two models we consider can be written as

$$\begin{aligned} \partial_t \rho &= -\operatorname{div}(\rho [\nabla(f_1(z, b(\rho))) - \alpha \log(\rho/\bar{\rho}) - W * \rho]) , \\ b(\rho) &:= \operatorname{argmin}_{\bar{x}} \int f_1(z, \bar{x}) d\rho(z) + \int f_2(\bar{x}, z') d\bar{\rho}(z') + \frac{\beta}{2} \|\bar{x} - x_0\|^2 \end{aligned}$$

for (5.5), and

$$\begin{aligned} \frac{d}{dt} x &= -\nabla_x \left(\int f_1(z, x) r(x) (dz) + \int f_2(x, z') d\bar{\rho}(z') + \frac{\beta}{2} \|x - x_0\|^2 \right) , \\ r(x) &:= \operatorname{argmax}_{\hat{\rho} \in \mathcal{P}} \int f_1(z, x) d\hat{\rho}(z) - \alpha KL(\hat{\rho}|\bar{\rho}) - \frac{1}{2} \int \hat{\rho} W * \hat{\rho} . \end{aligned}$$

for (5.6).

5.A.2 Definitions and notation

Here, and in what follows, I_d denotes the $d \times d$ identity matrix, and id denotes the identity map. The energy functionals we are considering are usually defined on the set of probability measures on $\mathbb{R}^d$, denoted by $\mathcal{P}(\mathbb{R}^d)$. If we consider the subset $\mathcal{P}_2(\mathbb{R}^d)$ of probability measures with bounded second moment,

$$\mathcal{P}_2(\mathbb{R}^d) := \{\rho \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|z\|^2 d\rho(z) < \infty\},$$

then we can endow this space with the Wasserstein-2 metric.

Definition 5.15 (Wasserstein-2 Metric). *The Wasserstein-2 metric between two probability measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ is given by*

$$W_2(\mu, \nu)^2 = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|z - z'\|_2^2 d\gamma(z, z')$$

where Γ is the set of all joint probability distributions with marginals μ and ν , i.e. $\mu(dz) = \int \gamma(dz, z') dz'$ and $\nu(dz') = \int \gamma(z, dz') dz$.

The restriction to $\mathcal{P}_2(\mathbb{R}^d)$ ensures that W_2 is always finite. Then the space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ is indeed a metric space. We will make use of the fact that W_2 metrizes narrow convergence of probability measures. To make this statement precise, let us introduce two common notions of convergence for probability measures, which are a subset of the finite signed Radon measures $\mathcal{M}(\mathbb{R}^d)$.

Definition 5.16. *Consider a sequence $(\mu_n) \in \mathcal{M}(\mathbb{R}^d)$ and a limit $\mu \in \mathcal{M}(\mathbb{R}^d)$.*

- **(Narrow topology)** *The sequence (μ_n) converges narrowly to μ , denoted by $\mu_n \rightarrow \mu$, if for all continuous bounded functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$,*

$$\int_{\mathbb{R}^d} f(z) d\mu_n(z) \rightarrow \int_{\mathbb{R}^d} f(z) d\mu(z).$$

- **(Weak-* topology)** *The sequence (μ_n) converges weakly-* to μ , denoted by $\mu_n \xrightarrow{*} \mu$, if for all continuous functions vanishing at infinity (i.e. $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that for all $\epsilon > 0$ there exists a compact set $K_\epsilon \subset \mathbb{R}^d$ such that $|f(z)| < \epsilon$ on $\mathbb{R}^d \setminus K_\epsilon$), we have*

$$\int_{\mathbb{R}^d} f(z) d\mu_n(z) \rightarrow \int_{\mathbb{R}^d} f(z) d\mu(z).$$

Let us denote the set of continuous functions on $\mathbb{R}^d$ vanishing at infinity by $C_0(\mathbb{R}^d)$, and the set of continuous bounded functions by $C_b(\mathbb{R}^d)$. Note that narrow convergence immediately implies that $\mu_n(\mathbb{R}^d) \rightarrow \mu(\mathbb{R}^d)$ as the constant function is in $C_b(\mathbb{R}^d)$. This is not necessarily true for weak-*

convergence. We will later make use of the Banach-Alaoglu theorem [3], which gives weak-* compactness of the unit ball in a dual space. Note that $\mathcal{M}(\mathbb{R}^d)$ is indeed the dual of $C_0(\mathbb{R}^d)$ endowed with the sup-norm, and $\mathcal{P}(\mathbb{R}^d)$ is the unit ball in $\mathcal{M}(\mathbb{R}^d)$ using the dual norm. Moreover, if we can ensure that mass does not escape to infinity, the two notions of convergence in Definition 5.16 are in fact equivalent.

Lemma 5.17. *Consider a sequence $(\mu_n) \in \mathcal{M}(\mathbb{R}^d)$ and a measure $\mu \in \mathcal{M}(\mathbb{R}^d)$. Then $\mu_n \rightharpoonup \mu$ if and only if $\mu_n \xrightarrow{*} \mu$ and $\mu_n(\mathbb{R}^d) \rightarrow \mu(\mathbb{R}^d)$.*

This follows directly from Definition 5.16. Here, the condition $\mu_n(\mathbb{R}^d) \rightarrow \mu(\mathbb{R}^d)$ is equivalent to tightness of (μ_n) , and follows from Markov's inequality [117] if we can establish uniform bounds on the second moments, i.e. we want to show that there exists a constant $C > 0$ independent of n such that

$$\int \|z\|^2 d\mu_n(z) < C \quad \forall n \in \mathbb{N}. \quad (5.11)$$

Definition 5.18 (Tightness of probability measures). *A collection of measures $(\mu_n) \in \mathcal{M}(\mathbb{R}^d)$ is tight if for all $\epsilon > 0$ there exists a compact set $K_\epsilon \subset \mathbb{R}^d$ such that $|\mu_n|(\mathbb{R}^d \setminus K_\epsilon) < \epsilon$ for all $n \in \mathbb{N}$, where $|\mu|$ denotes the total variation of μ .*

Another classical result is that the Wasserstein-2 metric metrizes narrow convergence and weak-* convergence of probability measures, see for example [194, Theorem 5.11] or [204, Theorem 7.12].

Lemma 5.19. *Let $\mu_n, \mu \in \mathcal{P}_2(\mathbb{R}^d)$. Then $W_2(\mu_n, \mu) \rightarrow 0$ if and only if*

$$\mu_n \rightharpoonup \mu \quad \text{and} \quad \int_{\mathbb{R}^d} \|z\|^2 d\mu_n(z) \rightarrow \int_{\mathbb{R}^d} \|z\|^2 d\mu(z).$$

Remark 5.20. *Note that $\mu_n \rightharpoonup \mu$ can be replaced by $\mu_n \xrightarrow{*} \mu$ in the above statement thanks to the fact that the limit μ is a probability measure with mass 1, see Lemma 5.17.*

Next, we consider two measures $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ that are *atomless*, i.e $\mu(\{z\}) = 0$ for all $z \in \mathbb{R}^d$. By Brenier's theorem [20] (also see [204, Theorem 2.32]) there exists a unique measurable map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $T_{\#}\mu = \nu$, and $T = \nabla\psi$ for some convex function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$. Here, the *push-forward* operator $\nabla\psi_{\#}$ is defined as

$$\int_{\mathbb{R}^d} f(z) d\nabla\psi_{\#}\rho_0(z) = \int_{\mathbb{R}^d} f(\nabla\psi(z)) d\rho_0(z)$$

for all Borel-measurable functions $f : \mathbb{R}^d \mapsto \mathbb{R}_+$. If $\rho_1 = \nabla\psi_{\#}\rho_0$, we denote by $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ the *displacement interpolant* between ρ_0 and ρ_1 . We are now ready to introduce the

notion of displacement convexity, which is the same as geodesic convexity in the geodesic space $(\mathcal{P}_2(\mathbb{R}^d), W_2)$. We will state the definition here for atomless measures, but it can be relaxed to any pair of measures in $\mathcal{P}_2$ using optimal transport plans instead of transport maps. In what follows, we will use s to denote the interpolation parameter for geodesics, and t to denote time related to solutions of (5.4), (5.5) and (5.6).

Definition 5.21 (Displacement Convexity). *A functional $G : \mathcal{P} \mapsto \mathbb{R}$ is displacement convex if for all ρ_0, ρ_1 that are atomless we have*

$$G(\rho_s) \leq (1-s)G(\rho_0) + sG(\rho_1),$$

where $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ is the displacement interpolant between ρ_0 and ρ_1 . Further, $G : \mathcal{P} \mapsto \mathbb{R}$ is uniformly displacement convex with constant $\eta > 0$ if

$$G(\rho_s) \leq (1-s)G(\rho_0) + sG(\rho_1) - s(1-s)\frac{\eta}{2}W_2(\rho_0, \rho_1)^2,$$

where $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ is the displacement interpolant between ρ_0 and ρ_1 .

Remark 5.22. *In other words, G is displacement convex (concave) if the function $G(\rho_s)$ is convex (concave) with $\rho_s = [(1-s)\text{id} + s\nabla\psi]_{\#}\rho_0$ being the displacement interpolant between ρ_0 and ρ_1 . Contrast this with the classical notion of convexity (concavity) for G , where we require that the function $G((1-s)\rho_0 + s\rho_1)$ is convex (concave).*

In fact, if the energy G is twice differentiable along geodesics, then the condition $\frac{d^2}{ds^2}G(\gamma_s) \geq 0$ along any geodesic $(\rho_s)_{s \in [0,1]}$ between ρ_0 and ρ_1 is sufficient to obtain displacement convexity. Similarly, when $\frac{d^2}{ds^2}G(\rho_s) \geq \eta W_2(\rho_0, \rho_1)^2$, then G is uniformly displacement convex with constant $\eta > 0$. For more details, see [162] and [204, Chapter 5.2].

5.A.3 Steady states

The main goal in our theoretical analysis is to characterize the asymptotic behavior for the models (5.4), (5.5) and (5.6) as time goes to infinity. The steady states of these equations are the natural candidates to be asymptotic profiles for the corresponding equations. Thanks to the gradient flow structure, we expect to be able to make a connection between ground states of the energy functionals, and the steady states of the corresponding gradient flow dynamics. More precisely, any minimizer or maximizer is in particular a critical point of the energy, and therefore satisfies that the first variation is constant on disconnected components of its support. If this ground state also has enough regularity (weak differentiability) to be a solution to the equation, it immediately follows that it is in fact a steady state.

To make this connection precise, we first introduce what exactly we mean by a steady state.

Definition 5.23 (Steady states for (5.4)). *Given $\rho_\infty \in L^1_+(\mathbb{R}^d) \cap L^\infty_{loc}(\mathbb{R}^d)$ with $\|\rho_\infty\|_1 = 1$ and $\mu_\infty \in \mathcal{P}_2(\mathbb{R}^d)$, then $(\rho_\infty, \mu_\infty)$ is a steady state for the system (5.4) if $\rho_\infty \in W^{1,2}_{loc}(\mathbb{R}^d)$, $\nabla W * \rho_\infty \in L^1_{loc}(\mathbb{R}^d)$, ρ_∞ is absolutely continuous with respect to $\tilde{\rho}$, and $(\rho_\infty, \mu_\infty)$ satisfy*

$$\nabla_z \left(\int f_1(z, x) d\mu_\infty(x) + \alpha \log \left(\frac{\rho_\infty(z)}{\tilde{\rho}(z)} \right) + W * \rho_\infty(z) \right) = 0 \quad \forall z \in \text{supp}(\rho_\infty), \quad (5.12a)$$

$$\nabla_x \left(\int f_1(z, x) d\rho_\infty(z) + \int f_2(z, x) d\tilde{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 \right) = 0 \quad \forall x \in \text{supp}(\mu_\infty) \quad (5.12b)$$

in the sense of distributions.

Here, $L^1_+(\mathbb{R}^d) := \{\rho \in L^1(\mathbb{R}^d) : \rho \geq 0\}$.

Definition 5.24 (Steady states for (5.5)). *Let $\rho_\infty \in L^1_+(\mathbb{R}^d) \cap L^\infty_{loc}(\mathbb{R}^d)$ with $\|\rho_\infty\|_1 = 1$. Then ρ_∞ is a steady state for the system (5.5) if $\rho_\infty \in W^{1,2}_{loc}(\mathbb{R}^d)$, $\nabla W * \rho_\infty \in L^1_{loc}(\mathbb{R}^d)$, ρ_∞ is absolutely continuous with respect to $\tilde{\rho}$, and ρ_∞ satisfies*

$$\nabla_z \left(f_1(z, b(\rho_\infty)) - \alpha \log \left(\frac{\rho_\infty(z)}{\tilde{\rho}(z)} \right) - W * \rho_\infty(z) \right) = 0 \quad \forall z \in \mathbb{R}^d, \quad (5.13)$$

in the sense of distributions, where $b(\rho_\infty) := \arg\min_x G_c(\rho_\infty, x)$.

Definition 5.25 (Steady states for (5.6)). *The vector $x_\infty \in \mathbb{R}^d$ is a steady state for the system (5.6) if it satisfies*

$$\nabla_x G_d(x_\infty) = 0.$$

In fact, with the above notions of steady state, we can obtain improved regularity for ρ_∞ .

Lemma 5.26. *Let Assumptions 5.5-5.7 hold. Then the steady states ρ_∞ for (5.4) and (5.5) are continuous.*

Proof. We present here the argument for equation (5.5) only. The result for (5.4) follows in exactly the same way by replacing $f_1(z, b(\rho_\infty))$ with $-\int f_1(z, x) d\mu_\infty(x)$.

Thanks to our assumptions, we have $f_1(\cdot, b(\rho_\infty)) + \alpha \log \tilde{\rho}(\cdot) \in C^1$, which implies that $\nabla(f_1(\cdot, b(\rho_\infty)) + \alpha \log \tilde{\rho}(\cdot)) \in L^\infty_{loc}$. By the definition of a steady state, $\rho_\infty \in L^1 \cap L^\infty_{loc}$ and thanks to Assumption 5.7 we have $W \in C^2$, which implies that $\nabla W * \rho_\infty \in L^\infty_{loc}$. Let

$$h(z) := \rho_\infty(z) \nabla [f_1(z, b(\rho_\infty)) + \alpha \log \tilde{\rho}(z) - (W * \rho_\infty)(z)] .$$

Then by the aforementioned regularity, we obtain $h \in L^1_{loc} \cap L^\infty_{loc}$. By interpolation, it follows that $h \in L^p_{loc}$ for all $1 < p < \infty$. This implies that $\operatorname{div}(h) \in W^{-1,p}_{loc}$. Since ρ_∞ is a weak $W^{1,2}_{loc}$ -solution of (5.13), we have

$$\Delta \rho_\infty = \operatorname{div}(h) ,$$

and so by classic elliptic regularity theory we conclude $\rho_\infty \in W^{1,p}_{loc}$. Finally, applying Morrey's inequality, we have $\rho_\infty \in C^{0,k}$ where $k = \frac{p-d}{p}$ for any $d < p < \infty$. Therefore $\rho_\infty \in C(\mathbb{R}^d)$ (after possibly being redefined on a set of measure zero). $\square$

5.B Proof of Theorem 5.9

For ease of notation, we write $G_a : \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \mapsto [0, \infty]$ as

$$G_a((\rho, \mu)) = \alpha KL(\rho|\tilde{\rho}) + \mathcal{V}(\rho, \mu) + \mathcal{W}(\rho) ,$$

where we define

$$\begin{aligned} \mathcal{V}(\rho, \mu) &= \iint f_1(z, x) d\rho(z) d\mu(x) + \int V(x) d\mu(x) , \\ \mathcal{W}(\rho) &= \frac{1}{2} \iint W(z_1 - z_2) d\rho(z_1) d\rho(z_2) , \end{aligned}$$

with potential given by $V(x) := \int f_2(z, x) d\tilde{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2$.

In order to prove the existence of a unique ground state for G_a , a natural approach is to consider the corresponding Euler-Lagrange equations

$$\alpha \log \frac{\rho(z)}{\tilde{\rho}(z)} + \int f_1(z, x) d\mu(x) + (W * \rho)(z) = c_1[\rho, \mu] \quad \text{for all } z \in \operatorname{supp}(\rho) , \quad (5.14a)$$

$$\int f_1(z, x) d\rho(z) + V(x) = c_2[\rho, \mu] \quad \text{for all } x \in \operatorname{supp}(\mu) , \quad (5.14b)$$

where c_1, c_2 are constants that may differ on different connected components of $\operatorname{supp}(\rho)$ and $\operatorname{supp}(\mu)$. These equations are not easy to solve explicitly, and we are therefore using general non-constructive techniques from calculus of variations. We first show continuity and convexity properties for the functional G_a (Lemma 5.27 and Proposition 5.28), essential properties that will allow us to deduce existence and uniqueness of ground states using the direct method in the calculus of variations (Proposition 5.30). Using the Euler-Lagrange equation 5.14, we then prove properties on the support of the ground state (Corollary 5.32). To obtain convergence results, we apply the HWI method: we first show a general 'interpolation' inequality between the energy, the energy dissipation and the metric (Proposition 5.34); this fundamental inequality will then imply a generalized logarithmic Sobolev inequality (Corollary 5.36) relating the energy to the energy dissipation, and a generalized

Talagrand inequality (Corollary 5.37) that allows to translate convergence in energy into convergence in metric. Putting all these ingrediends together will then allow us to conclude for the statements in Theorem 5.9.

Lemma 5.27 (Lower semi-continuity). *Let Assumptions 5.5-5.7 hold. Then the functional $G_a : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ is lower semi-continuous with respect to the weak-* topology.*

Proof. We split the energy G_a into three parts: (i) $KL(\rho|\tilde{\rho})$, (ii) the interaction energy $\mathcal{W}$, and (iii) the potential energy $\mathcal{V}$. For (i), lower semi-continuity has been shown in [184]. For (ii), we can directly apply [194, Proposition 7.2] using Assumption 5.7. For (iii), note that V and f_1 are lower semi-continuous and bounded below thanks to Assumption 5.5, and so the result follows from [194, Proposition 7.1]. $\square$

Proposition 5.28 (Uniform displacement convexity). *Let $\alpha, \beta > 0$. Fix $\gamma_0, \gamma_1 \in \mathcal{P}_2 \times \mathcal{P}_2$ and let Assumptions 5.5-5.7 hold. Along any geodesic $(\gamma_s)_{s \in [0,1]} \in \mathcal{P}_2 \times \mathcal{P}_2$ connecting γ_0 to γ_1 , we have for all $s \in [0, 1]$*

$$\frac{d^2}{ds^2} G_a(\gamma_s) \geq \lambda_a \overline{W}(\gamma_0, \gamma_1)^2, \quad \lambda_a := \lambda_1 + \min(\lambda_2 + \beta, \alpha \tilde{\lambda}). \quad (5.15)$$

As a result, the functional $G_a : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ is uniformly displacement convex with constant $\lambda_a > 0$.

Proof. Let γ_0 and γ_1 be two probability measures with bounded second moments. Denote by $\phi, \psi : \mathbb{R}^d \rightarrow \mathbb{R}$ the optimal Kantorovich potentials pushing ρ_0 onto ρ_1 , and μ_0 onto μ_1 , respectively:

$$\begin{aligned} \rho_1 &= \nabla \phi_{\#} \rho_0 \quad \text{such that} \quad W_2(\rho_0, \rho_1)^2 = \int_{\mathbb{R}^d} \|z - \nabla \phi(z)\|^2 d\rho_0(z), \\ \mu_1 &= \nabla \psi_{\#} \mu_0 \quad \text{such that} \quad W_2(\mu_0, \mu_1)^2 = \int_{\mathbb{R}^d} \|x - \nabla \psi(x)\|^2 d\mu_0(x). \end{aligned}$$

The now classical results in [20] guarantee that there always exists convex functions ϕ, ψ that satisfy the conditions above. Then the path $(\gamma_s)_{s \in [0,1]} = (\rho_s, \mu_s)_{s \in [0,1]}$ defined by

$$\begin{aligned} \rho_s &= [(1-s) \text{id} + s \nabla_z \phi]_{\#} \rho_0, \\ \mu_s &= [(1-s) \text{id} + s \nabla_x \psi]_{\#} \mu_0 \end{aligned}$$

is a $\overline{W}$ -geodesic from γ_0 to γ_1 .

The first derivative of $\mathcal{V}$ along geodesics in the Wasserstein metric is given by

$$\begin{aligned}
\frac{d}{ds}\mathcal{V}(\gamma_s) &= \frac{d}{ds} \left[\iint f_1((1-s)z + s\nabla\phi(z), (1-s)x + s\nabla\psi(x)) d\rho_0(z)d\mu_0(x) \right. \\
&\quad \left. + \int V((1-s)x + s\nabla\psi(x)) d\mu_0(x) \right] \\
&= \iint \nabla_x f_1((1-s)z + s\nabla\phi(z), (1-s)x + s\nabla\psi(x)) \cdot (\nabla\psi(x) - x) d\rho_0(z)d\mu_0(x) \\
&\quad + \iint \nabla_z f_1((1-s)z + s\nabla\phi(z), (1-s)x + s\nabla\psi(x)) \cdot (\nabla\phi(z) - z) d\rho_0(z)d\mu_0(x) \\
&\quad + \int \nabla_x V((1-s)x + s\nabla\psi(x)) \cdot (\nabla\psi(x) - x) d\mu_0(x),
\end{aligned}$$

and taking another derivative we have

$$\begin{aligned}
\frac{d^2}{ds^2}\mathcal{V}(\gamma_s) &= - \iint \begin{bmatrix} (\nabla\psi(x) - x) \\ (\nabla\phi(z) - z) \end{bmatrix}^T \cdot D_s(z, x) \cdot \begin{bmatrix} (\nabla\psi(x) - x) \\ (\nabla\phi(z) - z) \end{bmatrix} d\rho_0(z)d\mu_0(x) \\
&\quad + \iint (\nabla\psi(x) - x)^T \cdot \nabla_x^2 V((1-s)x + s\nabla\psi(x)) \cdot (\nabla\psi(x) - x) d\rho_0(z)d\mu_0(x) \\
&\geq \lambda_1 \overline{W}(\gamma_0, \gamma_1)^2 + (\lambda_2 + \beta) W_2(\mu_0, \mu_1)^2,
\end{aligned}$$

where we denoted $D_s(z, x) := \text{Hess}(f_1)((1-s)z + s\nabla\phi(z), (1-s)x + s\nabla\psi(x))$, and the last inequality follows from Assumption 5.5 and the optimality of the potentials ϕ and ψ .

Following [50, 204] and using Assumption 5.6, the second derivatives of the diffusion term and the interaction term along geodesics are given by

$$\frac{d^2}{ds^2} KL(\rho_s | \tilde{\rho}) \geq \alpha \tilde{\lambda} W_2(\rho_0, \rho_1)^2, \quad \frac{d^2}{ds^2} \mathcal{W}(\rho_s) \geq 0. \quad (5.16)$$

Putting the above estimates together, we obtain (5.15). $\square$

Remark 5.29. *Alternatively, one could assume strong convexity of W , which would improve the lower-bound on the second derivative along geodesics.*

Proposition 5.30. *(Ground state) Let Assumptions 5.5-5.7 hold for $\alpha, \beta > 0$. Then the functional $G_a : \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathbb{R}^d) \rightarrow [0, \infty]$ admits a unique minimizer $\gamma_* = (\rho_*, \mu_*)$, and it satisfies $\rho_* \in \mathcal{P}_2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$, $\mu_* \in \mathcal{P}_2(\mathbb{R}^d)$, and ρ_* is absolutely continuous with respect to $\tilde{\rho}$.*

Proof. We show existence of a minimizer of G_a using the direct method in the calculus of variations. Denote by $\gamma = (\rho, \mu) \in \mathcal{P} \times \mathcal{P} \subset \mathcal{M} \times \mathcal{M}$ a pair of probability measures as a point in the product space of Radon measures. Since $G_a \geq 0$ on $\mathcal{P} \times \mathcal{P}$ (see Assumption 5.5) and not identically $+\infty$ everywhere, there exists a minimizing sequence $(\gamma_n) \in \mathcal{P} \times \mathcal{P}$. Note that

(γ_n) is in the closed unit ball of the dual space of continuous functions vanishing at infinity $(C_0(\mathbb{R}^d) \times C_0(\mathbb{R}^d))^*$ endowed with the dual norm $\|\gamma_n\|_* = \sup \frac{|\int f d\rho_n + \int g d\mu_n|}{\|(f,g)\|_\infty}$ over $f, g \in C_0(\mathbb{R}^d)$ with $\|(f,g)\|_\infty := \|f\|_\infty + \|g\|_\infty \neq 0$. By the Banach-Alaoglu theorem [192, Thm 3.15] there exists a limit $\gamma_* = (\rho_*, \mu_*) \in \mathcal{M} \times \mathcal{M} = (C_0 \times C_0)^*$ and a convergent subsequence (not relabelled) such that $\gamma_n \xrightarrow{*} \gamma_*$. In fact, since $KL(\rho_* | \tilde{\rho}) < \infty$ it follows that ρ_* is absolutely continuous with respect to $\tilde{\rho}$, implying $\rho_* \in L^1(\mathbb{R}^d)$ thanks to Assumption 5.6. Further, μ_* has bounded second moment, else we would have $\inf_{\gamma \in \mathcal{P} \times \mathcal{P}} G_a(\gamma) = \infty$ which yields a contradiction. It remains to show that $\int d\rho_* = \int d\mu_* = 1$ to conclude that $\gamma_* \in \mathcal{P} \times \mathcal{P}$. To this aim, it is sufficient to show tightness of (ρ_n) and (μ_n) , preventing the escape of mass to infinity as we have $\int d\rho_n = \int d\mu_n = 1$ for all $n \geq 1$. Tightness follows from Markov's inequality [117] if we can establish uniform bounds on the second moments, i.e. we want to show that there exists a constant $C > 0$ independent of n such that

$$\int \|z\|^2 d\rho_n(z) + \int \|x\|^2 d\mu_n(x) < C \quad \forall n \in \mathbb{N}. \quad (5.17)$$

To establish (5.17), observe that thanks to Assumption 5.6, there exists a constant $c_0 \in \mathbb{R}$ (possibly negative) such that $-\log \tilde{\rho}(z) \geq c_0 + \frac{\tilde{\lambda}}{4} \|z\|^2$ for all $z \in \mathbb{R}^d$. Then

$$\frac{\alpha \tilde{\lambda}}{4} \int \|z\|^2 d\rho_n \leq -\alpha c_0 - \alpha \int \log \tilde{\rho}(z) d\rho_n$$

Therefore, using $\int d\rho_n = \int d\mu_n = 1$ and writing $\zeta := \min\{\frac{\alpha \tilde{\lambda}}{4}, \frac{\beta}{2}\} > 0$, we obtain the desired uniform upper bound on the second moments of the minimizing sequence,

$$\begin{aligned} \zeta \iint (\|z\|^2 + \|x\|^2) d\rho_n d\mu_n &\leq -\alpha c_0 - \alpha \int \log \tilde{\rho}(z) d\rho_n + \beta \int \|x - x_0\|^2 d\mu_n + \beta \|x_0\|^2 \\ &\leq -\alpha c_0 + \beta \|x_0\|^2 + G_a(\gamma_n) \\ &\leq -\alpha c_0 + \beta \|x_0\|^2 + G_a(\gamma_1) < \infty. \end{aligned}$$

This concludes the proof that the limit γ_* satisfies indeed $\gamma_* \in \mathcal{P} \times \mathcal{P}$, and indeed $\rho_* \in \mathcal{P}_2(\mathbb{R}^d)$ as well. Finally, γ_* is indeed a minimizer of G_a thanks to weak-* lower-semicontinuity of G_a following Lemma 5.27.

Next we show uniqueness using a contradiction argument. Suppose $\gamma_* = (\rho_*, \mu_*)$ and $\gamma'_* = (\rho'_*, \mu'_*)$ are minimizers of G_a . For $s \in [0, 1]$, define $\gamma_s := ((1-s)\text{id} + sT, (1-s)\text{id} + sS)_\# \gamma_*$, where $T, S : \mathbb{R}^d \mapsto \mathbb{R}^d$ are the optimal transport maps such that $\rho'_* = T_\# \rho_*$ and $\mu'_* = S_\# \mu_*$. By Proposition 5.28 the energy G_a is uniformly displacement convex, and so we have

$$G_a(\gamma_s) \leq (1-s)G_a(\gamma_*) + sG_a(\gamma'_*) = G_a(\gamma_*).$$

If $\gamma_* \neq \gamma'_*$ and $s \in (0, 1)$, then strict inequality holds by applying similar arguments as in [162, Proposition 1.2]. However, the strict inequality $G_a(\gamma_s) < G_a(\gamma_*)$ for $\gamma_* \neq \gamma'_*$ is a contradiction to the minimality of γ_* . Hence, the minimizer is unique. $\square$

Remark 5.31. If $\lambda_1 > 0$, then the strict convexity of f_1 can be used to deduce uniqueness, and the assumptions on $-\log \tilde{\rho}$ can be weakened from strict convexity to convexity.

Corollary 5.32. Let Assumptions 5.5-5.7 hold. Any minimizer $\gamma_* = (\rho_*, \mu_*)$ of G_a is a steady state for equation (5.4) according to Definition 5.23 and satisfies $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$.

Proof. By Proposition 5.30, we have $\rho_*, \mu_* \in \mathcal{P}_2$, as well as $\rho_* \in L_+^1$, $\|\rho_*\|_1 = 1$, and that ρ_* is absolutely continuous with respect to $\tilde{\rho}$. Since $W \in C^2(\mathbb{R}^d)$, it follows that $\nabla W * \rho_* \in L_{loc}^1$. In order to show that γ_* is a steady state for equation (5.4), it remains to prove that $\rho_* \in W_{loc}^{1,2} \cap L_{loc}^\infty$. As γ_* is a minimizer, it is in particular a critical point, and therefore satisfies equations (5.14). Rearranging, we obtain (for a possible different constant $c_1[\rho_*, \mu_*] \neq 0$) from (5.14a) that

$$\rho_*(z) = c_1[\rho_*, \mu_*] \tilde{\rho}(z) \exp \left[-\frac{1}{\alpha} \left(\int f_1(z, x) \mu_*(x) + W * \rho_*(z) \right) \right] \quad \text{on } \text{supp}(\rho_*). \quad (5.18)$$

Then for any compact set $K \subset \mathbb{R}^d$,

$$\sup_{z \in K} \rho_*(z) \leq c_1[\rho_*, \mu_*] \sup_{z \in K} \tilde{\rho}(z) \sup_{z \in K} \exp \left(-\frac{1}{\alpha} \left(\int f_1(z, x) \mu_*(x) \right) \right) \sup_{z \in K} \exp \left(-\frac{1}{\alpha} W * \rho_* \right).$$

As $f_1 \geq 0$ by Assumption 5.5 and $W \geq 0$ by Assumption 5.7, the last two terms are finite. The first supremum is finite thanks to continuity of $\tilde{\rho}$. Therefore $\rho_* \in L_{loc}^\infty$. To show that $\rho_* \in W_{loc}^{1,2}$, note that for any compact set $K \subset \mathbb{R}^d$, we have $\int_K |\rho_*(z)|^2 dz < \infty$ as a consequence of $\rho_* \in L_{loc}^\infty$. Moreover, defining $T[\gamma](z) := -\frac{1}{\alpha} \left(\int f_1(z, x) \mu(x) + W * \rho(z) \right) \leq 0$, we have

$$\begin{aligned} \int_K |\nabla \rho_*|^2 dz &= c_1[\rho_*, \mu_*]^2 \int_K |\nabla \tilde{\rho} + \tilde{\rho} \nabla T[\gamma_*]|^2 \exp(2T[\gamma_*]) dz \\ &\leq 2c_1[\rho_*, \mu_*]^2 \int_K |\nabla \tilde{\rho}|^2 \exp(2T[\gamma_*]) dz + 2c_1[\rho_*, \mu_*]^2 \int_K |\nabla T[\gamma_*]|^2 \tilde{\rho}^2 \exp(2T[\gamma_*]) dz, \end{aligned}$$

which is bounded noting that $\exp(2T[\gamma_*]) \leq 1$ and that $T[\gamma_*](\cdot)$, $\nabla T[\gamma_*](\cdot)$ and $\nabla \tilde{\rho}$ are in L_{loc}^∞ , where we used that $f_1, (\cdot, x), W(\cdot), \tilde{\rho}(\cdot) \in C^1(\mathbb{R}^d)$ by Assumptions 5.5-5.7. We conclude that $\rho_* \in W_{loc}^{1,2}$, and indeed (ρ_*, μ_*) solves (5.12) in the sense of distributions as a consequence of (5.14).

Next, we show that $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$ using again the relation (5.18). Firstly, note that $\text{supp}(\rho_*) \subset \text{supp}(\tilde{\rho})$ since ρ_* is absolutely continuous with respect to $\tilde{\rho}$. Secondly, we claim that $\exp \left[-\frac{1}{\alpha} \left(\int f_1(z, x) \mu_*(x) + W * \rho_*(z) \right) \right] > 0$ for all $z \in \mathbb{R}^d$. In other words, we claim that $\int f_1(z, x) \mu_*(x) < \infty$ and $W * \rho_*(z) < \infty$ for all $z \in \mathbb{R}^d$. Indeed, for the first term, fix any $z \in \mathbb{R}^d$ and choose $R > 0$ large enough such that $z \in B_R(0)$. Then, thanks to continuity of f_1 according to Assumption 5.5, we have

$$\int f_1(z, x) \mu_*(x) \leq \sup_{z \in B_R(0)} \int f_1(z, x) \mu_*(x) < \infty.$$

For the second term, note that by Assumption 5.7, we have for any $z \in \mathbb{R}^d$ and $\epsilon > 0$,

$$\begin{aligned} W(z) &\leq W(0) + \nabla W(z) \cdot z \leq W(0) + \frac{1}{2\epsilon} \|\nabla W(z)\|^2 + \frac{\epsilon}{2} \|z\|^2 \\ &\leq W(0) + \frac{D^2}{2\epsilon} (1 + \|z\|)^2 + \frac{\epsilon}{2} \|z\|^2 \leq W(0) + \frac{D^2}{\epsilon} + \left(\frac{D^2}{\epsilon} + \frac{\epsilon}{2} \right) \|z\|^2 \\ &= W(0) + \frac{D}{\sqrt{2}} + \sqrt{2}D \|z\|^2, \end{aligned}$$

where the last equality follows by choosing the optimal $\epsilon = \sqrt{2}D$. We conclude that

$$\begin{aligned} W * \rho_*(z) &\leq W(0) + \frac{D}{\sqrt{2}} + \sqrt{2}D \int \|z - \tilde{z}\|^2 \rho_*(\tilde{z}) \\ &\leq W(0) + \frac{D}{\sqrt{2}} + 2\sqrt{2}D \|z\|^2 + 2\sqrt{2}D \int \|\tilde{z}\|^2 \rho_*(\tilde{z}), \end{aligned} \quad (5.19)$$

which is finite for any fixed $z \in \mathbb{R}^d$ thanks to the fact that $\rho_* \in \mathcal{P}_2(\mathbb{R}^d)$. Hence, $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$. $\square$

Remark 5.33. *If we have in addition that $\tilde{\rho} \in L^\infty(\mathbb{R}^d)$, then the minimizer ρ_* of G_a is in $L^\infty(\mathbb{R}^d)$ as well. This follows directly by bounding the right-hand side of (5.18).*

The following inequality is referred to as HWI inequality and represents the key result to obtain convergence to equilibrium.

Proposition 5.34 (HWI inequality). *Define the dissipation functional*

$$D_a(\gamma) := \iint |\nabla_{x,z} \delta_\gamma G_a(z, x)|^2 d\gamma(z, x).$$

Assume $\alpha, \beta > 0$ and let λ_a as defined in (5.15). Let $\gamma_0, \gamma_1 \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $G_a(\gamma_0), G_a(\gamma_1), D_a(\gamma_0) < \infty$. Then

$$G_a(\gamma_0) - G_a(\gamma_1) \leq \overline{W}(\gamma_0, \gamma_1) \sqrt{D_a(\gamma_0)} - \frac{\lambda_a}{2} \overline{W}(\gamma_0, \gamma_1)^2 \quad (5.20)$$

Proof. For simplicity, consider γ_0, γ_1 that have smooth Lebesgue densities of compact support. The general case can be recovered using approximation arguments. Let $(\gamma_s)_{s \in [0,1]}$ denote a $\overline{W}$ -geodesic between γ_0, γ_1 . Following similar arguments as in [50] and [176, Section 5] and making use of the calculations in the proof of Proposition 5.28, we have

$$\left. \frac{d}{ds} G_a(\gamma_s) \right|_{s=0} \geq \iint \begin{bmatrix} \xi_1(z) \\ \xi_2(x) \end{bmatrix} \cdot \begin{bmatrix} (\nabla \phi(z) - z) \\ (\nabla \psi(x) - x) \end{bmatrix} d\gamma_0(z, x),$$

where

$$\begin{aligned}\xi_1[\gamma_0](z) &:= \alpha \nabla_z \log \left(\frac{\rho_0(z)}{\tilde{\rho}(z)} \right) + \int \nabla_z f_1(z, x) d\mu_0(x) + \int \nabla_z W(z - z') d\rho_0(z'), \\ \xi_2[\gamma_0](x) &:= \int \nabla_x f_1(z, x) d\rho_0(z) + \nabla_x V(x).\end{aligned}$$

Note that the dissipation functional can then be written as

$$D_a(\gamma_0) = \iint \left(|\xi_1(z)|^2 + |\xi_2(x)|^2 \right) d\gamma_0(z, x).$$

Using the double integral Cauchy-Schwarz inequality [200], we obtain

$$\begin{aligned}\left. \frac{d}{ds} G_a(\gamma_s) \right|_{s=0} &\geq - \left(\sqrt{\iint \left\| \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} \right\|_2^2 d\gamma_0} \right) \left(\sqrt{\iint \left\| \begin{bmatrix} \nabla \phi(z) - z \\ \nabla \psi(x) - x \end{bmatrix} \right\|_2^2 d\gamma_0} \right) \\ &= -\sqrt{D_a(\gamma_0)} \sqrt{\int \|\nabla \phi(z) - z\|^2 d\rho_0 + \int \|\nabla \psi(x) - x\|^2 d\mu_0} \\ &= -\sqrt{D_a(\gamma_0)} \overline{W}(\gamma_0, \gamma_1).\end{aligned}$$

Next, we compute a Taylor expansion of $G_a(\gamma_s)$ when considered as a function in s and use the bound on $\frac{d^2}{ds^2} G_a$ from (5.15):

$$\begin{aligned}G_a(\gamma_1) &= G_a(\gamma_0) + \left. \frac{d}{ds} G_a(\gamma_s) \right|_{s=0} + \int_0^1 (1-t) \left(\left. \frac{d^2}{ds^2} G_a(\gamma_s) \right|_{s=t} \right) dt \\ &\geq G_a(\gamma_0) - \sqrt{D_a(\gamma_0)} \overline{W}(\gamma_0, \gamma_1) + \frac{\lambda_a}{2} \overline{W}(\gamma_0, \gamma_1)^2.\end{aligned}$$

□

Remark 5.35. *The HWI inequality in Proposition 5.34 immediately implies uniqueness of minimizers for G_a in the set $\{\gamma \in \mathcal{P} \times \mathcal{P} : D_a(\gamma) < +\infty\}$. Indeed, if γ_0 is such that $D_a(\gamma_0) = 0$, then for any other γ_1 in the above set we have $G_a(\gamma_0) \leq G_a(\gamma_1)$ with equality if and only if $\overline{W}(\gamma_0, \gamma_1) = 0$.*

Corollary 5.36 (Generalized Log-Sobolev inequality). *Denote by γ_* the unique minimizer of G_a . With λ_a as defined in (5.15), any product measure $\gamma \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $G(\gamma), D_a(\gamma) < \infty$ satisfies*

$$D_a(\gamma) \geq 2\lambda_a G_a(\gamma | \gamma_*). \quad (5.21)$$

Proof. This statement follows immediately from Proposition 5.34. Indeed, let $\gamma_1 = \gamma_*$ and $\gamma_0 = \gamma$ in (5.20). Then

$$\begin{aligned}G_a(\gamma | \gamma_*) &\leq \overline{W}(\gamma, \gamma_*) \sqrt{D_a(\gamma)} - \frac{\lambda_a}{2} \overline{W}(\gamma, \gamma_*)^2 \\ &\leq \max_{t \geq 0} \left(\sqrt{D_a(\gamma)} t - \frac{\lambda_a}{2} t^2 \right) = \frac{D_a(\gamma)}{2\lambda_a}.\end{aligned}$$

□

Corollary 5.37 (Talagrand inequality). *Denote by γ_* the unique minimizer of G_a . With λ_a as defined in (5.15), it holds*

$$\overline{W}(\gamma, \gamma_*)^2 \leq \frac{2}{\lambda_a} G_a(\gamma | \gamma_*)$$

for any $\gamma \in \mathcal{P}_2 \times \mathcal{P}_2$ such that $G_a(\gamma) < \infty$.

Proof. This is also a direct consequence of Proposition 5.34 by setting $\gamma_0 = \gamma_*$ and $\gamma_1 = \gamma$. Then $G_a(\gamma_*) < \infty$ and $D_a(\gamma_*) = 0$, and the result follows. $\square$

Proof of Theorem 5.9. The entropy term $\int \rho \log \rho$ produces diffusion in ρ for the corresponding PDE in (5.4). As a consequence, solutions ρ_t to (5.4) and minimizers ρ^* for G_a have to be L^1 functions. As there is no diffusion for the evolution of μ_t , solutions may have a singular part. In fact, for initial condition $\mu_0 = \delta_{x_0}$, the corresponding solution will be of the form $\mu_t = \delta_{x(t)}$, where $x(t)$ solves the ODE (5.3) with initial condition x_0 . This follows from the fact that the evolution for μ_t is a transport equation (also see Section 5.A.1 for more details). Results (a) and (b) are the statements in Proposition 5.30, Corollary 5.32 and Corollary 5.37. To obtain (c), we differentiate the energy G_a along solutions γ_t to the equation (5.4):

$$\begin{aligned} \frac{d}{dt} G_a(\gamma_t) &= \int \delta_\rho G_a[\gamma_t](z) \partial_t \rho_t dz + \int \delta_\mu G_a[\gamma_t](x) \partial_t \mu_t dx \\ &= - \int \|\nabla_z \delta_\rho G_a[\gamma_t](z)\|^2 d\rho_t(z) - \int \|\nabla_x \delta_\mu G_a[\gamma_t](x)\|^2 d\mu_t(x) \\ &= -D_a(\gamma_t) \leq -2\lambda_a G_a(\gamma_t | \gamma_*), \end{aligned}$$

where the last bound follows from Corollary 5.36. Applying Gronwall's inequality, we immediately obtain decay in energy,

$$G_a(\gamma_t | \gamma_*) \leq e^{-2\lambda_a t} G_a(\gamma_0 | \gamma_*).$$

Finally, applying Talagrand's inequality (Corollary 5.37), the decay in energy implies decay in the product Wasserstein metric,

$$\overline{W}(\gamma_t, \gamma_*) \leq c e^{-\lambda_a t}$$

where $c > 0$ is a constant only depending on γ_0, γ_* and the parameter λ_a . $\square$

5.C Proof of Theorem 5.11

In the case of competing objectives, we rewrite the energy $G_c(\rho, x) : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \mapsto [-\infty, \infty]$ as follows:

$$G_c(\rho, x) = \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 - P(\rho),$$

where

$$P(\rho) := \alpha KL(\rho|\tilde{\rho}) + \frac{1}{2} \int \rho W * \rho.$$

Note that for any fixed $\rho \in \mathcal{P}$, the energy $G_c(\rho, \cdot)$ is strictly convex in x , and therefore has a unique minimizer. Define the best response by

$$b(\rho) := \underset{\bar{x}}{\operatorname{argmin}} G_c(\rho, \bar{x})$$

and denote $G_b(\rho) := G_c(\rho, b(\rho))$. We begin with auxiliary results computing the first variations of the best response b and then the different terms in $G_b(\rho)$ using Definition 5.1.

Lemma 5.38 (First variation of the best response). *The first variation of the best response of the classifier at ρ (if it exists) is*

$$\delta_\rho b[\rho](z) = -Q(\rho)^{-1} \nabla_x f_1(z, b(\rho)) \quad \text{for almost every } z \in \mathbb{R}^d,$$

where $Q(\rho) \succeq (\beta + \lambda_1 + \lambda_2) \mathbf{I}_d$ is a symmetric matrix, constant in z and x , defined as

$$Q(\rho) := \beta \mathbf{I}_d + \int \nabla_x^2 f_1(z, b(\rho)) d\rho(z) + \int \nabla_x^2 f_2(z, b(\rho)) d\bar{\rho}(z).$$

In particular, we then have for any $\psi \in C_c^\infty(\mathbb{R}^d)$ with $\int \psi dz = 0$ that

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left\| b[\rho + \epsilon\psi] - b[\rho] - \epsilon \int \delta_\rho b[\rho](z) \psi(z) dz \right\| = 0.$$

Proof. Let $\psi \in C_c^\infty(\mathbb{R}^d)$ with $\int \psi dz = 0$ and fix $\epsilon > 0$. Any minimizer of $G_c(\rho + \epsilon\psi, x)$ for fixed ρ must satisfy

$$\nabla_x G_c(\rho + \epsilon\psi, b(\rho + \epsilon\psi)) = 0.$$

Differentiating in ϵ , we obtain

$$\begin{aligned} & \int \delta_\rho \nabla_x G_c[\rho + \epsilon\psi, b(\rho + \epsilon\psi)] \psi(z) dz \\ & + \nabla_x^2 G_c(\rho + \epsilon\psi, b(\rho + \epsilon\psi)) \int \delta_\rho b[\rho + \epsilon\psi](z) \psi(z) dz = 0. \end{aligned} \tag{5.22}$$

Next, we explicitly compute all terms involved in (5.22). Computing the derivatives yields

$$\begin{aligned} \nabla_x G_c(\rho, x) &= \int \nabla_x f_1(z, x) d\rho(z) + \int \nabla_x f_2(z, x) d\bar{\rho}(z) + \beta(x - x_0) \\ \delta_\rho \nabla_x G_c[\rho, x](z) &= \nabla_x f_1(z, x) \\ \nabla_x^2 G_c(\rho, x) &= \int \nabla_x^2 f_1(z, x) d\rho(z) + \int \nabla_x^2 f_2(z, x) d\bar{\rho}(z) + \beta \mathbf{I}_d. \end{aligned}$$

Note that $\nabla_x^2 G_c$ is invertible by Assumption 5.5, which states that f_1 and f_2 have positive-definite Hessians. Inverting this term and substituting these expressions into (5.22) for $\epsilon = 0$ gives

$$\begin{aligned} \int \delta_\rho b[\rho](z) \psi(z) dz &= - \left[\beta \mathbf{I}_d + \int \nabla_x^2 f_1(z, b(\rho)) d\rho(z) \right. \\ &\quad \left. + \int \nabla_x^2 f_2(z, b(\rho)) d\bar{\rho}(z) \right]^{-1} \int \nabla_x f_1(z, b(\rho)) \psi(z) dz \\ &= - \int Q(\rho)^{-1} \nabla_x f_1(z, b(\rho)) \psi(z) dz. \end{aligned}$$

Finally, the lower bound on $Q(\rho)$ follows thanks to Assumption 5.5. $\square$

Remark 5.39. *If we include the additional assumption that $f_i \in C^3(\mathbb{R}^d \times \mathbb{R}^d; [0, \infty))$ for $i = 1, 2$, then the Hessian of $b[\rho]$ is well-defined. More precisely, the Hessian is given by*

$$\frac{d^2}{d\epsilon^2} b[\rho + \epsilon\psi]|_{\epsilon=0} = Q(\rho)^{-1} \left(\frac{d}{d\epsilon} Q(\rho + \epsilon\psi)|_{\epsilon=0} + \int \nabla_x^2 f_1(z, b[\rho]) \psi(z) dz \right) Q(\rho)^{-1} u[\rho, \psi]$$

where $u[\rho, \psi] = \int \nabla_x f_1(z, b[\rho]) \psi(z) dz$ and

$$\begin{aligned} \frac{d}{d\epsilon} Q_{ij}(\rho + \epsilon\psi)|_{\epsilon=0} &= - \int \partial_{x_i} \partial_{x_j} \nabla_x f_1(z, b[\rho]) \psi(z) \rho(z) dz Q(\rho)^{-1} u[\rho, \psi] \\ &\quad - \int \partial_{x_i} \partial_{x_j} \nabla_x f_2(z, b[\rho]) \psi(z) \bar{\rho}(z) dz Q(\rho)^{-1} u[\rho, \psi] + \int \partial_{x_i} \partial_{x_j} f_1(z, b[\rho]) \psi(z) dz. \end{aligned}$$

Therefore, we can Taylor expand $b[\rho]$ up to second order and control the remainder term of order ϵ^2 .

Lemma 5.40 (First variation of G_b). *The first variation of G_b is given by*

$$\delta_\rho G_b[\rho](z) = h_1(z) + h_2(z) + \beta h_3(z) - \delta_\rho P[\rho](z),$$

where

$$\begin{aligned} h_1(z) &:= \frac{\delta}{\delta\rho} \left(\int f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}) \right) (z) = \left\langle \int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}), \frac{\delta b}{\delta\rho}[\rho](z) \right\rangle + f_1(z, b(\rho)), \\ h_2(z) &:= \frac{\delta}{\delta\rho} \left(\int f_2(\tilde{z}, b(\rho)) d\bar{\rho}(\tilde{z}) \right) (z) = \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\bar{\rho}(\tilde{z}), \frac{\delta b}{\delta\rho}[\rho](z) \right\rangle, \\ h_3(z) &:= \frac{1}{2} \frac{\delta}{\delta\rho} \|b(\rho) - x_0\|^2 = \left\langle b(\rho) - x_0, \frac{\delta b}{\delta\rho}[\rho](z) \right\rangle, \end{aligned}$$

and

$$\delta_\rho P[\rho](z) = \alpha \log(\rho(z)/\bar{\rho}(z)) + (W * \rho)(z).$$

Proof. We begin with general expressions for Taylor expansions of $b : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ and $f_i(z, b(\cdot)) : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ for $i = 1, 2$ around ρ . Let $\psi \in \mathcal{T}$ with $\mathcal{T} = \{\psi : \int \psi(z) dz = 0\}$. Then

$$b(\rho + \epsilon\psi) = b(\rho) + \epsilon \int \frac{\delta b}{\delta \rho}[\rho](z')\psi(z')dz' + O(\epsilon^2) \quad (5.23)$$

and

$$f_i(z, b(\rho + \epsilon\psi)) = f_i(z, b(\rho)) + \epsilon \left\langle \nabla_x f_i(z, b(\rho)), \int \frac{\delta b}{\delta \rho}[\rho](z')\psi(z')dz' \right\rangle + O(\epsilon^2). \quad (5.24)$$

We compute explicitly each of the first variations:

(i) Using (5.24), we have

$$\begin{aligned} \int \psi(z) h_1(z) dz &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[\int f_1(z, b(\rho + \epsilon\psi))(\rho(z) + \epsilon\psi(z)) dz - \int f_1(z, b(\rho))\rho(z) dz \right] \\ &= \left\langle \int \nabla_x f_1(z, b(\rho)) d\rho(z), \int \frac{\delta b(\rho)}{\delta \rho}[\rho](z')\psi(z') dz' \right\rangle + \int f_1(z, b(\rho))\psi(z) dz \\ &= \int \left\langle \int \nabla_x f_1(z, b(\rho)) d\rho(z), \frac{\delta b(\rho)}{\delta \rho}[\rho](z') \right\rangle \psi(z') dz' + \int f_1(z, b(\rho))\psi(z) dz \\ &\Rightarrow h_1(z) = \left\langle \int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle + f_1(z, b(\rho)). \end{aligned}$$

(ii) Similarly, using again (5.24),

$$\begin{aligned} \int \psi(z) h_2(z) dz &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left[\int f_2(z, b(\rho + \epsilon\psi)) d\bar{\rho}(z) - \int f_2(z, b(\rho)) d\bar{\rho}(z) dz \right] \\ &= \int \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\bar{\rho}(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle \psi(z) dz \\ &\Rightarrow h_2(z) = \left\langle \int \nabla_x f_2(\tilde{z}, b(\rho)) d\bar{\rho}(\tilde{z}), \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle. \end{aligned}$$

(iii) Finally, from (5.23) it follows that

$$\begin{aligned} \int \psi(z) h_3(z) dz &= \lim_{\epsilon \rightarrow 0} \frac{1}{2\epsilon} \left[\langle b(\rho + \epsilon\psi) - x_0, b(\rho + \epsilon\psi) - x_0 \rangle - \langle b(\rho) - x_0, b(\rho) - x_0 \rangle \right] \\ &= \int \left\langle b(\rho) - x_0, \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle \psi(z) dz \\ &\Rightarrow h_3(z) = \left\langle b(\rho) - x_0, \frac{\delta b}{\delta \rho}[\rho](z) \right\rangle. \end{aligned}$$

Finally, the expression for $\delta_\rho P[\rho]$ follows by direct computation □

Lemma 5.41. Denote $G_b(\rho) := G_c(\rho, b(\rho))$ with $b(\rho)$ given by (5.5). Then $\delta_\rho G_b[\rho] = \delta_\rho G_c[\rho]|_{x=b(\rho)}$.

Proof. We start by computing $\delta_\rho G_c(\cdot, x)[\rho](z)$ for any $z, x \in \mathbb{R}^d$:

$$\delta_\rho G_c(\cdot, x)[\rho](z) = f_1(z, x) - \delta_\rho P[\rho](z). \quad (5.25)$$

Next, we compute $\delta_\rho G_b$. Using Lemma 5.40, the first variation of G_b is given by

$$\begin{aligned} \delta_\rho G_b[\rho](z) &= h_1(z) + h_2(z) + \beta h_3(z) - \delta_\rho P[\rho](z) \\ &= - \left\langle \left[\int \nabla_x f_1(\tilde{z}, b(\rho)) d\rho(\tilde{z}) + \int \nabla_x f_2(\tilde{z}, b(\rho)) d\bar{\rho}(\tilde{z}) + \beta(b(\rho) - x_0) \right], \delta_\rho b[\rho](z) \right\rangle \\ &\quad + f_1(z, b(\rho)) - \delta_\rho P[\rho](z). \end{aligned}$$

Note that

$$\nabla_x G_c(\rho, x) = \int \nabla_x f_1(\tilde{z}, x) d\rho(\tilde{z}) + \int \nabla_x f_2(\tilde{z}, x) d\bar{\rho}(\tilde{z}) + \beta(x - x_0), \quad (5.26)$$

and by the definition of the best response $b(\rho)$, we have $\nabla_x G_x(\rho, x)|_{x=b(\rho)} = 0$. Substituting into the expression for $\delta_\rho G_b$ and using (5.25), we obtain

$$\delta_\rho G_b[\rho](z) = f_1(z, b(\rho)) - \delta_\rho P[\rho](z) = \delta_\rho G_c(\cdot, x)[\rho](z) \Big|_{x=b(\rho)}.$$

This concludes the proof. $\square$

Lemma 5.42 (Uniform boundedness of the best response). *Let Assumption 5.5 hold. Then for any $\rho \in \mathcal{P}(\mathbb{R}^d)$, we have*

$$\|b(\rho)\|^2 \leq \|x_0\|^2 + \frac{2(a_1 + a_2)}{\beta}.$$

Proof. By definition of the best response $b(\rho)$, we have

$$\int \nabla_x f_1(z, b(\rho)) d\rho_t + \int \nabla_x f_2(z, b(\rho)) d\bar{\rho}(z) + \beta(b(\rho) - x_0) = 0.$$

To show that $b(\rho)$ is uniformly bounded, we take the inner product of the above expression with $b(\rho)$ itself

$$\beta \|b(\rho)\|^2 = \beta x_0 \cdot b(\rho) - \int \nabla_x f_1(z, b(\rho)) \cdot b(\rho) d\rho(z) - \int \nabla_x f_2(z, b(\rho)) \cdot b(\rho) d\bar{\rho}(z).$$

Using Assumption 5.5 to bound the two integrals, together with using Young's inequality to bound the first term on the right-hand side, we obtain

$$\beta \|b(\rho)\|^2 \leq \frac{\beta}{2} \|x_0\|^2 + \frac{\beta}{2} \|b(\rho)\| + a_1 + a_2,$$

which concludes the proof after rearranging terms. $\square$

Lemma 5.43 (Upper semi-continuity). *Let Assumptions 5.5-5.7 hold. The functional $G_c : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, +\infty]$ is upper semi-continuous when $\mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d$ is endowed with the product topology of the weak-* topology and the Euclidean topology. Moreover, the functional $G_b : \mathcal{P}(\mathbb{R}^d) \rightarrow [-\infty, +\infty]$ is upper semi-continuous with respect to the weak-* topology.*

Proof. The functional $G_c : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [-\infty, +\infty]$ is continuous in the second variable thanks to Assumption 5.5. Similarly, $\int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z)$ is continuous in ρ thanks to [194, Proposition 7.1] using the continuity of f_1 and f_2 . Further, $-P$ is upper semi-continuous using [184] and [194, Proposition 7.2] thanks to Assumptions 5.6 and 5.7. This concludes the continuity properties for G_c .

The upper semi-continuity of G_b then follows from a direct application of a version of Berge's maximum theorem [6, Lemma 16.30]. Let $R := \|x_0\|^2 + \frac{2(a_1+a_2)}{\beta} > 0$. We define $\varphi : (\mathcal{P}(\mathbb{R}^d), W_2) \rightarrow \mathbb{R}^d$ as the correspondence that maps any $\rho \in \mathcal{P}(\mathbb{R}^d)$ to the closed ball $\overline{B_R(0)} \subset \mathbb{R}^d$. Then the graph of φ is $\text{Gr } \varphi = \mathcal{P}(\mathbb{R}^d) \times \{\overline{B_R(0)}\}$. With this definition of φ , the range of φ is compact and φ is continuous with respect to weak-* convergence, and so it is in particular upper hemicontinuous. Thanks to Lemma 5.42, the best response function $b(\rho)$ is always contained in $\overline{B_R(0)}$ for any choice of $\rho \in \mathcal{P}(\mathbb{R}^d)$. As a result, maximizing $-G_c(\rho, x)$ in x over $\mathbb{R}^d$ for a fixed $\rho \in \mathcal{P}(\mathbb{R}^d)$ reduces to maximizing it over $\overline{B_R(0)}$. Using the notation introduced above, we can restrict G_c to $G_c : \text{Gr } \varphi \rightarrow \mathbb{R}$ and write

$$G_b(\rho) := \max_{\hat{x} \in \varphi(\rho)} -G_c(\rho, \hat{x}).$$

Because $G_c(\rho, x)$ is upper semi-continuous when $\mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d$ is endowed with the product topology of the weak-* topology and the Euclidean topology, [6, Lemma 16.30] guarantees that $G_b(\cdot)$ is upper semi-continuous in the weak-* topology. $\square$

Proposition 5.44. *Let $\alpha, \beta > 0$ and assume Assumptions 5.5-5.10 hold with the parameters satisfying $\alpha\tilde{\lambda} > \Lambda_1$. Fix $\rho_0, \rho_1 \in \mathcal{P}(\mathbb{R}^d)$. Along any geodesic $(\rho_s)_{s \in [0,1]} \in \mathcal{P}_2(\mathbb{R}^d)$ connecting ρ_0 to ρ_1 , we have for all $s \in [0, 1]$*

$$\frac{d^2}{ds^2} G_b(\rho_s) \leq -\lambda_b W_1(\rho_0, \rho_1)^2, \quad \lambda_b := \alpha\tilde{\lambda} - \Lambda_1, . \quad (5.27)$$

As a result, the functional $G_b : \mathcal{P}_2(\mathbb{R}^d) \rightarrow [-\infty, +\infty]$ is uniformly displacement concave with constant $\lambda_b > 0$.

Proof. Consider any $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$. Then any W_2 -geodesic $(\rho_s)_{s \in [0,1]}$ connecting ρ_0 with ρ_1 solves the following system of geodesic equations:

$$\begin{cases} \partial_s \rho_s + \text{div}(\rho_s v_s) = 0, \\ \partial_s(\rho_s v_s) + \text{div}(\rho_s v_s \otimes v_s) = 0, \end{cases} \quad (5.28)$$

where $\rho_s : \mathbb{R}^d \rightarrow \mathbb{R}$ and $v_s : \mathbb{R}^d \mapsto \mathbb{R}^d$. The first derivative of G_b along geodesics can be computed explicitly as

$$\begin{aligned} \frac{d}{ds} G_b(\rho_s) &= \int \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz - \frac{d}{ds} P(\rho_s) \\ &\quad + \left\langle \left[\int \nabla_x f_1(z, x) d\rho_s(z) + \int \nabla_x f_2(z, x) d\bar{\rho}(z) + \beta(x - x_0) \right] \Big|_{x=b(\rho_s)}, \frac{d}{ds} b(\rho_s) \right\rangle. \end{aligned}$$

The left-hand side of the inner product is zero by definition of the best response $b(\rho_s)$ to ρ_s , see (5.26). Therefore

$$\frac{d}{ds} G_b(\rho_s) = \int \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz - \frac{d}{ds} P(\rho_s).$$

Differentiating a second time, using (5.28) and integration by parts, we obtain

$$\frac{d^2}{ds^2} G_b(\rho_s) = L_1(\rho_s) + L_2(\rho_s) - \frac{d^2}{ds^2} P(\rho_s),$$

where

$$\begin{aligned} L_1(\rho_s) &:= \int \nabla_z^2 f_1(z, b(\rho_s)) \cdot (v_s \otimes v_s) \rho_s dz = \int \langle v_s, \nabla_z^2 f_1(z, b(\rho_s)) \cdot v_s \rangle \rho_s dz, \\ L_2(\rho_s) &:= \int \frac{d}{ds} b(\rho_s) \cdot \nabla_x \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) \rho_s(z) dz. \end{aligned}$$

From (5.16), we have that

$$\frac{d^2}{ds^2} \tilde{P}(\rho_s) \geq \alpha \tilde{\lambda} W_2(\rho_0, \rho_1)^2,$$

and thanks to Assumption 5.10 it follows that

$$L_1(s) \leq \Lambda_1 W_2(\rho_0, \rho_1)^2.$$

This leaves L_2 to bound; we first consider the term $\frac{d}{ds} b(\rho_s)$:

$$\begin{aligned} \frac{d}{ds} b(\rho_s) &= \int \delta_\rho b[\rho_s](\tilde{z}) \partial_s \rho_s(d\tilde{z}) = - \int \delta_\rho b[\rho_s](\tilde{z}) \operatorname{div}(\rho_s v_s) d\tilde{z} \\ &= \int \nabla_z \delta_\rho b[\rho_s](\tilde{z}) \cdot v_s(\tilde{z}) d\rho_s(\tilde{z}). \end{aligned}$$

Defining $u(\rho_s) \in \mathbb{R}^d$ by

$$u(\rho_s) := \int \nabla_x \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) d\rho_s(z),$$

using the results from Lemma 5.38 for $\nabla_z \delta_\rho b[\rho_s]$, Assumption 5.5 and the fact that $Q(\rho)$ is constant in z and x , we have

$$\begin{aligned} L_2(\rho_s) &= - \iint [Q(\rho_s)^{-1} \nabla_x \nabla_z f_1(\tilde{z}, b(\rho_s)) \cdot v_s(\tilde{z})] \cdot \nabla_x \nabla_z f_1(z, b(\rho_s)) \cdot v_s(z) d\rho_s(z) d\rho_s(\tilde{z}) \\ &= - \langle u(\rho_s), Q(\rho_s)^{-1} u(\rho_s) \rangle \leq 0 \end{aligned}$$

Combining all terms together, we obtain

$$\frac{d^2}{ds^2} G_b(\rho_s) \leq -(\alpha\tilde{\lambda} - \Lambda_1) W_2(\rho_0, \rho_1)^2.$$

□

Remark 5.45. Under some additional assumptions on the functions f_1 and f_2 , we can obtain an improved convergence rate. In particular, assume that for all $z, x \in \mathbb{R}^d$,

- there exists a constant $\Lambda_2 \geq \lambda_2 \geq 0$ such that $\nabla_x^2 f_2(z, x) \preceq \Lambda_2 \mathbf{I}_d$;
- there exists a constant $\sigma \geq 0$ such that $\|\nabla_x \nabla_z f_1(z, x)\| \geq \sigma$.

Then we have $-Q(\rho_s)^{-1} \preceq -1/(\beta + \Lambda_1 + \Lambda_2) \mathbf{I}_d$. Using Lemma 5.38, we then obtain a stronger bound on L_2 as follows:

$$\begin{aligned} L_2(\rho_s) &\leq -\frac{1}{\beta + \Lambda_1 + \Lambda_2} \|u(\rho_s)\|^2 \\ &\leq -\frac{1}{\beta + \Lambda_1 + \Lambda_2} \int \|\nabla_x \nabla_z f_1(z, b(\rho_s))\|^2 d\rho_s(z) \int \|v_s(z)\|^2 d\rho_s(z) \\ &\leq -\frac{\sigma^2}{\beta + \Lambda_1 + \Lambda_2} W_2(\rho_0, \rho_1)^2. \end{aligned}$$

This means we can improve the convergence rate in (5.27) to $\lambda_b := \alpha\tilde{\lambda} + \frac{\sigma^2}{\beta + \Lambda_1 + \Lambda_2} - \Lambda_1$.

Proposition 5.46 (Ground state). *Let Assumptions 5.5-5.10 hold for $\alpha\tilde{\lambda} > \Lambda_1 \geq 0$ and $\beta > 0$. Then there exists a unique maximizer ρ_* for the functional G_b over $\mathcal{P}(\mathbb{R}^d)$, and it satisfies $\rho_* \in \mathcal{P}_2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ and ρ_* is absolutely continuous with respect to $\tilde{\rho}$.*

Proof. Uniqueness of the maximizer (if it exists) is guaranteed by the uniform concavity provided by Lemma 5.44. To show existence of a maximizer, we use the direct method in the calculus of variations, requiring the following key properties for G_b : (1) boundedness from above, (2) upper semi-continuity, and (3) tightness of any minimizing sequence. To show (1), note that $\nabla_z^2(f_1(z, x) + \alpha \log \tilde{\rho}(z)) \preceq -(\alpha\tilde{\lambda} - \Lambda_1) \mathbf{I}_d$ for all $z, x \in \mathbb{R}^d \times \mathbb{R}^d$ by Assumptions 5.6 and 5.10, and so

$$f_1(z, x) + \alpha \log \tilde{\rho}(z) \leq c_0(x) - \frac{(\alpha\tilde{\lambda} - \Lambda_1)}{4} |z|^2 \quad \forall (z, x) \in \mathbb{R}^d \times \mathbb{R}^d \quad (5.29)$$

with $c_0(x) := f_1(0, x) + \alpha \log \tilde{\rho}(0) + \frac{1}{\alpha \tilde{\lambda} - \Lambda_1} \|\nabla_z [f_1(0, x) + \alpha \log \tilde{\rho}(0)]\|^2$. Therefore,

$$\begin{aligned} G_b(\rho) &= \int [f_1(z, b(\rho)) + \alpha \log \tilde{\rho}(z)] d\rho(z) + \int f_2(z, b(\rho)) d\bar{\rho}(z) + \frac{\beta}{2} \|b(\rho) - x_0\|^2 \\ &\quad - \alpha \int \rho \log \rho - \int \rho W * \rho \\ &\leq c_0(b(\rho)) + \int f_2(z, b(\rho)) d\bar{\rho}(z) + \frac{\beta}{2} \|b(\rho) - x_0\|^2. \end{aligned}$$

To estimate each of the remaining terms on the right-hand side, denote $R := \|x_0\|^2 + \frac{2(a_1+a_2)}{\beta}$ and recall that $\|b(\rho)\| \leq R$ for any $\rho \in \mathcal{P}(\mathbb{R}^d)$ thanks to Lemma 5.42. By continuity of f_1 and $\log \tilde{\rho}$, there exists a constant $c_1 \in \mathbb{R}$ such that

$$\sup_{x \in B_R(0)} c_0(x) = \sup_{x \in B_R(0)} \left[f_1(0, x) + \alpha \log \tilde{\rho}(0) + \frac{1}{\alpha \tilde{\lambda} - \Lambda_1} \|\nabla_z (f_1(0, x) + \alpha \log \tilde{\rho}(0))\|^2 \right] \leq c_1. \quad (5.30)$$

The second term is controlled by c_2 thanks to Assumption 5.10. And the third term can be bounded directly to obtain

$$G_b(\rho) \leq c_1 + c_2 + \beta(R^2 + \|x_0\|^2).$$

This concludes the proof of (1). Statement (2) was shown in Lemma 5.43. Then we obtain a minimizing sequence $(\rho_n) \in \mathcal{P}(\mathbb{R}^d)$ which is in the closed unit ball of $C_0(\mathbb{R}^d)^*$ and so the Banach-Alaoglu theorem [192, Theorem 3.15] there exists a limit ρ_* in the Radon measures and a subsequence (not relabeled) such that $\rho_n \xrightarrow{*} \rho_*$. In fact, ρ_* is absolutely continuous with respect to $\tilde{\rho}$ as otherwise $G_b(\rho_*) = -\infty$, which contradicts that $G_b(\cdot) > -\infty$ somewhere. We conclude that $\rho_* \in L^1(\mathbb{R}^d)$ since $\tilde{\rho} \in L^1(\mathbb{R}^d)$ by Assumption 5.6. To ensure $\rho_* \in \mathcal{P}(\mathbb{R}^d)$, we require (3) tightness of the minimizing sequence (ρ_n) . By Markov's inequality [117] it is sufficient to establish a uniform bound on the second moments:

$$\int \|z\|^2 d\rho_n(z) < C \quad \forall n \in \mathbb{N}. \quad (5.31)$$

To see this we proceed in a similar way as in the proof of Proposition 5.28. Defining

$$K(\rho) := - \int [f_1(z, b(\rho)) + \alpha \log \tilde{\rho}(z)] d\rho(z) + \alpha \int \rho \log \rho dz + \frac{1}{2} \int \rho W * \rho dz,$$

we have $K(\rho) = -G_b(\rho) + \int f_2(z, b(\rho)) d\bar{\rho}(z) + \frac{\beta}{2} \|b(\rho) - x_0\|^2$. Then using again the bound on $b(\rho)$ from Lemma 5.42,

$$\begin{aligned} K(\rho) &\leq -G_b(\rho) + \sup_{x \in B_R(0)} \int f_2(z, x) d\bar{\rho}(z) + \beta(R^2 + \|x_0\|^2) \\ &\leq -G_b(\rho) + c_2 + \beta(R^2 + \|x_0\|^2), \end{aligned}$$

where the last inequality is thanks to Assumption 5.10. Hence, using the estimates (5.29) and (5.30) from above, and noting that the sequence (ρ_n) is minimizing $(-G_b)$, we have

$$\begin{aligned} \frac{(\alpha\tilde{\lambda} - \Lambda_1)}{4} \int \|z\|^2 d\rho_n(z) &\leq c_0(b(\rho_n)) - \int [f_1(z, b(\rho_n)) + \alpha \log \tilde{\rho}(z)] d\rho_n(z) \\ &\leq c_1 + K(\rho_n) \leq c_1 - G_b(\rho_n) + c_2 + \beta \left(R^2 + \|x_0\|^2 \right) \\ &\leq c_1 - G_b(\rho_1) + c_2 + \beta \left(R^2 + \|x_0\|^2 \right) < \infty. \end{aligned}$$

which uniformly bounds the second moments of (ρ_n) . This concludes the proof for the estimate (5.31) and also ensures that $\rho_* \in \mathcal{P}_2(\mathbb{R}^d)$. □

Corollary 5.47. *Any maximizer ρ_* of G_b is a steady state for equation (5.5) according to Definition 5.24, and satisfies $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$.*

Proof. To show that ρ_* is a steady state we can follow exactly the same argument as in the proof of Corollary 5.32, just replacing $-\frac{1}{\alpha} \int f_1(z, x) \mu_*(x)$ with $+\frac{1}{\alpha} \int f_1(z, b(\rho_*))$. It remains to show that $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$. As ρ_* is a maximizer, it is in particular a critical point, and therefore satisfies that $\delta_{\rho} G_b[\rho_*](z)$ is constant on all connected components of $\text{supp}(\rho_*)$. Thanks to Lemma 5.41, this means there exists a constant $c[\rho_*]$ (which may be different on different components of $\text{supp}(\rho_*)$) such that

$$f_1(z, b(\rho_*)) - \alpha \log \left(\frac{\rho_*(z)}{\tilde{\rho}(z)} \right) - W * \rho_*(z) = c[\rho_*] \quad \text{on } \text{supp}(\rho_*).$$

Rearranging, we obtain (for a possible different constant $c[\rho_*] \neq 0$)

$$\rho_*(z) = c[\rho_*] \tilde{\rho}(z) \exp \left[\frac{1}{\alpha} (f_1(z, b(\rho_*)) - W * \rho_*(z)) \right] \quad \text{on } \text{supp}(\rho_*). \quad (5.32)$$

Firstly, note that $\text{supp}(\rho_*) \subset \text{supp}(\tilde{\rho})$ since ρ_* is absolutely continuous with respect to $\tilde{\rho}$. Secondly, note that $\exp \frac{1}{\alpha} f_1(z, b(\rho_*)) \geq 1$ for all $z \in \mathbb{R}^d$ since $f_1 \geq 0$. Finally, we claim that $\exp \left(-\frac{1}{\alpha} W * \rho_*(z) \right) > 0$ for all $z \in \mathbb{R}^d$. In other words, we claim that $W * \rho_*(z) < \infty$ for all $z \in \mathbb{R}^d$. This follows by exactly the same argument as in Corollary 5.32, see equation (5.19). We conclude that $\text{supp}(\rho_*) = \text{supp}(\tilde{\rho})$. □

Remark 5.48. *If we have in addition that $\tilde{\rho} \in L^\infty(\mathbb{R}^d)$ and $f_1(\cdot, x) \in L^\infty(\mathbb{R}^d)$ for all $x \in \mathbb{R}^d$, then the maximizer ρ_* of G_b is in $L^\infty(\mathbb{R}^d)$ as well. This follows directly by bounding the right-hand side of (5.32).*

With the above preliminary results, we can now show the HWI inequality, which implies again a Talagrand-type inequality and a generalized logarithmic Sobolev inequality.

Proposition 5.49 (HWI inequalities). *Define the dissipation functional*

$$D_b(\gamma) := \iint |\nabla_z \delta_\rho G_b[\rho](z)|^2 d\rho(z) .$$

Assume $\alpha, \beta > 0$ such that $\alpha\tilde{\lambda} > \Lambda_1 + \sigma^2$, and let λ_b as defined in (5.27). Denote by ρ_* the unique maximizer of G_b .

(HWI) Let $\rho_0, \rho_1 \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho_0), G_b(\rho_1), D_b(\rho_0) < \infty$. Then

$$G_b(\rho_0) - G_b(\rho_1) \leq \overline{W}(\rho_0, \rho_1) \sqrt{D_b(\rho_0)} - \frac{\lambda_b}{2} W_2(\rho_0, \rho_1)^2 \quad (5.33)$$

(logSobolev) Any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G(\rho), D_b(\rho) < \infty$ satisfies

$$D_b(\rho) \geq 2\lambda_b G_a(\rho | \rho_*) . \quad (5.34)$$

(Talagrand) For any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$ such that $G_b(\rho) < \infty$, we have

$$W_2(\rho, \rho_*)^2 \leq \frac{2}{\lambda_b} G_b(\rho | \rho_*) . \quad (5.35)$$

Proof. The proof for this result follows analogously to the arguments presented in the proofs of Proposition 5.34, Corollary 5.36 and Corollary 5.37, using the preliminary results established in Proposition 5.44 and Proposition 5.46. $\square$

Proof of Theorem 5.11. Following the same approach as in the proof of Theorem 5.9, the results in Theorem 5.11 immediately follow by combining Proposition 5.46, Corollary 5.47 and Proposition 5.49 applied to solutions of the PDE (5.5). $\square$

5.D Proof of Theorem 5.12

The proof for this theorem uses similar strategies as that of Theorem 5.11, but considers the evolution of an ODE rather than a PDE. Recall that for any $x \in \mathbb{R}^d$ the best response $r(x)(\cdot) \in \mathcal{P}(\mathbb{R}^d)$ in (5.6) is defined as

$$r(x) := \operatorname{argmax}_{\hat{\rho} \in \mathcal{P}} G_c(\hat{\rho}, x) ,$$

where the energy $G_c(\rho, x) : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \mapsto [-\infty, \infty]$ is given by

$$G_c(\rho, x) = \int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 - \alpha KL(\rho | \bar{\rho}) - \frac{1}{2} \int \rho W * \rho .$$

Lemma 5.50. *Let Assumptions 5.6- 5.10 hold and assume $\alpha\tilde{\lambda} > \Lambda_1$. Then for each $x \in \mathbb{R}^d$ there exists a unique maximizer $\rho_* := r(x)$ solving $\operatorname{argmax}_{\hat{\rho} \in \mathcal{P}_2} G_c(\hat{\rho}, x)$. Further, $r(x) \in L^1(\mathbb{R}^d)$, $\operatorname{supp}(r(x)) = \operatorname{supp}(\tilde{\rho})$, and there exists a function $c : \mathbb{R}^d \mapsto \mathbb{R}$ such that the best response $\rho_*(z) = r(x)(z)$ solves the Euler-Lagrange equation*

$$\delta_\rho G_c[\rho_*, x](z) := \alpha \log \rho_*(z) - (f_1(z, x) + \alpha \log \tilde{\rho}(z)) + (W * \rho_*)(z) = c(x) \quad (5.36)$$

for all $(z, x) \in \operatorname{supp}(\tilde{\rho}) \times \mathbb{R}^d$.

Proof. Equivalently, consider the minimization problem for $F(\rho) = -\int f_1(z, x) d\rho(z) + \alpha KL(\rho | \tilde{\rho}) + \frac{1}{2} \int \rho W * \rho$ with some fixed x . Note that we can rewrite $F(\rho)$ as

$$F(\rho) = \alpha \int \rho \log \rho dz + \int V(z, x) d\rho(z) + \frac{1}{2} \int \rho W * \rho$$

where $V(z, x) := -(f_1(z, x) + \alpha \log \tilde{\rho}(z))$ is strictly convex in z for fixed x by Assumptions 5.6 and 5.10. Together with Assumption 5.7, we can directly apply the uniqueness and existence result from [50, Theorem 2.1 (i)].

The result on the support of $r(x)$ and the expression for the Euler-Lagrange equation follows by exactly the same arguments as in Corollary 5.32 and Corollary 5.47. $\square$

Lemma 5.51. *The density of the best response $r(x)$ is continuous on $\mathbb{R}^d$ for any fixed $x \in \mathbb{R}^d$.*

Proof. Instead of solving the Euler-Lagrange equation (5.36), we can also obtain the best response $r(x)$ as the long-time asymptotics for the following gradient flow:

$$\partial_t \rho = \operatorname{div} (\rho \nabla \delta_\rho F[\rho]) . \quad (5.37)$$

Following Definitions 5.23 and 5.24, we can characterize the steady states ρ_∞ of the PDE (5.37) by requiring that $\rho_\infty \in L_+^1(\mathbb{R}^d) \cap L_{loc}^\infty(\mathbb{R}^d)$ with $\|\rho_\infty\|_1 = 1$ such that $\rho_\infty \in W_{loc}^{1,2}(\mathbb{R}^d)$, $\nabla W * \rho_\infty \in L_{loc}^1(\mathbb{R}^d)$, ρ_∞ is absolutely continuous with respect to $\tilde{\rho}$, and ρ_∞ satisfies

$$\nabla_z \left(-f_1(z, x) + \alpha \log \left(\frac{\rho_\infty(z)}{\tilde{\rho}(z)} \right) + W * \rho_\infty(z) \right) = 0 \quad \forall z \in \mathbb{R}^d, \quad (5.38)$$

in the sense of distributions. Noting that because the energy functional $F(\rho)$ differs from $G_a(\rho, \mu)$ only in the sign of $f_1(z, x)$ if viewing $G_a(\rho, \mu)$ as a function of ρ only. Note that $F(\rho)$ is still uniformly displacement convex in ρ due to Assumption 5.10. Then the argument to obtain that $\rho_\infty \in C(\mathbb{R}^d)$ follows exactly as that of Lemma 5.26. $\square$

Lemma 5.52. *Let $i \in \{1, \dots, d\}$. If the energy $H_i : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ given by*

$$H_i(\rho, x) := \frac{\alpha}{2} \int \frac{\rho(z)^2}{r(x)(z)} dz + \frac{1}{2} \int \rho W * \rho - \int \partial_{x_i} f_1(z, x) d\rho(z), \quad (5.39)$$

admits a critical point at $x \in \mathbb{R}^d$, then the best response $r(x) \in \mathcal{P}(\mathbb{R}^d)$ is differentiable in the i th coordinate direction at $x \in \mathbb{R}^d$. Further, the critical point of H_i is in the subdifferential $\partial_{x_i} r(x)$.

Proof. First, note that $DF[r(x)](x)(u) = 0$ for all directions $u \in C_c^\infty(\mathbb{R}^d)$ and for all $x \in \mathbb{R}^d$ thanks to optimality of $r(x)$. Here, DF denotes the Fréchet derivative of F , associating to every $\rho \in \mathcal{P}(\mathbb{R}^d)$ the bounded linear operator $DF[\rho] : C_c^\infty \rightarrow \mathbb{R}$

$$DF[\rho](u) := \int \delta_\rho F[\rho](z) u(z) dz,$$

and we note that $F(\rho)$ depends on x through the potential V . Fixing an index $i \in \{1, \dots, d\}$, and differentiating the optimality condition with respect to x_i we obtain

$$\partial_{x_i} DF[r(x)](x)(u) + D^2 F[r(x)](x)(u, \partial_{x_i} r(x)) = 0 \quad \forall u \in C_c^\infty(\mathbb{R}^d). \quad (5.40)$$

Both terms can be made more explicit using the expressions for the Fréchet derivative of F :

$$\partial_{x_i} DF[r(x)](x)(u) = - \int \partial_{x_i} f_1(z, x) u(z) dz,$$

and for the second term note that the second Fréchet derivative of F at $\rho \in \mathcal{P}(\mathbb{R}^d)$ along directions $u, v \in C_c^\infty(\mathbb{R}^d)$ such that $(\text{supp}(u) \cup \text{supp}(v)) \subset \text{supp}(\rho)$ is given by

$$D^2 F[\rho](x)(u, v) = \alpha \int \frac{u(z)v(z)}{\rho(z)} dz + \iint W(z - \tilde{z}) u(z) v(\tilde{z}) dz d\tilde{z}.$$

In other words, assuming $\text{supp}(r(x)) = \text{supp}(\tilde{\rho}) = \mathbb{R}^d$, relation (5.40) can be written as

$$\alpha \int \frac{\partial_{x_i} r(x)}{r(x)(z)} u(z) dz + \int (W * \partial_{x_i} r(x))(z) u(z) v dz - \int \partial_{x_i} f_1(z, x) u(z) dz = 0,$$

For ease of notation, given $r(x) \in \mathcal{P}(\mathbb{R}^d)$, we define the function $g : \mathcal{P}(\mathbb{R}^d) \rightarrow L_{loc}^1(\mathbb{R}^d)$ by

$$g[\rho](z) := \alpha \frac{\rho(z)}{r(x)(z)} + W * \rho - \partial_{x_i} f_1(z, x).$$

The question whether the partial derivative $\partial_{x_i} r(x)$ exists then reduces to the question whether there exists some $\rho_* \in \mathcal{P}(\mathbb{R}^d)$ such that $\rho = \rho_*$ solves the equation

$$g[\rho](z) = c \quad \text{for almost every } z \in \mathbb{R}^d.$$

and for some constant $c > 0$. This is precisely the Euler-Lagrange condition for the functional H_i defined in (5.39), which has a solution thanks to the assumption of Lemma 5.52. $\square$

We observe that the first term in H_i is precisely (up to a constant) the χ^2 -divergence with respect to $r(x)$,

$$\int \left(\frac{\rho}{r(x)} - 1 \right)^2 r(x) \, dz = \int \frac{\rho^2}{r(x)} \, dz - 1.$$

Depending on the shape of the best response $r(x)$, the χ^2 -divergence may not be displacement convex. Similarly, the last term $-\int \partial_{x_i} f_1(z, x) d\rho(z)$ in the energy H_i is in fact displacement concave due to the convexity properties of f_1 in z . The interaction term is displacement convex thanks to Assumption 5.7. As a result, the overall convexity properties of H_i are not known in general. Proving the existence of a critical for H_i under our assumptions on $f_1, f_2, \tilde{\rho}$ and W would be an interesting result in its own right, providing a new functional inequality that expands on the literature of related functional inequalities such as the related Hardy-Littlewood-Sobolev inequality [150].

It remains to show that H_i indeed admits a critical point. Next, we provide examples of additional assumptions that would guarantee for Lemma 5.52 to apply.

Lemma 5.53. *If either $C := \sup_{z \in \mathbb{R}^d} |W(z)| < \infty$, or*

$$C := \sup_{z \in \mathbb{R}^d} |\alpha \log(r(x)(z)/\tilde{\rho}(z)) + f_1(z, x) + c| < \infty,$$

then for each $x \in \mathbb{R}^d$ and for large enough $\alpha > 0$, the best response $r(x)$ is differentiable with the gradient coordinate $\partial_{x_i} r(x)$ given by the unique coordinate-wise solutions of the Euler-Lagrange condition for H_i .

Proof. We will show this result using the Banach Fixed Point Theorem for the mapping $T_i : L^1(\mathbb{R}^d) \rightarrow L^1(\mathbb{R}^d)$ for each fixed $i \in \{1, \dots, d\}$ given by

$$T_i(\rho) = -\frac{r(x)(z)}{\alpha} [(W * \rho)(z) - \partial_{x_i} f_1(z, x) + c],$$

noting that $\rho_* = T_i(\rho_*)$ is the Euler-Lagrange condition for a critical point of H_i . It remains to show that T_i is a contractive mapping. For the first assumption, note that

$$\begin{aligned} \|T_i(\rho) - T_i(\rho')\|_1 &= \frac{1}{\alpha} \int r(x) |W * (\rho - \rho')| \, dz \\ &\leq \frac{1}{\alpha} \iint r(x)(z) W(z - \hat{z}) |\rho(\hat{z}) - \rho'(\hat{z})| \, d\hat{z} \, dz \\ &\leq \frac{\|W\|_\infty}{\alpha} \left(\int r(x)(z) \, dz \right) \left(\int |\rho(\hat{z}) - \rho'(\hat{z})| \, d\hat{z} \right) \leq \frac{C}{\alpha} \|\rho - \rho'\|_1. \end{aligned}$$

Similarly, for the second assumption we estimate

$$\begin{aligned}
\|T_i(\rho) - T_i(\rho')\|_1 &= \frac{1}{\alpha} \int r(x) |W * (\rho - \rho')| dz \\
&\leq \frac{1}{\alpha} \iint r(x)(z) W(z - \hat{z}) |\rho(\hat{z}) - \rho'(\hat{z})| d\hat{z} dz \\
&= \frac{1}{\alpha} \int (W * r(x))(z) |\rho(z) - \rho'(z)| dz \\
&\leq \frac{1}{\alpha} \|W * r(x)\|_\infty \|\rho - \rho'\|_1
\end{aligned}$$

which requires a bound on $\|W * r(x)\|_\infty$. Using

$$\|W * r(x)\|_\infty = \sup_{z \in \mathbb{R}^d} |\alpha \log(r(x)/\tilde{\rho}) + f_1(z, x) + c| = C < \infty,$$

we conclude that T_i is a contraction map for large enough α . In both cases, we can then apply the Banach Fixed-Point Theorem to conclude that $\nabla_x r(x)$ exists and is unique. $\square$

Lemma 5.54. *Let $r(x)$ as defined in (5.6). If $r(x)$ is differentiable in x , then we have $\nabla_x G_d(x) = (\nabla_x G_c(\rho, x))|_{\rho=r(x)}$.*

Proof. We start by computing $\nabla_x G_d(x)$. We have

$$\begin{aligned}
\nabla_x G_d(x) &= \nabla_x (G_c(r(x), x)) = \int \delta_\rho [G_c(\rho, x)]|_{\rho=r(x)}(z) \nabla_x r(x)(z) dz + (\nabla_x G_c(\rho, x))|_{\rho=r(x)} \\
&= c(x) \nabla_x \int r(x)(z) dz + (\nabla_x G_c(\rho, x))|_{\rho=r(x)} = (\nabla_x G_c(\rho, x))|_{\rho=r(x)},
\end{aligned}$$

where we used that $r(x)$ solves the Euler-Lagrange equation (5.36) and that $r(x) \in \mathcal{P}(\mathbb{R}^d)$ for any $x \in \mathbb{R}^d$ so that $\int r(x)(z) dz$ is independent of x . $\square$

Lemma 5.55. *Let Assumption 5.5 hold. Then $G_d : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is strongly convex with constant $\lambda_d := \lambda_1 + \lambda_2 + \beta > 0$.*

Proof. The energy $G_c(\rho, x)$ is strongly convex in x due to our assumptions on f_1 , f_2 , and the regularizing term $\|x - x_0\|_2^2$. This means that for any $\rho \in \mathcal{P}$,

$$G_c(\rho, x) \geq G_c(\rho, x') + \nabla_x G_c(\rho, x')^\top (x - x') + \frac{\lambda_d}{2} \|x - x'\|_2^2.$$

Selecting $\rho = r(x')$, we have

$$G_c(r(x'), x) \geq G_c(r(x'), x') + \nabla_x G_c(r(x'), x')^\top (x - x') + \frac{\lambda_d}{2} \|x - x'\|_2^2.$$

Since $G_c(r(x'), x) \leq G_c(r(x), x)$ by definition of $r(x)$, we obtain the required convexity condition:

$$G_d(x) = G_c(r(x), x) \geq G_c(r(x'), x') + \nabla_x G_c(r(x'), x')^\top (x - x') + \frac{\lambda_d}{2} \|x - x'\|_2^2.$$

$\square$

Proof of Theorem 5.12. For any reference measure $\rho_0 \in \mathcal{P}$, we have

$$G_d(x) \geq G_c(\rho_0, x) \geq -\alpha KL(\rho_0 | \tilde{\rho}) - \frac{1}{2} \int \rho_0 W * \rho_0 + \frac{\beta}{2} \|x - x_0\|^2$$

and therefore, G_d is coercive. Together with the strong convexity provided by Lemma 5.55, we obtain the existence of a unique minimizer $x_\infty \in \mathbb{R}^d$. Convergence in norm now immediately follows also using Lemma 5.55: for solutions $x(t)$ to (5.6), we have

$$\frac{1}{2} \frac{d}{dt} \|x(t) - x_\infty\|^2 = -(G_d(x(t)) - G_d(x_\infty)) \cdot (x(t) - x_\infty) \leq -\lambda_d \|x(t) - x_\infty\|^2.$$

A similar result holds for convergence in entropy using the Polyák-Łojasiewicz convexity inequality

$$\frac{1}{2} \|\nabla G_d(x)\|_2^2 \geq \lambda_d (G_d(x) - G_d(x_\infty)),$$

which is itself a direct consequence of strong convexity provided in Lemma 5.55. Then

$$\frac{d}{dt} (G_d(x(t)) - G_d(x_\infty)) = \nabla_x G_d(x(t)) \cdot \dot{x}(t) = -\|\nabla_x G_d(x(t))\|^2 \leq -2\lambda_d (G_d(x(t)) - G_d(x_\infty)),$$

and so the result in Theorem 5.12 follows. $\square$

5.E Additional Simulation Results

We simulate a number of additional scenarios to illustrate extensions beyond the setting with provable guarantees and in the settings for which we have results but no numerical implementations in the main paper. First, we simulate the aligned objectives setting in one dimension, corresponding to (5.4). Then we consider two settings which are not covered in our theory: (1) the previously-fixed distribution $\bar{\rho}$ is also time varying, and (2) the algorithm does not have access to the full distributions of ρ and $\bar{\rho}$ and instead samples from them to update. Lastly, we illustrate a classifier with the population attributes in two dimensions, which requires a different finite-volume implementation [47, Section 2.2] than the one dimension version of the PDE due to flux in two dimensions.

5.E.1 Aligned Objectives

Here we show numerical simulation results for the aligned objectives case, where the population and distribution have the same cost function. In this setting, the dynamics are of the form

$$\begin{aligned} \partial_t \rho &= \operatorname{div} (\rho \nabla_z \delta_\rho G_a[\rho, \mu]) \\ &= \operatorname{div} \left(\rho \nabla_z \left(\int f_1(z, x) d\mu(x) + \alpha \log(\rho/\tilde{\rho}) + W * \rho \right) \right) \\ \frac{d}{dt} x &= -\nabla_x \left(\int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 \right) \end{aligned}$$

where f_1 and f_2 are as defined in section 5.4.1, and $W = \frac{1}{20}(1 + z)^{-1}$, a consensus kernel. Note that W does not satisfy Assumption 5.7, but we still observe convergence in the simulation. This is expected; in other works such as [50], the assumptions on W are relaxed and convergence results proven given sufficient convexity of other terms. The regularizer $\tilde{\rho}$ is set to ρ_0 , which models a penalty for the effort required of individuals to alter their attributes. The coefficient weights are $\alpha = 0.1$ and $\beta = 1$, with discretization parameters $dz = 0.1$, $dt = 0.01$.




Figure 5.4: The dynamics include a consensus kernel, which draws neighbors in z -space closer together. We see that the population moves to make the classifier performer better, as the two distributions become more easily separable by the linear classifier.

In Figure 5.4, we observe the strategic distribution separating itself from the stationary distribution, improving the performance of the classifier and also improving the performance of the population itself. The strategic distribution and classifier appear to be stationary by time $t = 40$.

5.E.2 Multiple Dynamical Populations

We also want to understand the dynamics when both populations are strategic and respond to the classifier. In this example, we numerically simulate this and in future work we hope to prove additional results regarding convergence. This corresponds to modeling the previously-fixed distribution $\bar{\rho}$ as time-dependent; let this distribution be $\tau \in \mathcal{P}_2$. We consider the case where ρ is competitive with x and τ is aligned with x , with dynamics given by

$$\begin{aligned} \partial_t \rho &= -\operatorname{div}(\rho \nabla_z (f_1(z, x) - \alpha \log(\rho/\tilde{\rho}) - W * \rho)) \\ \partial_t \tau &= \operatorname{div}(\tau \nabla_z (f_2(z, x) + \alpha \log(\tau/\tilde{\tau}) + W * \tau)) \\ \frac{d}{dt} x &= -\nabla_x \left(\int f_1(z, x) d\rho(z) + \int f_2(z, x) d\tau(z) + \frac{\beta}{2} \|x - x_0\|^2 \right). \end{aligned}$$

We use $W = 0$ and f_1, f_2 as in section 5.4.1 and the same discretization parameters as in Section 5.E.1. In Figure 5.5, we observe that the τ population moves to the right, assisting the classifier in maintaining accurate scoring. In contrast, ρ also moves to the right, rendering the right tail to be classified incorrectly, which is desirable for individuals in the ρ population but not




Figure 5.5: The population ρ aims to be classified with the τ population, while the classifier moves to delineate between the two. We observe that τ adjusts to improve the performance of the classifier while ρ competes against it. The distributions are plotted at time $t = 0$, corresponding to $\tilde{\rho}$ and $\tilde{\tau}$, and time $t = 20$, corresponding to ρ and τ .

desirable for the classifier. While we leave analyzing the long-term behavior mathematically for future work, the distributions and classifier appear to converge by time $t = 20$.

5.E.3 Sampled Gradients

In real-world applications of classifiers, the algorithm may not know the exact distribution of the population, relying on sampling to estimate it. In this section we explore the effects of the classifier updating based on an approximated gradient, which is computed by sampling the true underlying distributions ρ and $\bar{\rho}$. We use the same parameters for the population dynamics as in section 5.4.1, and for the classifier we use the approximate gradient

$$\nabla_x L(z, x_t) \approx \frac{1}{n} \sum_{i=1}^n (\nabla_x f_1(z_i, x_t) + \nabla_x f_2(\bar{z}_i, x_t)) + \beta(x_t - x_0), \quad z_i \sim \rho_t, \quad \bar{z}_i \sim \bar{\rho}_t.$$

First, we simulate the dynamics with the classifier and the strategic population updating at the same rate, using $\alpha = 0.05$, $\beta = 1$, and the same consensus kernel as used previously, with the same discretization parameters as in 5.E.1. In Figure 5.6, we observe no visual difference between the two results with $n = 4$ versus $n = 40$ samples, which suggests that not many samples are needed to estimate the gradient.

Next, we consider the setting where the classifier is best-responding to the strategic population.

Unlike the first setting, we observe in Figure 5.7 a noticeable difference between the evolution of ρ_t with $n = 4$ versus $n = 40$ samples. This is not surprising because optimizing with a very poor estimate of the cost function at each time step would cause x_t to vary wildly, and this method fails to take advantage of correct "average" behavior that gradient descent provides.




Figure 5.6: When the classifier is updating at the same rate as the population, we do not see a significant change in the evolution of both species, suggesting that as long as the gradient estimate for the classifier is correct on average, the estimate itself does not need to be particularly accurate.




Figure 5.7: When the classifier is best-responding to the population, we observe that using $n=4$ samples leads to different behavior for both the classifier and the population, compared with a more accurate estimate using $n=40$ samples.

5.E.4 Two-dimensional Distributions

In practice, individuals may alter more than one of their attributes in response to an algorithm, for example, both cancelling a credit card and also reporting a different income in an effort to change a credit score. We model this case with $z \in \mathbb{R}^2$ and $x \in \mathbb{R}^2$, and simulate the results for the setting where the classifier and the population are evolving at the same rate. While this setting is not covered in our theory, it interpolates between the two timescale extremes.

We consider the following classifier:

$$\begin{aligned} f_1(z, x) &= \frac{1}{2} \left(1 - \frac{1}{1 + \exp x^\top z} \right) \\ f_2(z, x) &= \frac{1}{2} \left(\frac{1}{1 + \exp x^\top z} \right) \end{aligned} \tag{5.41}$$

with $W = 0$. Again, the reference distribution $\tilde{\rho}$ corresponds to the initial shape of the distribution, instituting a penalty for deviating from the initial distribution. We use $\alpha = 0.5$ and $\beta = 1$ for the

penalty weights, run for $t = 4$ with $dt = 0.005$ and $dx = dy = 0.2$ for the discretization. In this case, the strategic population is competing with the classifier, with dynamics given by

$$\begin{aligned}\partial_t \rho &= -\operatorname{div}(\rho \nabla_z (f_1(z, x) - \alpha \log(\rho/\bar{\rho}))) \\ \frac{d}{dt}x &= -\nabla_x \left(\int f_1(z, x) d\rho(z) + \int f_2(z, x) d\bar{\rho}(z) + \frac{\beta}{2} \|x - x_0\|^2 \right)\end{aligned}$$




Figure 5.8: We use (5.41) for the classifier functions, using a Gaussian initial condition and regularizer for ρ . We see the distribution moving toward the region with higher probability of misclassification.

In Figure 5.8, we observe the strategic population increasing mass toward the region of higher probability of being labeled "1" while the true underlying label is zero, with the probability plotted at time $t = 4$. This illustrates similar behavior to the one-dimensional case, including the distribution splitting into two modes, which is another example of polarization induced by the classifier. Note that while in this example, $x \in \mathbb{R}^2$ and we use a linear classifier; we could have $x \in \mathbb{R}^d$ with $d > 2$ and different functions for f_1 and f_2 which yield a nonlinear classifier; our theory in the timescale-separated case holds as long as the convexity and smoothness assumptions on f_1 and f_2 are satisfied.

*Chapter 6***COMPUTING OPTIMAL TRANSPORT PLANS VIA MIN-MAX GRADIENT FLOWS**

L. Conger, F. Hoffmann, R. Baptista, and E. Mazumdar. Computing Optimal Transport Plans via Min-Max Gradient Flows. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 2210–2217, Dec. 2025. doi: 10.1109/CDC57313.2025.11312926. ISSN: 2576-2370

ABSTRACT

We pose the Kantorovich optimal transport problem as a min-max problem with a Nash equilibrium that can be obtained dynamically via a two-player game, providing a framework for approximating optimal couplings. We prove convergence of the timescale-separated gradient descent dynamics to the optimal transport plan, and implement the gradient descent algorithm with a particle method, where the marginal constraints are enforced weakly using the KL divergence, automatically selecting a dynamical adaptation of the regularizer. The numerical results highlight the different advantages of using the standard Kullback-Leibler (KL) divergence versus the reverse KL divergence with this approach, opening the door for new methodologies.

6.1 Introduction

Optimally transporting mass from one distribution μ to another ν is an important problem in many applications, ranging from sampling to numerically solving partial differential equations. In this work, we reformulate the Kantorovich optimal transport (OT) problem

$$\inf_{\rho \in \Gamma(\mu, \nu)} \iint c(x, y) d\rho(x, y),$$

as a min-max problem by rewriting the marginal constraints with KL divergences, the weight of which is maximized,

$$\inf_{\rho} \sup_{\Lambda \in \mathbb{R}} \iint c(x, y) d\rho(x, y) + \Lambda (\text{KL}(\rho_1 | \mu) + \text{KL}(\rho_2 | \nu)) .$$

In the above, $\Gamma(\mu, \nu)$ is the set of measures whose marginals are μ and ν , the function c denotes the optimal transport cost, and ρ_1 and ρ_2 are the marginals of the joint probability measure ρ . The Nash equilibrium of the min-max problem coincides with the minimizer of the OT problem, and this formulation allows us to develop a gradient flow particle method to dynamically approximate the optimal coupling.

Other works have proposed particle methods for solving OT problems, such as [178], which presents a finite-particle method for mass transport problems in the $\mathcal{W}$ metric, applied to particles diffusing on a sphere. Prior to that, [153] proposed a particle method for computing OT plans, and we build upon their numerical method by using a different, more efficient density estimation method and more advantageous particle initializations. [153] provides a Γ -convergence result that we extend to achieve a Danskin-type theorem (also known as the envelope theorem) for the best-response of the coupling to the soft constraint, when considering the min-max dynamics under timescale separation. We address an open question posed in that paper by proving that the energy functional is displacement convex, leading to gradient flow convergence guarantees.

Viewing the OT problem through the lens of a min-max flow in infinite dimensions allows us to build upon the results in [63], which provides conditions under which convergence of the min-max dynamics in the space of measures holds. In our setting, the energy functional defining the zero-sum game is no longer strongly convex in both players (with one player being the transport plan, and the other player being the regularization parameter enforcing the marginal constraints), and the estimates we prove here provide new results in this setting.

The most popular and best-developed method for computing OT plans is the Sinkhorn algorithm [73]. Our particle method is not expected to outperform state-of-the-art implementations of the Sinkhorn algorithm in computation time—due in large part to the amount of recent work on numerically tuning the Sinkhorn algorithm for better performance [8, 148]. The key contributions of our method

in comparison to Sinkhorn lies in the greater expressivity of our approximation to the optimal coupling. The Sinkhorn algorithm is a reweighting of points which approximate a probability density, which means the joint distribution achieved is constrained to have support on the samples from the input marginals. Because we use particles instead, we can construct joint distributions with different support than the samples from the target marginals, which is expected to yield more expressivity especially in high dimensions. Further, the entropy regularization term in the Sinkhorn algorithm encourages a joint density which factorizes with increasing regularization, but ours does not, allowing for the possibility of obtaining a solution closer to the true optimal plan, while still maintaining the computational benefits of regularization.

Recent works have explored the connection between the Sinkhorn algorithm and Wasserstein-2 gradient flows, including [132], which develops a continuous time Sinkhorn flow, encompassing the setting of [79], which formulates the notion of Sinkhorn as a Wasserstein-2 mirror descent and provides convergence results. [59] gives conditions on the cost function for exponential convergence of the dual sub-optimality gap for population-level Sinkhorn. Our work is complimentary in that we do not have the KL regularization term which encourages factorization, but we do have a KL penalty term on the marginals which, when taken to have infinite weight, enforces the marginal constraints exactly and guarantees convergence of the energy to the true optimal cost.

Other approaches for approximating the OT plan exist, with a vast literature. Alternative parametric methods learn potential functions in the OT problem, which define a parametric approximation to the transport plan [197, 137]. Works which compute OT maps using the Benamou-Brenier formulation of the Wasserstein-2 distance include [179], which provides a proximal splitting algorithm, and shows equivalence with the alternating direction method of multipliers. Similarly, [149] solves the discretized Schrödinger’s bridge problem and [48] provides a method for discretizing a gradient flow in time and solves the resulting subproblems in a way that could instead be applied to an OT problem with the Benamou-Brenier formulation. Our particle method and corresponding theoretical results take the approach of relaxing both marginal constraints simultaneously with convergence guarantees as the constraints are more tightly enforced.

Contributions: We pose the optimal transport problem as a min-max problem, allowing us to prove convergence of a timescale-separated min-max flow to the optimal transport plan, in infinite dimensions. To do so we draw on ideas and make connections to singularly perturbed dynamical systems and two-timescale dynamics [196]. In particular, we generalize ideas from that literature to infinite dimensions. Along the way, we address an open question in [153] by proving displacement convexity of the energy functional and convergence of the resulting gradient flow to the minimizer.

Our analysis allows us to derive a new particle-based method for approximating optimal transport plans that has *convergence guarantees* and which could ultimately be a powerful tool in the optimal

transport toolbox. Numerically, we illustrate how the L^2 marginal density error can be improved if the shape of the source and target are taken into account when selecting forward KL (mode-seeking) vs reverse KL (moment-matching) for the energy functional when enforcing the soft constraints on the marginals.

6.2 Problem Setup

6.2.1 Notation

Let $\mathcal{P}_2(\mathbb{R}^d)$ denote the space of probability measures with finite second moments on $\mathbb{R}^d$. We use the notation $\rho(x)dx$ for measures with or without Lebesgue density. For two measures $\mu \in \mathcal{P}(\mathbb{R}^{d_x})$ and $\nu \in \mathcal{P}(\mathbb{R}^{d_y})$, we denote by $\Gamma(\mu, \nu)$ the set of all joint distributions whose marginals are equal to μ and ν , i.e.,

$$\Gamma(\mu, \nu) = \left\{ \gamma \in \mathcal{P}(\mathbb{R}^{d_x} \times \mathbb{R}^{d_y}) \mid \mu = \pi_{1\#}\gamma, \nu = \pi_{2\#}\gamma \right\},$$

where $\pi_{i\#}\gamma$ denotes the pushforward of γ on the i^{th} marginal. The transport cost function is $c : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow [0, \infty)$, and if $c(x, y) = \|x - y\|^2$ is the quadratic cost, we denote the corresponding Wasserstein metric by $\mathcal{W}$, i.e.,

$$\mathcal{W}(\mu, \nu)^2 = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|x - y\|^2 d\gamma(x, y).$$

Then $\Gamma^*(\mu, \nu)$ is the set of optimal couplings $\gamma \in \Gamma(\mu, \nu)$ achieving the inf above. For a sequence of measures $(\rho^{(n)})$, narrow convergence is defined as convergence in duality with all continuous bounded functions and denoted $\rho^{(n)} \xrightarrow{*} \bar{\rho}$. Convergence in $\mathcal{W}$ is equivalent to narrow convergence plus convergence of the second moments of $\rho^{(n)}$.

6.2.2 Displacement Convexity

Definition 6.1 (Displacement Convexity). *A functional $G : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ is λ -displacement convex if for all ρ_0, ρ_1 that are atomless we have*

$$G(\rho_s) \leq (1 - s)G(\rho_0) + sG(\rho_1) + \frac{\lambda}{2} \mathcal{W}(\rho_0, \rho_1)^2,$$

where $\rho_s := ((1 - s)x + sy)_{\#}\gamma$ for $\gamma \in \Gamma^*(\rho_0, \rho_1)$.

When $\lambda > 0$, λ -displacement convexity is analogous to strong convexity for functions on Euclidean space, with the inequality using the Wasserstein-2 distance function instead of the two-norm. A sufficient condition for uniform displacement convexity with constant λ is

$$\int \langle x - x', \nabla_x \delta_\rho G[\rho](x) - \nabla_x \delta_\rho G[\bar{\rho}](x') \rangle d\gamma(x, x') \geq \lambda \mathcal{W}(\rho, \bar{\rho})^2 \quad (6.1)$$

for all $\gamma \in \Gamma^*(\rho_0, \rho_1)$. For details regarding convexity inequalities, see [203, Proposition 16.4].

6.2.3 Optimal Transport Problem

The Kantorovich formulation of the optimal transport problem, parameterized by a cost function $c : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}_{\geq 0}$, is solved by finding a joint distribution $\rho \in \mathcal{P}_2(\mathbb{R}^{d_x} \times \mathbb{R}^{d_y})$ (i.e., a coupling) which minimizes

$$\inf_{\rho \in \Gamma(\mu, \nu)} \iint c(x, y) d\rho(x, y). \quad (6.2)$$

The minimizer $\rho^* \in \Gamma^*(\mu, \nu)$ is the transport plan which moves mass between μ and ν in an optimal fashion. Rather than enforcing the marginal constraints exactly, we can solve this problem approximately by solving the relaxed problem

$$\inf_{\rho \in \mathcal{P}_2} E(\rho, \Lambda; \mu, \nu), \quad (6.3)$$

where E is the energy functional

$$E(\rho, \Lambda; \mu, \nu) = \iint c(x, y) d\rho(x, y) + \Lambda \text{KL}(\rho_1 | \mu) + \Lambda \text{KL}(\rho_2 | \nu),$$

with ρ_1 and ρ_2 the marginals of ρ , defined as

$$\rho_1(x) = \int \rho(x, y) dy, \quad \rho_2(y) = \int \rho(x, y) dx,$$

and $\text{KL}(\cdot | \cdot)$ is the Kullback-Leibler divergence, defined as $\text{KL}(\mu | \nu) := \int \mu(x) \log(\mu(x)/\nu(x)) dx$ if $\mu \ll \nu$ and $+\infty$ otherwise. For $\mu, \nu \in \mathcal{P}_2$, the KL divergence is always non-negative, and $\mu = \nu$ a.e. if and only if $\text{KL}(\mu | \nu) = 0$. We also consider the related energy functional

$$\tilde{E}(\rho, \Lambda; \mu, \nu) = \iint c(x, y) d\rho(x, y) + \Lambda \text{KL}(\mu | \rho_1) + \Lambda \text{KL}(\nu | \rho_2)$$

with the arguments in the KL-divergence reversed. We experiment with this choice of energy functional in our numerical results section, showing that it may have some advantages compared to the standard KL divergence. The optimizer of the unconstrained problem (6.3) is not the same as the solution to (6.2). However, the min-max optimization problems

$$\inf_{\rho \in \mathcal{P}_2} \sup_{\Lambda \in \mathbb{R}} E(\rho, \Lambda; \mu, \nu), \quad (6.4)$$

$$\inf_{\rho \in \mathcal{P}_2} \sup_{\Lambda \in \mathbb{R}} \tilde{E}(\rho, \Lambda; \mu, \nu). \quad (6.5)$$

have Nash equilibria given by the solution to (6.2). While both (6.4) and (6.5) have the same Nash equilibria, the KL divergence terms in (6.5) are not necessarily convex in ρ , even if μ and ν are log-concave. In order to solve (6.2) numerically, we will solve (6.4) using a min-max gradient flow

with respect to the metric $\mathcal{W}_2$ as

$$\begin{aligned}\partial_t \rho(t) &= -\frac{1}{\beta} \nabla_{\mathcal{W}, \rho} E(\rho(t), \Lambda(t); \mu, \nu) \\ \frac{d}{dt} \Lambda(t) &= \partial_\Lambda E(\rho(t), \Lambda(t); \mu, \nu),\end{aligned}\tag{6.6}$$

where $\nabla_{\mathcal{W}, \rho}$ denotes the gradient with respect to ρ in the $\mathcal{W}$ metric, and $0 < \beta < 1$ is a temperature parameter specifying how fast ρ updates relative to Λ . When β is close to zero, (6.6) is a two-timescale dynamical system.

We will prove convergence in the best-response setting, that is, when $\beta \rightarrow 0$. In this setting, ρ is instantly minimizing $E(\cdot, \Lambda; \mu, \nu)$, and the dynamics are given by

$$\begin{aligned}\frac{d}{dt} \Lambda(t) &= \partial_\Lambda E(b[\Lambda], \Lambda; \mu, \nu) \\ b[\Lambda] &:= \operatorname{argmin}_{\rho \in \mathcal{P}_2} E(\rho, \Lambda; \mu, \nu).\end{aligned}\tag{6.7}$$

The minimizer of $E(\cdot, \Lambda; \mu, \nu)$ was proven to exist for any fixed $\Lambda > 0$ in [153, Corollary 1], and therefore $b[\Lambda]$ is well-defined for any $0 < \Lambda < +\infty$. The other timescale-separation regime, when Λ instantly best-responds, is no longer an entropy-regularized problem because as $\Lambda \rightarrow +\infty$, the optimization functional E Γ -converges to the optimal transport problem (6.2) [153, Theorem 7].

6.3 Results

We use the following assumptions for the proofs of our key results; relaxation is likely possible, and we later show numerical examples that do not satisfy them.

Assumption 6.2 (Admissibility). *The cost functional $c(x, y) : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}_{\geq 0}$ is lower semicontinuous (lsc), and input measures $\mu \in \mathcal{P}_2(\mathbb{R}^{d_x})$ and $\nu \in \mathcal{P}_2(\mathbb{R}^{d_y})$ are such that $\iint c(x, y) d\mu(x) d\nu(y) < \infty$.*

Assumption 6.3 (Convexity). *The cost function c and input measures $\mu \in \mathcal{P}_2(\mathbb{R}^{d_x})$, $\nu \in \mathcal{P}_2(\mathbb{R}^{d_y})$ satisfy*

- $\operatorname{Hess}(c(x, y)) \succeq \tilde{\lambda} \mathbf{I}$ for all $x \in \mathbb{R}^{d_x}$, $y \in \mathbb{R}^{d_y}$,
- $\nabla_x^2 \log \mu(x) \preceq -\lambda \mathbf{I}$ for all $x \in \mathbb{R}^{d_x}$, and
- $\nabla_y^2 \log \nu(y) \preceq -\lambda \mathbf{I}$ for all $y \in \mathbb{R}^{d_y}$.

Assumption 6.4 (Existence of unique minimizer). *The functional $E(\cdot, \Lambda; \mu, \nu)$ has a unique minimizer for every $\Lambda > 0$.*

Assumption 6.4 is more general than Assumptions 6.2-6.3; for example, when Assumptions 6.2-6.3 are satisfied with $c(x, y) = h(x - y)$, then Assumption 6.4 follows directly [153, Corollary 1].

6.3.1 Displacement Convexity of Energy

For a fixed $\Lambda > 0$, consider the dynamics

$$\partial_t \rho(t) = -\nabla_{\mathcal{W}, \rho} E(\rho(t), \Lambda; \mu, \nu). \quad (6.8)$$

The gradient flow (6.8) with a fixed Λ was proposed in [153] as a method for approximately solving the optimal transport problem (6.2), together with a numerical scheme for (6.8). The displacement convexity and convergence to a steady state of the flow (6.8) were left as an open question. In this section, we provide proofs for both. In fact, the evolution of the marginals ρ_1, ρ_2 for ρ solving (6.8) is not a closed system, unless $\rho(x, y)$ factors. If ρ does factor, we can think of ρ_1, ρ_2 as solving an infinite dimensional two-player game (the aligned case in [63]) with the same energy $E(\rho_1 \rho_2, \Lambda; \mu, \nu)$. First we will prove displacement convexity of E with respect to ρ , and then convergence of the dynamics (6.8) to a unique steady state. To do so, we will first prove displacement convexity of $\mathcal{H}_1(\rho)$ and $\mathcal{H}_2(\rho)$, where

$$\mathcal{H}_i(\rho) := \int \pi_{i\#} \rho(x) \log \pi_{i\#} \rho(x) dx = \int \rho_i(x) \log \rho_i(x) dx.$$

Lemma 6.5. *The energy functionals $\mathcal{H}_1$ and $\mathcal{H}_2$ are zero-displacement-convex in the space of absolutely continuous measures $\mathcal{P}^{ac}$.*

Proof. In order to show that $\mathcal{H}_1(\rho)$ is displacement convex in ρ , we will show that

$$\begin{aligned} & \iint \left(\langle x - x', \nabla_x \delta_\rho \mathcal{H}_1[\rho](x, y) - \nabla_x \delta_\rho \mathcal{H}_1[\tilde{\rho}](x', y') \rangle \right. \\ & \left. + \langle y - y', \nabla_y \delta_\rho \mathcal{H}_1[\rho](x, y) - \nabla_y \delta_\rho \mathcal{H}_1[\tilde{\rho}](x', y') \rangle \right) d\gamma(z) \geq 0, \end{aligned}$$

for all absolutely continuous $\rho, \tilde{\rho} \in \mathcal{P}_2(\mathbb{R}^{d_x} \times \mathbb{R}^{d_y})$, with γ an optimal coupling between ρ and $\tilde{\rho}$, and $z := [x, x', y, y']$. The first variation of $\mathcal{H}_1(\rho)$ with respect to ρ is

$$\delta_\rho \mathcal{H}_1(\rho)(x, y) = \log \rho_1(x),$$

and therefore the second inner product in the convexity condition is zero. The convexity condition becomes

$$\int \langle x - x', \nabla_x (\log(\rho_1(x)) - \log(\tilde{\rho}_1(x'))) \rangle d\gamma(z) \geq 0.$$

The above holds because $\int \rho_1 \log \rho_1$ is zero-displacement convex with respect to ρ_1 ; therefore, $\mathcal{H}_1$ is also zero-displacement convex with respect to ρ . The result follows similarly for $\mathcal{H}_2(\rho)$. $\square$

From now on and for the remainder of Section 6.3, let Assumption 6.2 hold. Assumption 6.3 is needed for Subsection 6.3.1 and Assumption 6.4 is crucial for Sections 6.3.2 and 6.3.3.

Theorem 6.6 (Displacement Convexity). *Let Assumptions 6.2-6.3 hold. The energy functional $E(\cdot, \Lambda; \mu, \nu)$ is $\tilde{\lambda} + \Lambda\lambda$ displacement convex.*

Proof. The term $\iint c(x, y) d\rho(x, y)$ is $\tilde{\lambda}$ -displacement convex in ρ . To check the entropy regularization terms, we will show the convexity of $\text{KL}(\rho_1 | \mu) + \text{KL}(\rho_2 | \nu)$. Since

$$\text{KL}(\rho_1 | \mu) + \text{KL}(\rho_2 | \nu) = \mathcal{H}_1(\rho) + \mathcal{H}_2(\rho) - \iint \log(\mu(x)\nu(y)) d\rho(x, y),$$

and by Lemma 6.5, the first two terms are zero-displacement convex, it remains to check the last integral. By Assumption 6.3, the function $-\log(\mu(x)\nu(y))$ is strongly λ -convex in (x, y) , and therefore the integral is λ -displacement convex in ρ . $\square$

Theorem 6.6 is the key result that provides convergence guarantees for dynamics (6.8).

Theorem 6.7 (Convergence). *Let Assumptions 6.2-6.3 hold with $\tilde{\lambda} + \Lambda\lambda > 0$. There exists a unique minimizer $\bar{\rho}$ of (6.3), and any solution $\rho(t)$ to (6.8) satisfies*

$$\mathcal{W}(\rho(t), \bar{\rho}) \leq e^{-(\tilde{\lambda} + \Lambda\lambda)t} \mathcal{W}(\rho(t), \rho(0)),$$

Proof. Existence of $\bar{\rho}$ was shown in [153, Corollary 1] and uniqueness follows from strong convexity of E (Theorem 6.6). Further, E is lsc and non-negative since the KL-divergences and $\int c(x, y) d\rho(x, y)$ are each lsc and non-negative. Differentiating $\mathcal{W}(\rho(t), \bar{\rho})$ in time via [56, Thm 3.14] together with $(\tilde{\lambda} + \Lambda\lambda)$ -displacement convexity (Theorem 6.6) allows us to conclude the estimate. $\square$

6.3.2 Danskin's Theorem

Let Assumptions 6.2 and 6.4 hold for this section. Dropping μ, ν from the notation for brevity, denote $E_d(\Lambda) := E(b[\Lambda], \Lambda)$. In this section, we prove that

$$\frac{d}{d\Lambda} E_d(\Lambda) = \partial_\Lambda E(\rho, \Lambda)|_{\rho=b[\Lambda]},$$

that is, the gradient through the best response is equal to the gradient with the best response plugged in after differentiating, which is classically known for finite-dimensional systems as Danskin's theorem. This results in dynamics of the form

$$\frac{d}{dt} \Lambda(t) = \text{KL}(b_1[\Lambda(t)] | \mu) + \text{KL}(b_2[\Lambda(t)] | \nu),$$

which allows us to analyze the long-time behavior of Λ in this timescale separated setting. Denoting

$$V(\Lambda) := \text{KL}(b_1[\Lambda] \mid \mu) + \text{KL}(b_2[\Lambda] \mid \nu), \quad (6.9)$$

we can write

$$E_d(\Lambda) = \iint c(x, y) db[\Lambda](x, y) + \Lambda V(\Lambda),$$

and $\dot{\Lambda} = V(\Lambda)$. Danskin's theorem crucially relies on continuity of $V(\cdot)$ and $b(\cdot)$; for V this can be shown directly, and for b we prove this by showing convergence of second moments of $b[\Lambda_n]$ for converging sequences $(\Lambda_n)_n$ and proving Γ -convergence of $E(\cdot, \Lambda_n)$ w.r.t. $\mathcal{W}$.

Lemma 6.8 (Marginals of best response). *Fix $\Lambda > 0$ and denote $Z[\Lambda] := \iint e^{-c(x, y)/\Lambda} d\mu(x) d\nu(y)$. Then the product of the marginals of the best response $b[\Lambda]$ satisfies*

$$b_1[\Lambda](x)b_2[\Lambda](y) = \frac{1}{Z[\Lambda]} e^{-c(x, y)/\Lambda} \mu(x)\nu(y) =: \sigma[\Lambda](x, y), \quad (6.10)$$

where the normalization constant $Z = Z[\Lambda]$ is bounded,

$$0 < e^{-c^*/\Lambda} \leq Z[\Lambda] \leq 1 \quad \forall \Lambda > 0, \quad (6.11)$$

with c^* the optimal value of (6.2).

Proof. Expression (6.10) directly follows from the Euler-Lagrange condition for $E(\cdot, \Lambda)$ together with the mass constraint $\iint b_1[\Lambda](x)b_2[\Lambda](y) dx dy = 1$. From Assumption 6.2, the optimal transport problem (6.2) admits a solution $\rho^* \in \mathcal{P}_2$ with finite cost [9, Ch 6]. Let

$$c^* := \inf_{\rho \in \Gamma(\mu, \nu)} \iint c(x, y) d\rho(x, y) < \infty.$$

Then the optimal value of (6.3) is upper-bounded by c^*

$$\inf_{\rho \in \mathcal{P}_2} E(\rho, \Lambda) \leq c^*,$$

since the OT minimizer ρ^* for (6.2) results in $E(\rho^*, \Lambda) = c^*$. Rearranging (6.10), we have a.e. on $\text{supp}(b[\Lambda])$,

$$c(x, y) + \Lambda \left(\log \frac{b_1[\Lambda]}{\mu} + \log \frac{b_2[\Lambda]}{\nu} \right) = -\Lambda \log Z[\Lambda].$$

Integrating against $b[\Lambda]$ and using the previous estimate,

$$-\Lambda \log Z[\Lambda] = E_d(\Lambda) \leq c^*.$$

□

As a direct consequence of Lemma 6.8, we obtain explicit expressions for the marginals of the best response,

$$b_1[\Lambda](x) = \frac{1}{Z[\Lambda]} \mu(x) \int e^{-c(x,y)/\Lambda} d\nu(y) \quad (6.12)$$

$$b_2[\Lambda](y) = \frac{1}{Z[\Lambda]} \nu(y) \int e^{-c(x,y)/\Lambda} d\mu(x). \quad (6.13)$$

Lemma 6.9 (Convergence of Second Moments). *For any positive sequence $\Lambda_n \rightarrow \bar{\Lambda} > 0$, it holds that $\int \|z\|^2 db[\Lambda_n](z) \rightarrow \int \|z\|^2 db[\bar{\Lambda}](z)$ with $z = (x, y)$.*

Proof. It is sufficient to show convergence of the second moments for the marginals, which is equivalent to $\int \|z\|^2 d\sigma[\Lambda_n](z) \rightarrow \int \|z\|^2 d\sigma[\bar{\Lambda}](z)$ with σ defined in (6.10). By Lebesgue's dominated convergence theorem and the bound in (6.11), we have $Z[\Lambda_n] \rightarrow Z[\bar{\Lambda}]$ and convergence of the second moments follows. $\square$

Proposition 6.10 (Γ -convergence). *For any positive sequence $\Lambda_n \rightarrow \bar{\Lambda} > 0$ as $n \rightarrow \infty$, it holds*

$$E(\cdot, \Lambda_n) \xrightarrow{\Gamma} E(\cdot, \bar{\Lambda}) \text{ as } n \rightarrow \infty \text{ in } \mathcal{W}.$$

Proof. Given $\Lambda_n \rightarrow \bar{\Lambda}$, we have for any sequence $\rho^{(n)} \rightarrow \bar{\rho}$ in $\mathcal{W}$,

$$\liminf_{n \rightarrow \infty} E(\rho^{(n)}, \Lambda_n) \geq \liminf_{n \rightarrow \infty} \iint c(x, y) d\rho^{(n)}(x, y) + \liminf_{n \rightarrow \infty} \Lambda_n (\text{KL}(\rho_1^{(n)} | \mu) + \text{KL}(\rho_2^{(n)} | \nu))$$

(see [30, Chapter 1, (1.2)]), and lower bounds for each term are sufficient. For the first term, since $c \geq 0$, it holds by Fatou's Lemma for weakly converging measures [107, Theorem 2.4],

$$\liminf_{n \rightarrow \infty} \iint c(x, y) d\rho^{(n)}(x, y) \geq \iint c(x, y) d\bar{\rho}(x, y).$$

For the second term, the lower bound follows by lsc of the KL-divergence, and we obtain

$$\liminf_{n \rightarrow \infty} E(\rho^{(n)}, \Lambda_n) \geq E(\bar{\rho}, \bar{\Lambda}).$$

Next, we show for any measure $\bar{\rho}$ there exists a recovery sequence $(\rho^{(n)}) \in \mathcal{P}_2$ with $\mathcal{W}(\rho^{(n)}, \bar{\rho}) \rightarrow 0$ such that

$$E(\bar{\rho}, \bar{\Lambda}) \geq \limsup_{n \rightarrow \infty} E(\rho^{(n)}, \Lambda_n).$$

Indeed, selecting the constant sequence $\rho^{(n)} = \bar{\rho}$, the integral against $c(x, y)$ is already constant, and since $\Lambda_n \rightarrow \bar{\Lambda}$, the lim sup bound follows immediately. $\square$

Lemma 6.11 (Continuity of b). *For any sequence $(\Lambda_n)_n > 0$, $\Lambda_n \rightarrow \bar{\Lambda} > 0$ as $n \rightarrow \infty$, we have $b[\Lambda_n] \rightarrow b[\bar{\Lambda}]$ in $\mathcal{W}$.*

Proof. The boundedness of second moments, which follows from convergence of second moments via Lemma 6.9, provides tightness of $(b[\Lambda_n])_n$, and so $(b[\Lambda_n])_n$ is precompact in the narrow topology [9, Prokhorov's Theorem 5.1.3]. Thanks to convergence of second moments (Lemma 6.9), the sequence $(b[\Lambda_n])_n$ is also precompact in $\mathcal{W}$. Together with Γ -convergence of $E(\cdot, \Lambda_n)$ from Proposition 6.10, we conclude $b[\Lambda_n] \rightarrow b[\bar{\Lambda}]$ in $\mathcal{W}$ [30, Theorem 1.21]. $\square$

Lemma 6.12 (Continuity of V). *$V : (0, \infty) \rightarrow [0, \infty)$ defined in (6.9) is continuous.*

Proof. Substituting the expressions for the marginals of $b[\Lambda]$ obtained in (6.12)-(6.13), we can write

$$V = -2 \log Z + \frac{1}{Z} \int Z_1(x) \log Z_1(x) d\mu(x) + \frac{1}{Z} \int Z_2(y) \log Z_2(y) d\nu(y), \quad (6.14)$$

with $Z_1[\Lambda](x), Z_2[\Lambda](y) \in (0, 1)$ given by

$$Z_1(x) := \int e^{-c(x,y)/\Lambda} d\nu(y), \quad Z_2(y) := \int e^{-c(x,y)/\Lambda} d\mu(x).$$

For any positive sequence $\Lambda_n \rightarrow \bar{\Lambda}$, we have by Lebesgue's dominated convergence theorem that $Z[\Lambda_n]$, $Z_1[\Lambda_n](x)$ and $Z_2[\Lambda_n](y)$ converge (pointwise). By Jensen's inequality, $|Z_1[\Lambda](x) \log Z_1[\Lambda](x)| \leq \frac{1}{\Lambda} \int c(x,y) d\nu(y)$ for all $x \in \text{supp } \mu$, and thanks to Assumption 6.2 this upper bound is integrable w.r.t. $d\mu$. Again by Lebesgue's dominated convergence theorem,

$$\lim_{n \rightarrow \infty} \int Z_1[\Lambda_n](x) \log Z_1[\Lambda_n](x) d\mu(x) = \int Z_1[\bar{\Lambda}](x) \log Z_1[\bar{\Lambda}](x) d\mu(x),$$

and similarly for the term involving Z_2 . We conclude that $V(\Lambda_n) \rightarrow V(\bar{\Lambda})$. $\square$

Theorem 6.13 (Danskin). $\frac{d}{d\Lambda} E_d(\Lambda) = \partial_\Lambda E(\rho, \Lambda)|_{\rho=b[\Lambda]}$. *As a result, the dynamics (6.7) are given by $\dot{\Lambda} = V(\Lambda)$.*

Proof. Let $\bar{\Lambda} > 0$ and define $\Lambda_h := \bar{\Lambda} + h$ with $h \in \mathbb{R}$ small enough so that $\Lambda_h > 0$. Let $\rho^{(h)} := b(\Lambda_h)$. By continuity of $b(\cdot)$ w.r.t. $\mathcal{W}$ (Lemma 6.11), we have that $\rho^{(h)}$ and its marginals converge in $\mathcal{W}$ to $b[\bar{\Lambda}]$ and its marginals, respectively. We can then apply exactly the same argument as in [63, Proposition B.10] to conclude the proof, with the following adjustments: (i) all inequalities and uses of $\liminf$ and $\limsup$ are flipped as $b[\Lambda]$ is a minimizer not a maximizer, with the same result at the end, (ii) our functional $E(\rho, \Lambda)$ is linear in Λ , (iii) its derivative $\partial_\Lambda E(\rho, \Lambda)$ is only lsc in ρ

rather than continuous. Therefore, only the last of the four inequalities in [63, Proposition B.10] needs to be addressed in our setting: for $h < 0$,

$$\limsup_{h \uparrow 0} \frac{E_d(\Lambda_h) - E_d(\bar{\Lambda})}{h} \leq \limsup_{h \uparrow 0} \partial_{\Lambda} E(\rho, \bar{\Lambda})|_{\rho=b(\Lambda_h)} = \limsup_{h \uparrow 0} V(\Lambda_h) = V(\bar{\Lambda}),$$

by continuity of V (Lemma 6.12). $\square$

6.3.3 Long-Time Behavior

In regularized optimal transport problems, it is important for numerical stability that the weight of the regularizer does not grow to infinity too quickly. Under Assumptions 6.2 and 6.4, we prove that Λ grows at most $O(\sqrt{t})$, and that the KL divergence terms in V go to zero as $t \rightarrow \infty$, implying convergence of the best response to the optimal transport plan. By non-negativity of the KL-divergence (thanks to Gibbs' inequality), we have $V \geq 0$, and so $\Lambda(t)$ solving (6.7) is non-decreasing since $\dot{\Lambda} = V(\Lambda)$ by Theorem 6.13.

Theorem 6.14 (Regularization growth condition). *Under the dynamics (6.7), the regularization parameter Λ satisfies*

$$\Lambda(t) \leq \sqrt{2(c^*t + c_0)} \quad \forall t \geq 0,$$

where $c^* < \infty$ is the optimal value of (6.2) and $c_0 > 0$ is a constant dependent on $\Lambda(0) > 0$.

Proof. From $E_d(\Lambda) \leq c^*$ (proof of Lemma 6.8), we have $V(\Lambda) \leq c^*/\Lambda$, and so $\dot{\Lambda} \leq c^*/\Lambda$. By Grönwall's inequality,

$$\Lambda(t) \leq \sqrt{2(c^*t + c_0)}.$$

$\square$

To characterize the long-time behavior of E_d and V , we re-write V and its derivative in compact form using the measure σ defined in (6.10), and show Lipschitzness of V .

Lemma 6.15. *We have $V = -2 \log Z + \mathbb{E}_{\sigma}[\log(Z_1 Z_2)]$, and*

$$\frac{dV}{d\Lambda} = \frac{1}{\Lambda^2} \text{Cov}_{\sigma}(c, \log(Z_1 Z_2)) = \frac{1}{\Lambda^2} (\mathbb{E}_{\sigma}[c \log(Z_1 Z_2)] - \mathbb{E}_{\sigma}[c] \mathbb{E}_{\sigma}[\log(Z_1 Z_2)]).$$

Proof. Using σ defined in (6.10) and expression (6.14) for V , we can compute its derivative in Λ explicitly, using $\frac{dZ}{d\Lambda} = \frac{1}{\Lambda^2} \iint c e^{-c/\Lambda} d\mu d\nu$ and

$$\frac{dZ_1(x)}{d\Lambda} = \frac{1}{\Lambda^2} \int c e^{-c/\Lambda} d\nu, \quad \frac{dZ_2(y)}{d\Lambda} = \frac{1}{\Lambda^2} \int c e^{-c/\Lambda} d\mu.$$

$\square$

Lemma 6.16 (Lipschitzness of V). *There exists a constant $0 < c_0 < \infty$ such that*

$$\left\| \frac{dV}{d\Lambda} \right\|_{\infty} = \sup_{\Lambda \in [0, \infty)} \left| \frac{dV}{d\Lambda}(\Lambda) \right| < c_0.$$

Proof. Consider a sequence $\Lambda_n \rightarrow \infty$. As in the proof of Lemma 6.12, we have $Z[\Lambda_n], Z_1[\Lambda_n](x), Z_2[\Lambda_n](y) \rightarrow 1$ (pointwise) as $n \rightarrow \infty$. Thanks to Lebesgue's dominated convergence theorem again, we can pull the limit $n \rightarrow \infty$ into all integrals appearing in the expression of $dV/d\Lambda$ from Lemma 6.15. To do so, we use that $|\log Z_1[\Lambda](x)| \leq \int c dv/\Lambda$ (and similarly for Z_2) thanks to Jensen's inequality, and $ce^{-c/\Lambda}/\Lambda \leq 1/e$. Hence

$$\left| \frac{dV}{d\Lambda}(\Lambda_n) \right| \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Using a similar argument via Lebesgue's dominated convergence theorem, one can show that for any sequence $\tilde{\Lambda}_n \rightarrow \tilde{\Lambda}$ it holds $dV(\tilde{\Lambda}_n)/d\Lambda \rightarrow dV(\tilde{\Lambda})/d\Lambda$, and so $dV/d\Lambda$ is continuous. Finally, for our sequence $\Lambda_n \rightarrow \infty$, there exists $N > 0$ such that $|dV(\Lambda_n)/d\Lambda| < 1$ for all $n \geq N$, and by continuity $dV/d\Lambda$ achieves its maximum (denoted by c_1) on $[0, \Lambda_N]$. We conclude by taking $c_0 := \max\{1, c_1\}$. $\square$

Proposition 6.17 (Properties $\Lambda(t)$). *The ODE $\dot{\Lambda} = V(\Lambda)$ with $\Lambda(0) > 0$ has a unique solution $\Lambda(\cdot) \in C([0, \infty))$ that satisfies $V(\Lambda(t)) \rightarrow 0$ as $t \rightarrow \infty$. In other words, the best response $b[\Lambda(t)]$ satisfies the marginal constraints asymptotically.*

Proof. For $\Lambda(0) > 0$, define $C_0 := \frac{c^*}{\Lambda_0(\Lambda_0+1)}$. Since $\dot{\Lambda} = V(\Lambda) \geq 0$, we have

$$\Lambda(t) \geq \Lambda(0) = -\frac{1}{2} + \sqrt{\frac{1}{4} + \frac{c^*}{C_0}} \Rightarrow V(\Lambda) \leq \frac{c^*}{\Lambda} \leq C_0(1 + \Lambda).$$

for all $t \geq 0$. Thanks to continuity of V (Lemma 6.12), Lipschitzness of V (Lemma 6.16) and the growth condition from Theorem 6.14, we obtain global existence and uniqueness of a continuous solution $\Lambda(t)$ to the ODE $\dot{\Lambda} = V(\Lambda)$. Further, $\Lambda(t)$ is non-decreasing, so it must either remain bounded for all times, or diverge to $+\infty$. Thus, we can write $\Lambda(t) \rightarrow \Lambda_{\infty} \in [\Lambda(0), \infty]$, and by continuity,

$$V(\Lambda(t)) \rightarrow V_{\infty} := V(\Lambda_{\infty}) \in [0, \infty].$$

Note from Lemma 6.8 for all $t \geq 0$,

$$V(\Lambda(t)) = \frac{1}{\Lambda(t)} \left(E_d(\Lambda(t)) - \int c db[\Lambda(t)] \right) \leq \frac{c^*}{\Lambda(0)},$$

and so V is uniformly bounded from above. Thus, there are two possibilities: either (a) $V_\infty \in (0, \infty)$, or (b) $V_\infty = 0$. In Case (a), there exists $T > 0$ such that for all $t \geq T$,

$$\dot{\Lambda}(t) = V(\Lambda(t)) > \frac{V_\infty}{2} > 0.$$

Integrating the ODE, we obtain the lower bound

$$\Lambda(t) > \frac{\Lambda_\infty}{2}t + V_0 \quad \forall t \geq T,$$

where $V_0 = V_0(T) \in \mathbb{R}$. For large enough $t \geq T$, this is a contradiction to the upper bound in Theorem 6.14, so (a) cannot happen. We conclude $V(\Lambda(t)) \rightarrow 0$. $\square$

Corollary 6.18 (Convergence of OT Cost and Minimizer). *The best response $b[\Lambda(t)] \rightarrow \rho^*$ in $\mathcal{W}$, where ρ^* is an optimizer for (6.2), and $E_d(\Lambda) \rightarrow c^*$ as $t \rightarrow \infty$.*

Proof. Applying Theorem 6.13, the time derivative of E_d along solutions to (6.7) is

$$\frac{d}{dt}E_d(\Lambda) = V(\Lambda)^2.$$

This indicates that unless $V = 0$, the energy will increase. From Proposition 6.17, we have that the solution $\Lambda(t)$ to (6.7) converges to $\Lambda_\infty \in [\Lambda(0), \infty]$ as $t \rightarrow \infty$. By continuity of the best response (Lemma 6.11), $b[\Lambda(t)] \rightarrow b_\infty := b(\Lambda_\infty)$ in $\mathcal{W}$. Note that $V_\infty := V(\Lambda_\infty) = 0$ (Proposition 6.17) means $b_\infty \in \Gamma(\mu, \nu)$ and so $E_d(\Lambda_\infty) = \int c b_\infty \geq c^*$ since b_∞ is admissible for the OT problem (6.2). Since also $E_d(\Lambda_\infty) \leq c^*$ (Lemma 6.8), we have equality, and b_∞ is an optimal plan in $\Gamma(\mu, \nu)$. Finally, E_d is continuous thanks to the continuity of b and V , and so $E_d(\Lambda) \rightarrow E_d(\Lambda_\infty) = c^*$. $\square$

6.4 Numeric Implementation

In practice, Theorem 6.13 allows us to chose $\dot{\Lambda} = V(\Lambda)$, rather than taking a gradient of $b[\Lambda]$ with respect to Λ . In our numerical implementation¹, we chose very small β , which results in ρ updating much faster than Λ , and we expect this to approximate the dynamics (6.7) due to the timescale separation results in [28]. We adapt a method from [153] by approximating the distribution ρ by particles $(X_i, Y_i) \sim \rho$. We divide the points $(X_j, Y_j)_{j=1}^{2N}$ into two sets: for one set, denoted $(X_i^{(1)}, Y_i^{(1)})_{i=1}^N$, we initialize $X_i^{(1)} \sim \mu$ and $Y_i^{(1)} = X_i^{(1)} + \xi_i$, where $\xi_i \sim \mathcal{N}(0, \eta)$ with small η , and we evolve $Y_i^{(1)}$ while keeping $X_i^{(1)}$ fixed. For the other set, denoted $(X_i^{(2)}, Y_i^{(2)})$ we initialize $Y_i^{(2)} \sim \nu$

¹Code available <https://github.com/LEConger/OT-MinMax-GradientFlow>.

and set $X_i^{(2)} = Y_i^{(2)} + \xi_i$, and evolve $X_i^{(2)}$ while keeping $Y_i^{(2)}$ fixed, resulting in

$$\begin{aligned}\dot{X}_i^{(2)}(t) &= -\frac{1}{\beta} \nabla_x c(X_i^{(2)}(t), Y_i^{(2)}(t)) + \frac{1}{\beta} \Lambda(t) \nabla_x \log \frac{\rho_1(t, X_i^{(2)}(t))}{\mu(X_i^{(2)}(t))} \\ \dot{Y}_i^{(1)}(t) &= -\frac{1}{\beta} \nabla_y c(X_i^{(1)}(t), Y_i^{(1)}(t)) + \frac{1}{\beta} \Lambda(t) \nabla_y \log \frac{\rho_2(t, Y_i^{(1)}(t))}{\nu(Y_i^{(1)}(t))} \\ \dot{\Lambda}(t) &= \text{KL}(\rho_1(t) \mid \mu) + \text{KL}(\rho_2(t) \mid \nu).\end{aligned}$$

The densities ρ_1, ρ_2 are approximated at each time step by the binning estimator

$$\begin{aligned}\rho_1(x) &\approx \frac{1}{2N} \sum_{k=1}^{N_b} \mathbf{1}_{\Omega_k}\{x\} \sum_{m=1}^N (\mathbf{1}_{\Omega_k}\{X_m^{(1)}\} + \mathbf{1}_{\Omega_k}\{X_m^{(2)}\}) \\ \rho_2(y) &\approx \frac{1}{2N} \sum_{k=1}^{N_b} \mathbf{1}_{\Omega_k}\{y\} \sum_{m=1}^N (\mathbf{1}_{\Omega_k}\{Y_m^{(1)}\} + \mathbf{1}_{\Omega_k}\{Y_m^{(2)}\}),\end{aligned}$$

where the domain is divided into N_b regions, and the fraction of particles in each region Ω_k is counted. The key differences from [153] are (i) how we estimate the marginal densities $\rho_1(t), \rho_2(t)$ at each time step, and (ii) the choice of particle initialization and updates. For (i), our density estimator is significantly faster to compute because the complexity scales as $O(N_b)$ instead of $O(N^2)$, and we can select $N_b \ll N^2$. For (ii), the update strategy ensures that each marginal has at least half of the particles distributed according to the input distribution, and the correlation between X and Y particles in the initial condition aims to minimize cost functions of the form $c(x, y) = h(x - y)$. We also compute marginal estimates of μ and ν from samples rather than assuming closed-form knowledge, and including a time-varying Λ .

Method Comparison We compare our method with the method proposed in [153], for computing an approximate transport plan between $\mathcal{N}(0.4 \mathbf{I}_2, \Sigma)$ and $\mathcal{N}(0.6 \mathbf{I}_2, \Sigma)$ for $\Sigma = 0.02 \mathbf{I}_2$. To replicate the method in [153], we used 1000 points with $dt = 0.00002$ for $T = 50$ iterations and selected the τ parameter for best results. This results in a run time of 2 minutes and 14 seconds. In comparison, we implemented our method with $2N = 20,000$ points, $dt = 0.0005$, and $T = 2000$ iterations, which had a run time of only six seconds. In both methods we include additive Gaussian noise in the particle update with standard deviation $0.02\sqrt{dt}$ for smoothing; the results are shown in Figure 6.1. In Figures 6.1 and 6.4, we plot instances of the interpolant $\mu_s := ((1-s)x + sy)_{\#} \rho_\infty$ for different values of $s \in (0, 1)$, where ρ_∞ is the coupling at the final time obtained dynamically via the chosen numerical scheme.

The optimal transport cost c^* for $c(x, y) = \|x - y\|^2$ can be computed in closed form to be 0.080; numerically, our method achieves a cost of 0.0861 (Figure 6.2), with an L^2 error between the




Figure 6.1: The transport map generated with a kernel density estimator method from [153] (top) is limited by the number of points ($N = 1,000$), whereas estimating the density by counting voxels (bottom, our method) allows for a greater number of points (20,000) and faster computation time.




Figure 6.2: In the setting of two log-concave input measures, the coupling cost starts small due to the initialization of $X_i \approx Y_i$, and increases to approach the optimal coupling cost. The regularization penalty weight Λ increases more slowly over time, as expected due to the $O(\sqrt{t})$ estimate, and the KL divergence between the marginals and the inputs decreases as the particles drift to minimize the entropy regularization terms.

marginals and input marginals of 0.006 and KL divergence of 0.098. This shows the obtained transport plan is very close to the optimal plan. The computational speedup in comparison with [153] is due to the estimation of the kernel density.

Reversing KL Divergence, Non-Log-Concave Source and Target We experiment with non-log-concave reference measures, shown in Figure 6.3. We also test using the energy $\tilde{E}$, where the KL terms are reversed. In this gradient flow setting, $\text{KL}(\mu | \cdot)$ is not necessarily geodesically convex, even if the source marginal μ is log-concave. Interestingly, we observe better convergence using the reversed KL for a number of examples which coincides with the phenomenon that $\text{KL}(\cdot | \mu)$ exhibits "mode-seeking" behavior while $\text{KL}(\mu | \cdot)$ exhibits moment matching behavior [195]. In the particular example shown (Figure 6.3) where the density of μ has a ring-shape with a central




Figure 6.3: The source and target distributions for the experiment (b) are shown. Note that the input marginals are not log-concave, going beyond Assumption 6.3.

peak and ν is a mixture of four Gaussian distributions, using $\text{KL}(\mu \mid \rho_1)$ to flow points from ν to μ results in a smaller L^2 error norm, $\int |\rho_1(x) - \rho_1^*(x)|^2 dx$, compared with $\text{KL}(\rho_1 \mid \mu)$. We suspect this is due to the moment-matching behavior of $\text{KL}(\mu \mid \rho_1)$, since the estimated marginal achieved with $\text{KL}(\mu \mid \rho_1)$ preserves mass on the outer ring, which is further away from the initial conditions of the point $(X_i^{(2)})$ (as compared to being mode-seeking, which would focus the points more on the mode near the origin). For the points $(Y_i^{(1)})$ flowing from μ to ν , the mode-seeking behavior (corresponding to $\text{KL}(\rho_2 \mid \nu)$) is more advantageous because ν is very multi-modal. In contrast, when using $\text{KL}(\nu \mid \rho_2)$, we observe more outlier points which makes the resulting marginal further from the target distribution.

In our experiments, we found that using randomly selected left or right differences to estimate the gradients outperformed the centered first-order and centered second-order gradients. Using the left and right differences, we compare three settings:

- (I) Both the $X^{(1)}$ and $Y^{(2)}$ particles are updated via E .
- (II) Both the $X^{(1)}$ and $Y^{(2)}$ particles are updated via $\tilde{E}$.
- (III) The $X^{(1)}$ particles are updated via $\tilde{E}$ and the $Y^{(2)}$ particles are updated via E .

The simulation results, shown in Figure 6.4, show that evolution according to (I) and (III) give similar results, while (II) is slightly worse with some rogue particles in corners at the target ν . For all three setting, a large error comes from the source μ . The L^2 norm error and KL divergences between the obtained marginals and the true input marginals are listed in Table 6.1, showing that method (III) has the smallest L^2 error and ties for the smallest KL divergence error. Investigating when $\tilde{E}$ versus E gives better performance is a future direction of research, along with the displacement convexity properties of $\tilde{E}$.




Figure 6.4: We test our method on two input measures which are non-log-concave (Left side is μ and right side is ν). When running our method with two different entropy regularizers, using $\text{KL}(\rho | \mu, \nu)$ for the $Y^{(2)}$ particles and $\text{KL}(\mu, \nu | \rho)$ for the $X^{(2)}$ particles (labeled "both KL") is compared with using the same for all particles. While the displacement convexity of $\text{KL}(\cdot | \mu)$ is well known, the displacement convexity of $\text{KL}(\mu | \cdot)$ is not.

	L^2 error	KL	reverse KL	total KL
(I), E	0.030	0.14	0.13	0.27
(II), $\tilde{E}$	0.033	0.88	0.23	1.11
(III), both	0.028	0.14	0.13	0.27

Table 6.1: Using method (III) results in the smallest L^2 error among the three methods because it takes advantage of $\tilde{E}$ being a better flow in the direction from ν ($s = 1$) to μ ($s = 0$) and E being a better flow in the direction from μ ($s = 0$) to ν ($s = 1$). The large error for the $\tilde{E}$ flow is due to the poor performance for $\tilde{E}$ in the μ to ν direction, and this is avoided in method (III).

Acknowledgments

LC is supported by NDSEG and PIMCO fellowships. FH is supported by start-up funds at the California Institute of Technology and by NSF CAREER Award 2340762. RB is supported in part by a Vannevar Bush Faculty Fellowship (award N00014-22-1-2790) held by Andrew M. Stuart. EM is supported in part from NSF award 2240110.

PART I BIBLIOGRAPHY

- [1] A. Agarwal and T. Zhang. Minimax regret optimization for robust machine learning under distribution shift. In *Proceedings of Thirty Fifth Conference on Learning Theory*. Conference on Learning Theory, pages 2704–2729. PMLR, June 28, 2022.
- [2] A. Aggarwal, M. Mittal, and G. Battineni. Generative adversarial network: an overview of theory and applications. *International Journal of Information Management Data Insights*, 1(1):100004, 2021. ISSN: 2667-0968. DOI: <https://doi.org/10.1016/j.jjime.2020.100004>.
- [3] L. Alaoglu. Weak topologies of normed linear spaces. *The Annals of Mathematics*, 41(1):252, Jan. 1940. ISSN: 0003486X. DOI: 10.2307/1968829.
- [4] L. Alasio, M. Bruna, S. Fagioli, and S. Schulz. Existence and regularity for a system of porous medium equations with small cross-diffusion and nonlocal drifts. *Nonlinear Analysis*, 223:113064, Oct. 2022. ISSN: 0362-546X. DOI: 10.1016/j.na.2022.113064.
- [5] C. D. Aliprantis and K. C. Border. Correspondences. en. In *Infinite Dimensional Analysis: A Hitchhiker's Guide*, pages 555–590. Springer, Berlin, Heidelberg, 1999. ISBN: 9783540295877. DOI: 10.1007/3-540-29587-9_17.
- [6] C. D. Aliprantis and K. C. Border. Measurable correspondences. en. In *Infinite Dimensional Analysis: A Hitchhiker's Guide*, pages 591–620. Springer, Berlin, Heidelberg, 2006. ISBN: 9783540295877. DOI: 10.1007/3-540-29587-9_18.
- [7] A. Alsenafi and A. B. T. Barbaro. A multispecies cross-diffusion model for territorial development. *Mathematics*, 9(12):1428, Jan. 2021. ISSN: 2227-7390. DOI: 10.3390/math9121428.
- [8] J. Altschuler, J. Niles-Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. *NeurIPS*, 30, 2017.
- [9] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, second edition, 2008, pages x+334. ISBN: 978-3-7643-8721-1.
- [10] J. An and J. Lu. Convergence of two-timescale gradient descent ascent dynamics: finite-dimensional and mean-field perspectives, 2025. arXiv: 2501.17122.
- [11] N. An, K. Yang, Y. Chen, and L. Yang. Wasserstein distributionally robust optimization for train operation and freight assignment in a metro-based underground logistics system. *Computers & Industrial Engineering*, 192:110228, 2024. ISSN: 0360-8352. DOI: <https://doi.org/10.1016/j.cie.2024.110228>.
- [12] C. Appert-Rolland, P. Degond, and S. Motsch. Two-way multi-lane traffic model for pedestrians in corridors. en. *Networks and Heterogeneous Media*, 6(3):351–381, Aug. 2011. ISSN: 1556-1801. DOI: 10.3934/nhm.2011.6.351.

- [13] L. Aquilanti, S. Cacace, F. Camilli, and R. De Maio. A Mean Field Games Approach to Cluster Analysis. en. *Applied Mathematics & Optimization*, 84(1):299–323, Aug. 2021. ISSN: 1432-0606. DOI: 10.1007/s00245-019-09646-2.
- [14] C. Baiocchi and A. Capelo. *Variational and Quasivariational Inequalities*. John Wiley and Sons, Belfast, 1984.
- [15] D. Balagué, J. Carrillo, and Y. Yao. Confinement for repulsive-attractive kernels. *Discrete and Continuous Dynamical Systems-series B*, 2012.
- [16] D. Balagué, J. A. Carrillo, T. Laurent, and G. Raoul. Nonlocal interactions by repulsive-attractive potentials: Radial ins/stability. *Physica D: Nonlinear Phenomena*. Emergent Behaviour in Multi-particle Systems with Non-local Interactions, 260:5–25, Oct. 2013. ISSN: 0167-2789. DOI: 10.1016/j.physd.2012.10.002.
- [17] D. Bartl, S. Drapeau, J. Obloj, and J. Wiesel. Sensitivity analysis of Wasserstein distributionally robust optimization problems. en. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 477(2256):20210176, Dec. 2021. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2021.0176.
- [18] T. Başar and P. Bernhard. Appendix B: Danskin’s Theorem. en. In T. Başar and P. Bernhard, editors, *H_∞-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, Modern Birkhäuser Classics, pages 383–389. Birkhäuser, Boston, MA, 2008. ISBN: 9780817647575. DOI: 10.1007/978-0-8176-4757-5_10.
- [19] L. Beck, D. Matthes, and M. Zizza. Exponential convergence to equilibrium for coupled systems of nonlinear degenerate drift diffusion equations. *SIAM Journal on Mathematical Analysis*, 55(3):1766–1809, 2023. DOI: 10.1137/21M1466980. eprint: <https://doi.org/10.1137/21M1466980>.
- [20] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, Jan. 1, 2000. ISSN: 0945-3245. DOI: 10.1007/s002110050002.
- [21] J.-D. Benamou, G. Carlier, and M. Laborde. An augmented Lagrangian approach to Wasserstein gradient flows and applications. en. *ESAIM: Proceedings and Surveys*, 54:1–17, June 2016. ISSN: 2267-3059. DOI: 10.1051/proc/201654001.
- [22] A. L. Bertozzi, J. A. Carrillo, and T. Laurent. Blow-up in multidimensional aggregation equations with mildly singular interaction kernels*. *Nonlinearity*, 22(3):683, Feb. 2009. ISSN: 0951-7715. DOI: 10.1088/0951-7715/22/3/009.
- [23] A. L. Bertozzi, T. Laurent, and F. Lager. Aggregation and spreading via the newtonian potential: the dynamics of patch solutions. *Mathematical Models and Methods in Applied Sciences*, 22:1140005, supp01, Apr. 2012. ISSN: 0218-2025. DOI: 10.1142/S0218202511400057.
- [24] D. P. Bertsekas. *Control of Uncertain Systems with a Set-Membership Description of Uncertainty*. PhD thesis, MIT, Cambridge, MA, USA, 1971.
- [25] S. Bickel, M. Brückner, and T. Scheffer. Discriminative learning under covariate shift. *The Journal of Machine Learning Research*, 10:2137–2155, Dec. 1, 2009. ISSN: 1532-4435.

- [26] V. Bogachev. *Measure Theory*. en. Springer Berlin, Heidelberg, 1st edition, 2007.
- [27] V. S. Borkar. *Stochastic approximation: a dynamical systems viewpoint*, volume 48. Springer, 2009.
- [28] V. S. Borkar. Multiple Timescales. In *Stochastic Approximation*. Volume 48, pages 64–77. Hindustan Book Agency, Gurgaon, 2008. ISBN: 9788185931852. DOI: 10.1007/978-93-86279-38-5_6.
- [29] A. Braides. A handbook of gamma convergence. In *Handbook of Differential Equations. Stationary Partial Differential Equations*. Volume 3. Elsevier, 2006.
- [30] A. Braides. *Gamma-Convergence for Beginners*. en. Oxford University Press, July 2002. ISBN: 9780191709890. DOI: 10.1093/acprof:oso/9780198507840.001.0001.
- [31] H. Brezis. The Hille-Yosida Theorem. en. In H. Brezis, editor, *Functional Analysis, Sobolev Spaces and Partial Differential Equations*, pages 181–199. Springer, New York, NY, 2011. ISBN: 9780387709147. DOI: 10.1007/978-0-387-70914-7_7.
- [32] H. Brézis. *Opérateurs maximaux monotones et semi-groupes de contractions dans les espaces de Hilbert*. North-Holland Publishing Co., Amsterdam, 1973, pages vi+183. North-Holland Mathematics Studies, No. 5. Notas de Matemática (50).
- [33] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. en. S. S. Chern, B. Eckmann, P. De La Harpe, H. Hironaka, F. Hirzebruch, N. Hitchin, L. Hörmander, M.-A. Knus, A. Kupiainen, J. Lannes, G. Lebeau, M. Ratner, D. Serre, Y. G. Sinai, N. J. A. Sloane, J. Tits, M. Waldschmidt, S. Watanabe, M. Berger, J. Coates, and S. R. S. Varadhan, editors, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999. ISBN: 9783642083990 9783662124949. DOI: 10.1007/978-3-662-12494-9.
- [34] F. E. Browder. Nonlinear maximal monotone operators in Banach space. In *Mathematische Annalen*. Volume 175, pages 89–113. 1968.
- [35] F. E. Browder. Nonlinear monotone operators and convex sets in Banach spaces. *Bulletin of the American Mathematical Society*, 71(5):780–785, Sept. 1965. ISSN: 0002-9904, 1936-881X.
- [36] R. E. Bruck. On the weak convergence of an ergodic iteration for the solution of variational inequalities for monotone operators in Hilbert space. *Journal of Mathematical Analysis and Applications*, 61(1):159–164, Nov. 1977. ISSN: 0022-247X. DOI: 10.1016/0022-247X(77)90152-4.
- [37] A. Brunk, H. Egger, T. Oyediji, Y. Yang, and B.-X. Xu. On existence, uniqueness and stability of solutions to cahn–hilliard/allen–cahn systems with cross-kinetic coupling. *Nonlinear Analysis: Real World Applications*, 77:104051, 2024. ISSN: 1468-1218. DOI: <https://doi.org/10.1016/j.nonrwa.2023.104051>.
- [38] F. Bullo. *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.2 edition, 2024. ISBN: 9798836646806.

- [39] M. Burger, M. Di Francesco, S. Fagioli, and A. Stevens. Sorting phenomena in a mathematical model for two mutually attracting/repelling species. *SIAM Journal on Mathematical Analysis*, 50(3):3210–3250, Jan. 1, 2018. ISSN: 0036-1410. DOI: 10.1137/17M1125716.
- [40] M. Burger, L. He, and C.-B. Schönlieb. Cahn–hilliard inpainting and a generalization for grayvalue images. *SIAM Journal on Imaging Sciences*, 2(4):1129–1167, 2009. DOI: 10.1137/080728548.
- [41] M. Butano, C. Appert-Rolland, and D. Ullmo. Discounted Mean-Field Game model of a dense static crowd with variable information crossed by an intruder. en. *SciPost Physics*, 16(4):104, Apr. 2024. ISSN: 2542-4653. DOI: 10.21468/SciPostPhys.16.4.104.
- [42] A. Camacho and E. Conover. Manipulation of social program eligibility. *American Economic Journal: Economic Policy*, 3(2):41–65, 2011. ISSN: 1945-7731.
- [43] C. Cancès, D. Matthes, I. Medina, and B. Schmitzer. Continuum of coupled Wasserstein gradient flows, Nov. 2024. DOI: 10.48550/arXiv.2411.13969. arXiv:2411.13969.
- [44] J. Carrillo, A. Fernandez-Jimenez, and D. G’omez-Castro. Partial mass concentration for fast-diffusions with non-local aggregation terms. In Apr. 10, 2023.
- [45] J. A. Carrillo, M. DiFrancesco, A. Figalli, T. Laurent, and D. Slepcev. Global-in-time weak measure solutions and finite-time aggregation for nonlocal interaction equations. *Duke Mathematical Journal*, 156(2), Feb. 1, 2011. ISSN: 0012-7094. DOI: 10.1215/00127094-2010-211.
- [46] J. A. Carrillo, Y. Chen, D. Z. Huang, J. Huang, and D. Wei. Fisher-Rao Gradient Flow: Geodesic Convexity and Functional Inequalities, July 2024. DOI: 10.48550/arXiv.2407.15693. arXiv:2407.15693.
- [47] J. A. Carrillo, A. Chertock, and Y. Huang. A finite-volume method for nonlinear nonlocal equations with a gradient flow structure. *Communications in Computational Physics*, 17(1):233–258, Jan. 2015. ISSN: 1815-2406, 1991-7120. DOI: 10.4208/cicp.160214.010814a.
- [48] J. A. Carrillo, K. Craig, L. Wang, and C. Wei. Primal Dual Methods for Wasserstein Gradient Flows. en. *Foundations of Computational Mathematics*, 22(2):389–443, Apr. 2022. ISSN: 1615-3383. DOI: 10.1007/s10208-021-09503-1.
- [49] J. A. Carrillo, Y. Huang, and M. Schmidtchen. Zoology of a nonlocal cross-diffusion model for two species. *SIAM Journal on Applied Mathematics*, 78(2):1078–1104, Jan. 2018. ISSN: 0036-1399, 1095-712X. DOI: 10.1137/17M1128782.
- [50] J. A. Carrillo, R. J. McCann, and C. Villani. Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates. *Revista Matemática Iberoamericana*, 19(3):971–1018, Dec. 2003. ISSN: 0213-2230.
- [51] J. A. Carrillo, R. J. McCann, and C. Villani. Contractions in the 2-wasserstein length space and thermalization of granular media. *Archive for Rational Mechanics and Analysis*, 179(2):217–263, Feb. 1, 2006. ISSN: 1432-0673. DOI: 10.1007/s00205-005-0386-1.

- [52] J. A. Carrillo and G. Toscani. Contractive Probability Metrics and Asymptotic Behavior of Dissipative Kinetic Equations. *Rivista di Matematica della Universita di Parma*, 7(6):75–198, 2007.
- [53] J. A. Carrillo, M. Chipot, and Y. Huang. On global minimizers of repulsive-attractive power-law interaction energies. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028):20130399, Nov. 13, 2014. ISSN: 1364-503X, 1471-2962. DOI: 10.1098/rsta.2013.0399.
- [54] J. A. Carrillo, M. Di Francesco, A. Esposito, S. Fagioli, and M. Schmidtchen. Measure solutions to a system of continuity equations driven by Newtonian nonlocal interactions. en. *Discrete and Continuous Dynamical Systems*, 40(2):1191–1231, Feb. 2020. ISSN: 1078-0947. DOI: 10.3934/dcds.2020075.
- [55] G. Cavagnari, G. Savaré, and G. E. Sodini. A Lagrangian approach to totally dissipative evolutions in Wasserstein spaces, May 2023. DOI: 10.48550/arXiv.2305.05211. arXiv:2305.05211.
- [56] G. Cavagnari, G. Savaré, and G. E. Sodini. Dissipative probability vector fields and generation of evolution semigroups in Wasserstein spaces. en. *Probability Theory and Related Fields*, 185(3):1087–1182, Apr. 2023. ISSN: 1432-2064. DOI: 10.1007/s00440-022-01148-7.
- [57] G. H.-G. Chen and R. T. Rockafellar. Convergence Rates in Forward–Backward Splitting. en. *SIAM Journal on Optimization*, 7(2):421–444, May 1997. ISSN: 1052-6234, 1095-7189. DOI: 10.1137/S1052623495290179.
- [58] Y. Chen and T. Kolokolnikov. A minimal model of predator-swarm interactions. en. *Journal of The Royal Society Interface*, 11(94):20131208, May 2014. ISSN: 1742-5689, 1742-5662. DOI: 10.1098/rsif.2013.1208.
- [59] L. Chizat, A. Delalande, and T. Vaškevičius. Sharper exponential convergence rates for Sinkhorn’s algorithm in continuous settings. en. *Mathematical Programming*, 215(1):809–858, Jan. 2026. ISSN: 1436-4646. DOI: 10.1007/s10107-025-02242-z.
- [60] M. Cirant. Multi-population Mean Field Games systems with Neumann boundary conditions. *Journal de Mathématiques Pures et Appliquées*, 103(5):1294–1315, May 2015. ISSN: 0021-7824. DOI: 10.1016/j.matpur.2014.10.013.
- [61] L. Conger, F. Hoffmann, R. Baptista, and E. Mazumdar. Computing Optimal Transport Plans via Min-Max Gradient Flows. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 2210–2217, Dec. 2025. DOI: 10.1109/CDC57313.2025.11312926. ISSN: 2576-2370.
- [62] L. Conger, F. Hoffmann, E. Mazumdar, and L. Ratliff. Strategic Distribution Shift of Interacting Agents via Coupled Gradient Flows. en. *Advances in Neural Information Processing Systems*, 36:45971–46006, Dec. 2023.
- [63] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Coupled Wasserstein Gradient Flows for Min-Max and Cooperative Games, Nov. 2024. DOI: 10.48550/arXiv.2411.07403. arXiv:2411.07403 [math].

- [64] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff. Monotone Multispecies Flows, June 2025. doi: 10.48550/arXiv.2506.22947. arXiv:2506.22947.
- [65] L. Conger, F. Hoffmann, and G. Savaré. Dissipative variational evolutions in metric spaces, 2026.
- [66] R. Cottle, F. Giannessi, and J. Lions. *Variational Inequalities and Complementarity Problems: Theory and Applications*. en. John Wiley & Sons, Chichester-New York-Brisbane-Toronto, Jan. 1981. ISBN: 0-471-27610-3.
- [67] K. Craig and I. Topaloglu. Convergence of Regularized Nonlocal Interaction Energies. en. *SIAM Journal on Mathematical Analysis*, 48(1):34–60, Jan. 2016. ISSN: 0036-1410, 1095-7154. doi: 10.1137/15M1013882.
- [68] K. Craig, N. G. Trillos, and Đ. Nikolić. Vector valued optimal transport: from dynamic to static formulations, May 2025. doi: 10.48550/arXiv.2505.03670. arXiv:2505.03670.
- [69] M. G. Crandall. Nonlinear semigroups and evolution governed by accretive operators. In F. E. Browder, editor, *Nonlinear Functional Analysis and its Applications*. Volume 45, Proceedings of Symposia in Pure Mathematics, pages 305–338. American Mathematical Society, Providence, 1986.
- [70] M. G. Crandall and L. C. Evans. On the relation of the operator $\partial/\partial s + \partial/\partial \tau$ to evolution governed by accretive operators. *Israel J. Math.*, 21:261–278, 1975.
- [71] M. G. Crandall and T. M. Liggett. Generation of semi-groups of nonlinear transformations on general Banach spaces. *Amer. J. Math.*, 93:265–298, 1971.
- [72] W. Cukierski. Give Me Some Credit. en, 2011.
- [73] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *NeurIPS*, volume 26. Curran Associates, Inc., 2013.
- [74] J. M. Danskin. *The Theory of Max-Min and its Application to Weapons Allocation Problems*, volume 5 of *Econometrics and Operations Research*. Springer, Berlin, Heidelberg, 1967. ISBN: 9783642460944. doi: 10.1007/978-3-642-46092-0.
- [75] A. Davydov, S. Jafarpour, and F. Bullo. Non-euclidean contraction theory for robust nonlinear stability. *IEEE Transactions on Automatic Control*, 67(12):6667–6681, 2022. doi: 10.1109/TAC.2022.3183966.
- [76] E. De Giorgi, M. Degiovanni, A. Marino, and M. Tosques. Evolution equations for a class of nonlinear operators. *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Natur. (8)*, 75(1-2):1–8 (1984), 1983. ISSN: 0392-7881.
- [77] C. De Lellis. Ordinary differential equations with rough coefficients and the renormalization theorem of ambrosio. In *Séminaire Bourbaki*, volume 59 of 972, pages 175–204, Mar. 2007.
- [78] L. De Pascale and L. Nenna. A variational formulation of a Multi-Population Mean Field Games with non-local interactions, Sept. 2024. doi: 10.48550/arXiv.2408.03118. arXiv:2408.03118 [cs, math].

- [79] N. Deb, Y.-H. Kim, S. Pal, and G. Schiebinger. Wasserstein Mirror Gradient Flow as the limit of the Sinkhorn Algorithm, July 2023. DOI: 10.48550/arXiv.2307.16421. arXiv:2307.16421.
- [80] T. Debiec, M. Mandal, and M. Schmidtchen. From finite to continuous phenotypes in (visco-)elastic tissue growth models. *Journal of Differential Equations*, 437:113375, Aug. 2025. DOI: 10.1016/j.jde.2025.113375.
- [81] T. Debiec and M. Schmidtchen. Incompressible limit for a two-species tumour model with coupling through brinkman’s law in one dimension. *Acta Applicandae Mathematicae*, 169(1):593–611, Oct. 1, 2020. ISSN: 1572-9036. DOI: 10.1007/s10440-020-00313-1.
- [82] G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- [83] E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, June 2010. ISSN: 0030-364X, 1526-5463. DOI: 10.1287/opre.1090.0741.
- [84] C. Dellarocas. Strategic manipulation of internet opinion forums: implications for consumers and firms. *Management Science*, 52(10):1577–1593, Oct. 2006. ISSN: 0025-1909, 1526-5501. DOI: 10.1287/mnsc.1060.0567.
- [85] L. Desvillettes and C. Villani. On the spatially homogeneous landau equation for hard potentials part i : existence, uniqueness and smoothness. *Communications in Partial Differential Equations*, 25(1):179–259, Jan. 1, 2000. ISSN: 0360-5302. DOI: 10.1080/03605300008821512.
- [86] M. Di Francesco, A. Esposito, and S. Fagioli. Nonlinear degenerate cross-diffusion systems with nonlocal interaction. *Nonlinear Analysis*, 169:94–117, Apr. 2018. ISSN: 0362-546X. DOI: 10.1016/j.na.2017.12.003.
- [87] M. Di Francesco and S. Fagioli. A nonlocal swarm model for predator-prey interactions. en. *Mathematical Models and Methods in Applied Sciences*, 26(02):319–355, Feb. 2016. ISSN: 0218-2025, 1793-6314. DOI: 10.1142/S0218202516400042.
- [88] M. Di Francesco, S. Fagioli, and V. Iorio. Second order two-species systems with nonlocal interactions: existence and large damping limits. *Acta Appl. Math.*, 184:Paper No. 9, 41, 2023. ISSN: 0167-8019, 1572-9036. DOI: 10.1007/s10440-023-00564-8.
- [89] C. Domingo-Enrich, S. Jelassi, A. Mensch, G. Rotskoff, and J. Bruna. A mean-field analysis of two-player zero-sum games, May 6, 2021. DOI: 10.48550/arXiv.2002.06277. arXiv:2002.06277[cs,math,stat].
- [90] J. Dong, A. Roth, Z. Schutzman, B. Waggoner, and Z. S. Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018.
- [91] M. Doumic, S. Hecht, B. Perthame, and D. Peurichard. Multispecies cross-diffusions: From a nonlocal mean-field to a porous medium system without self-diffusion. *Journal of Differential Equations*, 389:228–256, Apr. 2024. DOI: 10.1016/j.jde.2024.01.017.

- [92] Q. Du, K. Wang, E. Zhang, and C. Zhong. A Particle Algorithm for Mean-Field Variational Inference, Dec. 2024. DOI: 10.48550/arXiv.2412.20385. arXiv:2412.20385.
- [93] J. C. Duchi and H. Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, June 2021. ISSN: 0090-5364, 2168-8966. DOI: 10.1214/20-AOS2004.
- [94] M. H. Duong and J. Tugaut. Coupled McKean-vlasov diffusions: wellposedness, propagation of chaos and invariant measures. *Stochastics*, 92(6):900–943, Aug. 17, 2020. ISSN: 1744-2508. DOI: 10.1080/17442508.2019.1677663.
- [95] B. Düring, P. Markowich, J.-F. Pietschmann, and M.-T. Wolfram. Boltzmann and Fokker-Planck equations modelling opinion formation in the presence of strong leaders. en. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2112):3687–3708, Dec. 2009. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2009.0239.
- [96] A. Durmus and É. Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [97] P. Dvurechensky and J.-J. Zhu. Analysis of kernel mirror prox for measure optimization. In S. Dasgupta, S. Mandt, and Y. Li, editors, *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 2350–2358. PMLR, May 2024.
- [98] I. Ekeland and R. Temam. *Analyse Convexe et Problèmes Variationnels*. Dunod, Gauthier-Villars, Paris, 1974.
- [99] C. Elbar and F. Santambrogio. A cross-diffusion system with independent drifts and fast diffusion, 2025. arXiv: 2510.07937.
- [100] E. E. Espejo, A. Stevens, and J. J. L. Velazquez. Simultaneous finite time blow-up in a two-species model for chemotaxis. *Analysis*, 29(3), Jan. 2009. ISSN: 0174-4747. DOI: 10.1524/anly.2009.1029.
- [101] L. Evans. *Partial Differential Equations*. American Mathematical Society, 2nd edition, 2010. ISBN: 978-0-8218-4974-3.
- [102] J. H. M. Evers, R. C. Fetecau, and T. Kolokolnikov. Equilibria for an Aggregation Model with Two Species. en. *SIAM Journal on Applied Dynamical Systems*, 16(4):2287–2338, Jan. 2017. ISSN: 1536-0040. DOI: 10.1137/16M1109527.
- [103] F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*, volume I–II. Springer, New York, 2003.
- [104] F. Facchinei and J.-S. Pang. *Nash equilibria: the variational approach*. In *Convex Optimization in Signal Processing and Communications*. D. P. Palomar and Y. C. Eldar, editors. Cambridge University Press, 2009, pages 443–493.
- [105] C. Falcó, R. E. Baker, and J. A. Carrillo. A local continuum model of cell-cell adhesion. *SIAM J. Appl. Math.*, 84(3):S17–S42, 2024. ISSN: 0036-1399, 1095-712X. DOI: 10.1137/22M1506079.

- [106] K. Fan. Some properties of convex sets related to fixed point theorems. *Mathematische Annalen*, 266(4):519–537, 1984.
- [107] E. A. Feinberg, P. O. Kasyanov, and Y. Liang. Fatou’s lemma for weakly converging measures under the uniform integrability condition, Mar. 27, 2019. DOI: 10.48550/arXiv.1807.07931. arXiv: 1807.07931 [math].
- [108] T. Fiez, B. Chasnov, and L. Ratliff. Implicit learning dynamics in stackelberg games: equilibria characterization, convergence analysis, and empirical study. In *International Conference on Machine Learning*, pages 3133–3144. PMLR, 2020.
- [109] T. Fiez, L. Ratliff, E. Mazumdar, E. Faulkner, and A. Narang. Global convergence to local minmax equilibrium in classes of nonconvex zero-sum games. *Advances in Neural Information Processing Systems*, 34:29049–29063, 2021.
- [110] T. Fiez and L. J. Ratliff. Local convergence analysis of gradient descent ascent with finite timescale separation. In *Proceedings of the International Conference on Learning Representation*, 2021.
- [111] D. Fonseca and M. Junca. Decision-dependent Distributionally Robust Optimization, Mar. 2023. DOI: 10.48550/arXiv.2303.03971. arXiv:2303.03971 [math].
- [112] M. D. Francesco and S. Fagioli. Measure solutions for non-local interaction PDEs with two species. *Nonlinearity*, 26(10):2777–2808, Oct. 1, 2013. ISSN: 0951-7715, 1361-6544. DOI: 10.1088/0951-7715/26/10/2777.
- [113] W. Gangbo, A. R. Mészáros, C. Mou, and J. Zhang. Mean field games master equations with nonseparable Hamiltonians and displacement monotonicity. en. *The Annals of Probability*, 50(6):2178–2217, Nov. 2022. ISSN: 0091-1798, 2168-894X. DOI: 10.1214/22-AOP1580.
- [114] B. Gao and L. Pavel. Continuous-Time Convergence Rates in Potential and Monotone Games. en. *SIAM Journal on Control and Optimization*, 60(3):1712–1731, June 2022. ISSN: 0363-0129, 1095-7138. DOI: 10.1137/20M1381873.
- [115] R. Gao and A. Kleywegt. Distributionally Robust Stochastic Optimization with Wasserstein Distance. en. *Mathematics of Operations Research*, 48(2):603–655, May 2023. ISSN: 0364-765X, 1526-5471. DOI: 10.1287/moor.2022.1275.
- [116] C. A. García Trillos and N. García Trillos. On adversarial robustness and the use of Wasserstein ascent-descent dynamics to enforce it. en. *Information and Inference: A Journal of the IMA*, 13(3):iaae018, July 2024. ISSN: 2049-8772. DOI: 10.1093/imaiai/iaae018.
- [117] B. K. Ghosh. Probability inequalities related to markov’s theorem. *The American Statistician*, 56(3):186–190, 2002. ISSN: 00031305.
- [118] V. Giunta, T. Hillen, M. Lewis, and J. R. Potts. Local and global existence for nonlocal multispecies advection-diffusion models. *SIAM Journal on Applied Dynamical Systems*, 21(3):1686–1708, Sept. 2022. ISSN: 1536-0040. DOI: 10.1137/21M1425992.
- [119] K. Glasner. Segregation and domain formation in non-local multi-species aggregation equations. *Phys. D*, 456:Paper No. 133936, 11, 2023. ISSN: 0167-2789, 1872-8022. DOI: 10.1016/j.physd.2023.133936.

- [120] I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [121] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [122] P. J. Graber and A. R. Mészáros. On monotonicity conditions for mean field games. *Journal of Functional Analysis*, 285(9):110095, Nov. 2023. ISSN: 0022-1236. DOI: 10.1016/j.jfa.2023.110095.
- [123] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. *International Journal of Computer Vision*, 60(3):225–240, 2004. ISSN: 1573-1405. DOI: 10.1023/B:VISI.0000036836.66311.97.
- [124] M. Hardt, N. Megiddo, C. Papadimitriou, and M. Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS ’16, pages 111–122, New York, NY, USA. Association for Computing Machinery, Jan. 14, 2016. ISBN: 9781450340571. DOI: 10.1145/2840728.2840730.
- [125] G. Heinze, J.-F. Pietschmann, and M. Schmidtchen. Nonlocal cross-interaction systems on graphs: nonquadratic finslerian structure and nonlinear mobilities, Jan. 11, 2022. DOI: 10.48550/arXiv.2107.11289. arXiv: 2107.11289 [math].
- [126] D. Horstmann and M. Lucia. Nonlocal elliptic boundary value problems related to chemotactic movement of mobile species. *RIMS Kokyuroku Bessatsu*:39–72, 2009.
- [127] N. Isobe and S. Shimoyama. On gradient descent-ascent flows in metric spaces, June 2025. DOI: 10.48550/arXiv.2506.20258. arXiv:2506.20258.
- [128] Z. Izzo, L. Ying, and J. Zou. How to learn when data reacts to your model: performative gradient descent. In *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning, pages 4641–4650. PMLR, July 1, 2021.
- [129] P. Jordan and J. V. Neumann. On Inner Products in Linear, Metric Spaces. *The Annals of Mathematics*, 36(3):719, July 1935. ISSN: 0003486X. DOI: 10.2307/1968653.
- [130] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck Equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [131] A. Jungel, S. Portisch, and A. Zurek. Nonlocal cross-diffusion systems for multi-species populations and networks. *Nonlinear Analysis*, 219:112800, June 1, 2022. ISSN: 0362-546X. DOI: 10.1016/j.na.2022.112800.
- [132] M. R. Karimi, Y.-P. Hsieh, and A. Krause. Sinkhorn Flow as Mirror Flow: A Continuous-Time Framework for Generalizing the Sinkhorn Algorithm. en. In *AISTATS*, pages 4186–4194. PMLR, Apr. 2024.

- [133] D. Kinderlehrer and G. Stampacchia. *An Introduction to Variational Inequalities and Their Applications*. en. Society for Industrial and Applied Mathematics, Jan. 2000. ISBN: 9780898719451. DOI: 10.1137/1.9780898719451.
- [134] B. Knaster, K. Kuratowski, and S. Mazurkiewicz. Ein Beweis des Fixpunktsatzes für n -dimensionale Simplexe. *Fundamenta Mathematicae*, 14(1):132–137, 1929. DOI: 10.4064/fm-14-1-132-137.
- [135] Y. Kobayashi. Difference approximation of Cauchy problems for quasi-dissipative operators and generation of nonlinear semigroups. *Journal of the Mathematical Society of Japan*, 27(4):640–665, 1975. DOI: 10.2969/jmsj/02740640.
- [136] H. Komiya. Elementary proof for Sion’s minimax theorem. *Kodai Mathematical Journal*, 11(1):5–7, 1988. DOI: 10.2996/kmj/1138038812.
- [137] A. Korotin, V. Egiazarian, A. Asadulaev, A. Safin, and E. Burnaev. Wasserstein-2 generative networks. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [138] G. Korpelevich. The Extragradient Method for Finding Saddle Points and Other Problems. *Ekonomika i matematicheskie metody*, 12(4):747–756, 1976.
- [139] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: theory and applications in machine learning. In S. Netessine, D. Shier, and H. J. Greenberg, editors, *Operations Research & Management Science in the Age of Analytics*, pages 130–166. INFORMS, Oct. 2019. ISBN: 9780990615330. DOI: 10.1287/educ.2019.0198.
- [140] D. Kwon, Y. Fan, and K. Lee. Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance. en. *Advances in Neural Information Processing Systems*, 35:20205–20217, Dec. 2022.
- [141] J. Lang, M. Spear, and S. F. Wu. Social manipulation of online recommender systems. In L. Bolc, M. Makowski, and A. Wierzbicki, editors, *Social Informatics*, Lecture Notes in Computer Science, pages 125–139, Berlin, Heidelberg. Springer, 2010. ISBN: 9783642165672. DOI: 10.1007/978-3-642-16567-2_10.
- [142] T. Laokul and B. Panyanak. A generalization of the (cn) inequality and its applications. *Carpathian Journal of Mathematics*, 36(1):81–90, 2020. ISSN: 15842851.
- [143] R.-A. Lascu, M. B. Majka, and L. Szpruch. A Fisher-Rao gradient flow for entropic mean-field min-max games. en. *Transactions on Machine Learning Research*, May 2024. ISSN: 2835-8856.
- [144] R. Lascu, M. B. Majka, and L. Szpruch. Mirror descent-ascent for mean-field min-max problems. *CoRR*, 2024. DOI: 10.48550/ARXIV.2402.08106.
- [145] J.-M. Lasry and P.-L. Lions. Mean field games. en. *Japanese Journal of Mathematics*, 2(1):229–260, Mar. 2007. ISSN: 1861-3624. DOI: 10.1007/s11537-007-0657-8.
- [146] Q. Lei, W. Hu, and J. D. Lee. Near-optimal linear regression under distribution shift. *ArXiv*, June 23, 2021.

- [147] G. Li and Y. Yao. Two-species competition model with chemotaxis: well-posedness, stability and dynamics. *Nonlinearity*, 35(3):1329, Feb. 3, 2022. ISSN: 0951-7715. DOI: 10.1088/1361-6544/ac4a8d.
- [148] M. Li, J. Yu, T. Li, and C. Meng. Importance sparsification for sinkhorn algorithm. *Journal of Machine Learning Research*, 24(247):1–44, 2023.
- [149] W. Li, P. Yin, and S. Osher. Computations of Optimal Transport Distance with Fisher Information Regularization. en. *Journal of Scientific Computing*, 75(3):1581–1595, June 2018. ISSN: 1573-7691. DOI: 10.1007/s10915-017-0599-0.
- [150] E. H. Lieb. Sharp constants in the hardy-littlewood-sobolev and related inequalities. *Annals of Mathematics*, 118(2):349–374, 1983. ISSN: 0003-486X. DOI: 10.2307/2007032.
- [151] F. Lin, X. Fang, and Z. Gao. Distributionally robust optimization: a review on theory and applications. *Numerical Algebra, Control and Optimization*, 12(1):159–212, 2022.
- [152] L. Liu, Y. Zhang, Z. Yang, R. Babanezhad, and Z. Wang. Infinite-dimensional optimization for zero-sum games via variational transport. In *Proceedings of the 38th International Conference on Machine Learning*. International Conference on Machine Learning, pages 7033–7044. PMLR, July 1, 2021.
- [153] S. Liu, H. Sun, and H. Zha. A Particle-Evolving Method for Approximating the Optimal Transport Plan. en. In F. Nielsen and F. Barbaresco, editors, *Geometric Science of Information*, pages 878–887, Cham. Springer International Publishing, 2021. ISBN: 9783030802097. DOI: 10.1007/978-3-030-80209-7_94.
- [154] Y. Lu. Two-scale gradient descent ascent dynamics finds mixed nash equilibria of continuous games: a mean-field perspective, Jan. 8, 2023. DOI: 10.48550/arXiv.2212.08791. arXiv: 2212.08791[cs,math,stat].
- [155] Y. Lu and P. Monmarché. Convergence of Time-Averaged Mean Field Gradient Descent Dynamics for Continuous Multi-Player Zero-Sum Games, May 2025. DOI: 10.48550/arXiv.2505.07642. arXiv:2505.07642.
- [156] Y.-A. Ma, Y. Chen, C. Jin, N. Flammarion, and M. I. Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885, 2019.
- [157] C. Ma and L. Ying. Provably convergent quasistatic dynamics for mean-field two-player zero-sum games. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [158] J. Maas, M. Rumpf, C. Schönlieb, and S. Simon. A generalized model for optimal transport of images including dissipation and density modulation. en. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1745–1769, 2015. DOI: 10.1051/m2an/2015043.
- [159] A. Mackey, T. Kolokolnikov, and A. L. Bertozzi. Two-species particle aggregation and stability of co-dimension one solutions. *Discrete and Continuous Dynamical Systems*, 19(5):1411–1436, 2014.

- [160] F. Marcellini. ODE-PDE models in traffic flow dynamics. en. *Bulletin of the Brazilian Mathematical Society, New Series*, 47(2):533–544, June 2016. ISSN: 1678-7714. DOI: 10.1007/s00574-016-0167-5.
- [161] E. Mazumdar, K. Panaganti, and L. Shi. Tractable Multi-Agent Reinforcement Learning through Behavioral Economics. en. In Oct. 2024.
- [162] R. J. McCann. A convexity principle for interacting gases. *Advances in Mathematics*, 128(1):153–179, June 1, 1997. ISSN: 0001-8708. DOI: 10.1006/aima.1997.1634.
- [163] R. J. McCann. A convexity principle for interacting gases. *Adv. Math.*, 128(1):153–179, 1997. ISSN: 0001-8708.
- [164] A. R. Mészáros and C. Mou. Mean Field Games Systems under Displacement Monotonicity. en. *SIAM Journal on Mathematical Analysis*, 56(1):529–553, Feb. 2024. ISSN: 0036-1410, 1095-7154. DOI: 10.1137/22M1534353.
- [165] A. R. Mészáros and G. Parker. Existence theory for a cross-diffusion system with independent drifts: mixing dynamics, 2025. arXiv: 2603.18770.
- [166] A. Mielke and T. Roubíček. *Rate-Independent Systems: Theory and Application*. Applied Mathematical Sciences. Springer New York, 2015. ISBN: 9781493927067.
- [167] A. Mielke. Geodesic convexity of the relative entropy in reversible Markov chains. en. *Calculus of Variations and Partial Differential Equations*, 48(1):1–31, Sept. 2013. ISSN: 1432-0835. DOI: 10.1007/s00526-012-0538-8.
- [168] J. Miller, C. Hsu, J. Troutman, J. Perdomo, T. Zrnic, L. Liu, Y. Sun, L. Schmidt, and M. Hardt. Whynot, 2020. DOI: 10.5281/zenodo.3875775.
- [169] J. Miller, K. Krauth, B. Recht, and L. Schmidt. The effect of natural distribution shift on question answering models. In *Proceedings of the 37th International Conference on Machine Learning*. International Conference on Machine Learning, pages 6905–6916. PMLR, Nov. 21, 2020.
- [170] J. P. Miller, J. C. Perdomo, and T. Zrnic. Outside the echo chamber: optimizing the performative risk. In *International Conference on Machine Learning*, pages 7710–7720. PMLR, 2021.
- [171] M. Muratori and G. Savaré. Gradient flows and evolution variational inequalities in metric spaces. I: Structural properties. *J. Funct. Anal.*, 278(4):108347, 67, 2020. ISSN: 0022-1236, 1096-0783. DOI: 10.1016/j.jfa.2019.108347.
- [172] A. Narang, E. Faulkner, D. Drusvyatskiy, M. Fazel, and L. Ratliff. Learning in Stochastic Monotone Games with Decision-Dependent Data. en. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 5891–5912. PMLR, May 2022.
- [173] A. Narang, E. Faulkner, D. Drusvyatskiy, M. Fazel, and L. J. Ratliff. Multiplayer Performative Prediction: Learning in Decision-Dependent Games. *Journal of Machine Learning Research*, 24(202):1–56, 2023. ISSN: 1533-7928.

- [174] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [175] R. H. Nochetto and G. Savaré. Nonlinear evolution governed by accretive operators in Banach spaces: error control and applications. *Math. Models Methods Appl. Sci.*, 16(3):439–477, 2006. ISSN: 0218-2025.
- [176] F. Otto and C. Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, June 1, 2000. ISSN: 0022-1236. DOI: 10.1006/jfan.1999.3557.
- [177] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm. Partial Differential Equations*, 26(1-2):101–174, 2001. ISSN: 0360-5302.
- [178] A. Pandolfi, L. Stainier, and M. Ortiz. An optimal-transport finite-particle method for mass diffusion. *Computer Methods in Applied Mechanics and Engineering*, 416:116385, Nov. 2023. ISSN: 0045-7825. DOI: 10.1016/j.cma.2023.116385.
- [179] N. Papadakis, G. Peyré, and E. Oudet. Optimal Transport with Proximal Splitting. en. *SIAM Journal on Imaging Sciences*, 7(1):212–238, Jan. 2014. ISSN: 1936-4954. DOI: 10.1137/130920058.
- [180] P. A. Pathak and T. Sönmez. School admissions reform in chicago and england: comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103(1):80–106, Feb. 2013. ISSN: 0002-8282. DOI: 10.1257/aer.103.1.80.
- [181] M. A. Peletier and A. Schlichting. Cosh gradient systems and tilting. *Nonlinear Analysis. Variational Models for Discrete Systems*, 231:113094, June 2023. ISSN: 0362-546X. DOI: 10.1016/j.na.2022.113094.
- [182] J. Perdomo, T. Zrnic, C. Mendler-Dünner, and M. Hardt. Performative prediction. In *Proceedings of the 37th International Conference on Machine Learning*. International Conference on Machine Learning, pages 7599–7609. PMLR, Nov. 21, 2020.
- [183] B. Perthame. *Parabolic Equations in Biology: Growth, reaction, movement and diffusion*. Lecture Notes on Mathematical Modelling in the Life Sciences. Springer International Publishing, Cham, 2015. ISBN: 9783319194998. DOI: 10.1007/978-3-319-19500-1.
- [184] E. Posner. Random coding strategies for minimum entropy. *IEEE Transactions on Information Theory*, 21(4):388–391, July 1975. ISSN: 1557-9654. DOI: 10.1109/TIT.1975.1055416.
- [185] J. Quinonero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence. Dataset shift in machine learning. MIT Press. URL: <https://mitpress.mit.edu/9780262545877/dataset-shift-in-machine-learning/>.
- [186] J. Rabin and N. Papadakis. Convex color image segmentation with optimal transport distances. In J.-F. Aujol, M. Nikolova, and N. Papadakis, editors, *Scale Space and Variational Methods in Computer Vision*, pages 256–269, Cham. Springer International Publishing, 2015. ISBN: 978-3-319-18461-6.
- [187] M. Ray, L. J. Ratliff, D. Drusvyatskiy, and M. Fazel. Decision-dependent risk minimization in geometrically decaying dynamic environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36 of number 7, pages 8081–8088, 2022.

- [188] A. Rees-Jones. Suboptimal behavior in strategy-proof mechanisms: evidence from the residency match. *Games and Economic Behavior*. Special Issue in Honor of Lloyd Shapley: Seven Topics in Game Theory, 108:317–330, Mar. 1, 2018. ISSN: 0899-8256. DOI: 10.1016/j.geb.2017.04.011.
- [189] R. T. Rockafellar. Monotone operators associated with saddle-functions and minimax problems. In *Nonlinear Functional Analysis (Proc. Sympos. Pure Math., Vol. XVIII, Part 1, Chicago, Ill., 1968)*. Volume XVIII, Part 1, Proc. Sympos. Pure Math. Pages 241–250. Amer. Math. Soc., Providence, RI, 1970.
- [190] R. T. Rockafellar. On the maximality of sums of nonlinear monotone operators. *Transactions of the American Mathematical Society*, 149(1):75–88, 1970.
- [191] J. B. Rosen. Existence and Uniqueness of Equilibrium Points for Concave N-Person Games. *Econometrica*, 33(3):520–534, 1965. ISSN: 0012-9682. DOI: 10.2307/1911749.
- [192] W. Rudin. *Functional analysis*. International series in pure and applied mathematics. McGraw-Hill, New York, 2nd ed edition, 1991. 424 pages. ISBN: 9780070542365.
- [193] F. Santambrogio. Dealing with Moment Measures via Entropy and Optimal Transport. *Journal of Functional Analysis*, 271(2):418–436, 2016. DOI: <https://doi.org/10.1016/j.jfa.2016.04.009>.
- [194] F. Santambrogio. *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*. en, volume 87 of *Progress in Nonlinear Differential Equations and Their Applications*. Springer International Publishing, Cham, 2015. ISBN: 9783319208275. DOI: 10.1007/978-3-319-20828-2.
- [195] D. Sanz-Alonso, A. Stuart, and A. Taeb. *Inverse Problems and Data Assimilation*. Cambridge University Press, 1st edition, July 2023. ISBN: 9781009414319. DOI: 10.1017/9781009414319.
- [196] S. Sastry. *Nonlinear Systems: Analysis, Stability, and Control*. Interdisciplinary Applied Mathematics. Springer New York, 1999.
- [197] V. Seguy, B. B. Damodaran, R. Flamary, N. Courty, A. Rolet, and M. Blondel. Large-scale optimal transport and mapping estimation. In *ICLR*, pages 1–15, 2018.
- [198] E. Shekhtman and S. Dean. Strategic Usage in a Multi-Learner Setting. en. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pages 2665–2673. PMLR, Apr. 2024.
- [199] S. Simons. Minimax Theorems and Their Proofs. en. In D.-Z. Du and P. M. Pardalos, editors, *Minimax and Applications*, pages 1–23. Springer US, Boston, MA, 1995. ISBN: 9781461335573. DOI: 10.1007/978-1-4613-3557-3_1.
- [200] J. M. Steele. *The Cauchy-Schwarz Master Class: An Introduction to the Art of Mathematical Inequalities*. Cambridge University Press, Apr. 26, 2004. 320 pages. ISBN: 9780521546775.
- [201] T. Tatarenko and M. Kamgarpour. Learning Nash Equilibria in Monotone Games. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 3104–3109, Dec. 2019. DOI: 10.1109/CDC40024.2019.9029659. ISSN: 2576-2370.

- [202] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. en. Springer Series in Statistics. Springer Cham, 2nd edition, 2023.
- [203] C. Villani. *Optimal transport, old and new*. Lecture Notes for the 2005 Saint-Flour summer school. Springer, 2007.
- [204] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003, pages xvi+370. ISBN: 0-8218-3312-X. DOI: 10.1090/gsm/058.
- [205] J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.
- [206] G. Wang and L. Chizat. Local convergence of mean-field Langevin dynamics: from gradient flows to linearly monotone games, Feb. 2026. DOI: 10.48550/arXiv.2602.11999. arXiv:2602.11999.
- [207] G. Wang and L. Chizat. An exponentially converging particle method for the mixed nash equilibrium of continuous games, Apr. 13, 2023. DOI: 10.48550/arXiv.2211.01280. arXiv: 2211.01280 [cs, math].
- [208] G. Wang and L. Chizat. Open problem: Convergence of single-timescale mean-field Langevin descent-ascent for two-player zero-sum games. *Proceedings of Machine Learning Research*. Conference on Learning Theory, 247, 2024.
- [209] O. Wiles, S. Goyal, F. Stimberg, S.-A. Rebuffi, I. Ktena, K. Dvijotham, and A. Cemgil. A fine-grained analysis on distribution shift. *ArXiv*, Oct. 21, 2021.
- [210] G. Wolansky. Chemotactic systems in the presence of conflicts: A new functional inequality. *Journal of Differential Equations*, 261(9):5119–5143, Nov. 2016. ISSN: 0022-0396. DOI: 10.1016/j.jde.2016.07.021.
- [211] G. Wolansky. Multi-components chemotactic system in the absence of conflicts. en. *European Journal of Applied Mathematics*, 13(6):641–661, Dec. 2002. ISSN: 1469-4425, 0956-7925. DOI: 10.1017/S0956792501004843.
- [212] Z. Wu and K. Sun. Distributionally robust optimization with wasserstein metric for multi-period portfolio selection under uncertainty. *Applied Mathematical Modelling*, 117:513–528, 2023. ISSN: 0307-904X. DOI: <https://doi.org/10.1016/j.apm.2022.12.037>.
- [213] B. J. Zhang and M. A. Katsoulakis. A mean-field games laboratory for generative modeling, Oct. 2023. DOI: 10.48550/arXiv.2304.13534. arXiv:2304.13534 [cs, stat].
- [214] L. Zhang, J. Yang, and R. Gao. A Short and General Duality Proof for Wasserstein Distributionally Robust Optimization. en. *Operations Research*:opre.2023.0135, July 2024. ISSN: 0030-364X, 1526-5463. DOI: 10.1287/opre.2023.0135.
- [215] T. Zrnic, E. Mazumdar, S. Sastry, and M. Jordan. Who leads and who follows in strategic classification? In *Advances in Neural Information Processing Systems*, volume 34, pages 15257–15269. Curran Associates, Inc., 2021.

Part II

DISTRIBUTED CONTROL

*Chapter 7***CHARACTERIZING CONTROLLABILITY AND OBSERVABILITY FOR
SYSTEMS WITH LOCALITY, COMMUNICATION, AND ACTUATION
CONSTRAINTS**

L. Conger, Y. Lin, A. Wierman, and E. Mazumdar. Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pages 8491–8497, Dec. 2024. DOI: 10.1109/CDC56724.2024.10886635. ISSN: 2576-2370

ABSTRACT

This chapter presents a closed-form notion of controllability and observability for systems with communication delays, actuation delays, and locality constraints. The formulation reduces to classical notions of controllability and observability in the unconstrained setting. As a consequence of our formulation, we show that the addition of locality and communication constraints may not affect the controllability and observability of the system, and we provide an efficient sufficient condition under which this phenomenon occurs. This contrasts with actuation and sensing delays, which cause a gradual loss of controllability and observability as the delays increase. We illustrate our results using linearized swing equations for the power grid, showing how actuation delay and locality constraints affect controllability.

7.1 Introduction

Sensor and actuator placement is a crucial problem in the design of real world systems in application areas ranging from power systems and traffic networks to environmental monitoring. Taking power systems as an example, devices such as generators and transformers can provide additional inputs to actuate the system, while other devices like meters report information about the system's state. A system designer may face the challenge of deploying a limited number of such devices to different locations in a power system with the goal of keeping the system stable or performing well in other control tasks. While sensor and actuator placement is an old problem that has been well studied [51, 41], the proliferation of low-cost sensors and the decentralization of control in large systems like power systems has led to new challenges. Many real-world systems are subject to delayed and information-constrained sensors that can significantly impact the quality of the control. However, we currently lack tools to rigorously characterize the performance loss under various sensor configurations. This leads us to study the problem of quantifying the quality of a set of controllers and sensors in systems subject to actuation and sensing delays as well as communication constraints.

While task-specific metrics like cost are commonly used to measure the quality of a set of sensors and actuators, there are many settings where the cost is unknown or time varying; here task-agnostic notions of controllability and observability of dynamical systems [41, 39] offer a suitable proxy. Controllability, defined as the determinant of the controllability Gramian [50], is the volume over which a control input with magnitude one can drive the state to the zero state. Likewise, observability, defined as the determinant of the observability Gramian, is the volume of the signal-to-noise ratio along any direction in the state space. These metrics are used to compare different sets of actuators and sensors for dynamical systems; a system with larger controllability, given fixed dynamics, is preferred to one with smaller controllability.

Although the notions of controllability and observability are well-defined for classic control problems, new challenges arise from the decentralized control of modern large-scale systems [107, 70, 67, 135], where the controllers may face different constraints on communication speed, actuation delays, and/or locality. Such constraints, which we refer to as *system level constraints*, cannot easily be captured in the classical formulations of controllability/observability based on the Gramian or control inputs [50]. Fortunately, a recently proposed framework called *system level synthesis (SLS)* [7] handles system level constraints efficiently, where the key technique is to formulate the original control problem in terms of closed-loop maps. Some recent works address controllability and observability in the specific setting of systems with delays [129, 144, 48]; SLS allows for a much broader class of constraints, including delays. Using SLS, a recent work [62] shows sufficient conditions under which locality constraints on model predictive control (MPC) leads to *no loss in*

optimality for any convex global cost function. However, even with SLS, it is unclear whether one can formally extend the notion of controllability and observability to include system level constraints faced by large-scale networked systems, and whether such measures can be computed efficiently to quantify the loss of performance under different combinations of system level constraints.

Contributions. We propose a formal definition of controllability and observability subject to system level constraints (Section 7.3). Our definitions use the SLS framework and we draw connections to classical distributed control, differential privacy, and speed-accuracy tradeoffs in neurons. We provide a provable closed-form solution for our definitions (Algorithm 1, Proposition 7.8), which strictly generalize the classic notions of controllability and observability (Theorem 7.6).

Importantly, we establish rank conditions for system level constraints under which there is no performance loss compared with the unconstrained system (Theorem 7.11). This rigorously characterizes the empirical observation that locality and communication constraints do not necessarily decrease controllability and observability. In contrast, actuation delay does not satisfy this condition, consistent with numerical results showing a gradual decrease in controllability as actuation delay increases. We perform numerical experiments on a synthetic example and a real-world model of linearized and discretized swing equations for the power grid (Section 7.5).

Notation. Let $N_x \in \mathbb{Z}_+$ and $N_u \in \mathbb{Z}_+$ be the state and input dimensions respectively. Let $x_0 := x(0)$ and $x_T := x(T)$. A single index such as $Z[i]$ indicates the i^{th} row or block of rows of a matrix, and a double index such as $A[i, j]$ indicates the i^{th} row and j^{th} column of a matrix. Let $A[:, j]$ or $A[i, :]$ denote all rows or columns of A respectively, and $\text{blkdiag}([A, B, C])$ denote a matrix with submatrices A , B , and C on the block diagonals. Brackets around a scalar indicates the set of indices $[a] := [1, \dots, a]$. A vector over a variable, such as $\vec{\phi}_x$, is the vectorization of matrix, where the columns are stacked into a vector. A hat symbol over a matrix indicates the vectorization of matrix multiplication, that is, $\hat{A}$ is defined such that if $Y = AX$ for two matrices X and Y , then $\vec{y} = \hat{A}\vec{x}$. We define $\text{supp } A$ to be an indicator matrix such that $(\text{supp } A)[i, j] = 0$ if $A[i, j]$ is constrained to be zero. $A^\dagger$ is the Moore-Penrose pseudoinverse of A .

7.2 Motivation & Background

We begin by introducing the problem of computing the controllability and observability of a dynamical system and illustrating the challenges that come from considering constraints in the classical formulation of the problem. We then show how system level synthesis (SLS) offers a framework for computing the controllability and observability of a system while incorporating system-level constraints.

7.2.1 Classical Controllability and Observability

We consider linear systems with dynamics given by

$$x(t+1) = Ax(t) + Bu(t), \quad y(t) = Cx(t), \quad (7.1)$$

where $x(t) \in \mathbb{R}^{N_x}$, $u(t) \in \mathbb{R}^{N_u}$, and $y(t) \in \mathbb{R}^{N_y}$. A system is said to be *controllable* in time T if it can be driven from any state x_0 to any terminal state x_T , which holds when the controllability Gramian W_c is full-rank. The controllability Gramian provides additional information about how easily the system is controlled in any direction, and the controllability volume is denoted V_c ,

$$W_c = \sum_{i=0}^{T-1} A^i B B^\top (A^\top)^i, \quad V_c = \det W_c. \quad (7.2)$$

The magnitude of each eigenvalue of W_c is the distance in the direction of the eigenvector that one unit of control provides (after accounting for the trajectory resulting from the initial condition alone). The determinant of W_c is then the state space volume that is reachable with one unit of control input. If the system is not controllable in a certain direction, the volume is zero. However, this computational method assumes that any sequence of control input can be selected; this assumption is not true when information among states is restricted, actuation is delayed, or controllers are dependent on only local information.

Similarly, for an unconstrained system, one can define the observability Gramian and observability respectively, as

$$W_o = \sum_{i=0}^{T-1} A^{\top i} C^\top C A^i, \quad V_o = \det W_o. \quad (7.3)$$

The magnitude of each eigenvalue of the Gramian W_o is the signal-to-noise ratio for a state in the direction of the corresponding eigenvector. The determinant of W_o corresponds to the ‘observability volume’ of the system. Similarly to controllability, this method assumes instantaneous and global observations for estimation. This does not apply to large-scale systems where state estimation depends only on local observation, and the observations are not immediate.

This paper proposes a new method for computing controllability and observability that accounts for delayed actuation and observation, local control and information sharing. Critical to the proposed approach is the system level synthesis (SLS) parameterization of the linear system. In what follows, we consider the state estimation problem as the dual of the control problem; in what follows, one can substitute

$$A \leftrightarrow A^\top, \quad B \leftrightarrow C^\top.$$

to obtain the corresponding observability framework.

7.2.2 System Level Synthesis (SLS)

The SLS parametrization of the state dynamics in (7.1) is

$$\phi_x(t+1) = A\phi_x(t) + B\phi_u(t), \quad (7.4)$$

where $x(t) = \phi_x(t)x_0$ and $u(t) = \phi_u(t)x_0$, $\phi_x(t) \in \mathbb{R}^{N_x \times N_x}$, $\phi_u(t) \in \mathbb{R}^{N_u \times N_x}$, and $\phi_x(0) = I$. Let the closed-loop maps ϕ_x , ϕ_u , and ϕ be defined as

$$\phi_x = \begin{bmatrix} \phi_x(0) \\ \vdots \\ \phi_x(T-1) \end{bmatrix}, \quad \phi_u = \begin{bmatrix} \phi_u(0) \\ \vdots \\ \phi_u(T-1) \end{bmatrix}, \quad \phi = \begin{bmatrix} \phi_x \\ \phi_u \end{bmatrix}.$$

A prior result from [7, Theorem 2] guarantees that the subspace defined by ϕ parameterizes all possible system responses; optimizing over the state feedback controller $u(t) = K(t)x(t)$ is equivalent to optimizing over ϕ .

This parameterization allows us to explicitly write constraints on ϕ_x and ϕ_u which encode the communication, actuation, and sensing constraints. To enforce that the i^{th} entry of x_0 does not affect the j^{th} entry of the state or input at time t , we set $\phi_x(t)[j, i] = 0$ or $\phi_u(t)[j, i] = 0$ respectively. We write the system level constraints as

$$S_x \vec{\phi}_x = 0, \quad S_u \vec{\phi}_u = 0, \quad (7.5)$$

where S_x (S_u resp.) has a single nonzero entry in each row, with the column corresponding to the entry of ϕ_x (ϕ_u resp.) that is constrained to be zero. We define the set of closed-loop maps that satisfy (7.5) as $\mathcal{S}$, such that $\phi \in \mathcal{S}$. Equipped with the ability to specify system level constraints, in the next section we present the definitions of controllability under these constraints.

7.3 Constrained Systems

Our focus in this paper is systems that are subject to actuation delay, communication delay and restrictions, and locality constraints. As discussed in the previous section, these constraints can be expressed as support constraints on the closed-loop maps ϕ_x and ϕ_u .

We formally define the different types of constraints considered in this work as follows. We assume the system is formed by a network of nodes, and each entry of the system state $x(t)$ corresponds to the local state of a node. The network is characterized by a directed graph whose adjacency matrix $\mathcal{G}[i, j]$ is defined such that $\mathcal{G}[i, j] = 1$ if $A[i, j] \neq 0$; $\mathcal{G}[i, j] = 0$ otherwise. We denote the graph distance between two nodes i and j as $\text{dist}_{\mathcal{G}}(i, j)$.

- **Actuation delay:** The system is subject to a τ -step actuation delay if the actuator can react to state information at least τ steps after, i.e., $\phi_u(t) = 0$, for $t < \tau$.

- **Communication speed:** The communication speed is v if the propagation of communication signals is v times the dynamics, i.e., $\text{supp } \phi_x(t) = \text{supp } \mathcal{G}^{\lfloor vt \rfloor}$ and $\text{supp } \phi_u(t) = \text{supp } (|B^\top| \text{supp } \phi_x(t))$, so that $x(t)[j]$ cannot depend on $x_0[i]$ if $\text{dist}_{\mathcal{G}}(i, j) > vt$. Likewise, $u(t)[i]$ cannot depend on $x_0[j]$ if $B[k, i] = 0$ for all k such that $\text{dist}_{\mathcal{G}}(j, k) \leq vt$.
- **Locality:** The system is subject to a locality constraint with radius r if the state of a node can only be affected by the states of its r -hop neighbors, i.e., $\phi_x(t)[j, i] = 0$ if $\text{dist}_{\mathcal{G}}(i, j) > r$, and a controller uses information from the nearest r -hops away from the states on which it acts, specified as $\text{supp } \phi_u = \text{supp } (|B^\top| \text{supp } \phi_x)$.

We formulate the definition of controllability and observability for constrained systems using SLS operators. We consider the magnitude of any x_0 such that one unit of control can drive the state to the origin, again scaled by the effect of the dynamics without any control. Using this notion of controllability from (7.2), we define the controllability with constraints in Definition 7.1.

Definition 7.1. *The controllability Y_c of (7.1) with constraints (7.5) is given by*

$$Y_c = \text{volume} \left\{ A^T x_0 : h_c(x_0) \leq 1 \right\}, \text{ where}$$

$$h_c(x_0) := \min_{\phi_u} x_0^\top \left(\sum_{i=0}^{T-1} \phi_u(i)^\top \phi_u(i) \right) x_0 \quad (7.6)$$

$$\text{s.t. } S_c \vec{\phi}_u + C_c = 0.$$

The function h_c is the squared norm of the input signal, and $A^\top x_0$ is the scaled set of initial conditions. The constraint for the terminal state condition, $x_T = 0$, as well as support constraints for ϕ_x and ϕ_u , are specified by the form of S_c and C_c . Recall that we use the hat symbol over a matrix to indicate the vectorization of matrix multiplication. The constraint matrices S_c and C_c are given by

$$S_c = \begin{bmatrix} S_u \\ D_1 \\ D_2 \end{bmatrix}, \quad C_c = \begin{bmatrix} 0 \\ v_1 \\ \hat{A}^T \vec{I} \end{bmatrix}, \quad v_1 = \begin{bmatrix} S_x[1] \vec{I} \\ \vdots \\ S_x[T-1] \hat{A}^{T-1} \vec{I} \end{bmatrix},$$

$$D_1 = \begin{bmatrix} S_x[1] \hat{B} & & & \\ S_x[2] \hat{A} \hat{B} & S_x[2] \hat{B} & & \\ \vdots & & \ddots & \\ S_x[T-1] \hat{A}^{T-2} \hat{B} & \dots & & \hat{B} \end{bmatrix},$$

$$D_2 = \begin{bmatrix} \hat{A}^{T-1} \hat{B} & \dots & \hat{B} \end{bmatrix}.$$

The observability is likewise stated in Definition 7.2.

Definition 7.2. The observability Y_o of (7.1) with constraints (7.5) is given by

$$Y_o = \text{volume} \left\{ A^{T^\top} x_0 : h_o(x_0) \leq 1 \right\}, \text{ where}$$

$$h_o(x_0) := \min_{\phi_u} x_0^\top \left(\sum_{i=0}^{T-1} \phi_u(i)^\top \phi_u(i) \right) x_0 \quad (7.7)$$

$$\text{s.t. } S_o \vec{\phi}_u + C_o = 0,$$

The matrices S_o and C_o again utilize the terminal constraint $x_T = 0$, with S_o and C_o defined analogously to S_c and C_c .

Example 7.3 (Distributed Control). *In large-scale systems with fast propagation of the dynamics, such as the power grid, the speed at which sensors send information can be on the same order of magnitude as the dynamics [114] and understanding controllability and observability with delays is an important problem [91]. This means that information about a distant state may not be available for a number of time steps, and controllers must be designed accounting for this lack of information. Additionally, if a large disturbance hits a system, ensuring that the disturbance will not propagate further than r neighbors ensures that distant states will not be affected by a disturbance about which they only have delayed information [7]. The sensor and actuator co-design problem with constraints [83] is important for addressing these new challenges; comparing the controllability and observability of sensors and actuators for large-scale systems with these delays and constraints helps designers understand how much faster or more densely-placed hardware improves performance.*

Example 7.4 (Sensorimotor control). *In neurophysiology, a nerve is made up of several heterogeneous communication channels formed by different types of axons. To achieve the control objective, the design of the communication channels faces a strict speed-accuracy trade-off [88] over channels with lower delay and less information or a channel with longer delay and more information. Our constrained observability problem (7.7) provides a method to compare the observability of different sets of axons that satisfy a fixed metabolic cost.*

Example 7.5 (Privacy as Observability). *The problem of preserving privacy in decentralized control of large-scale networked systems has received considerable attention recently, including privacy through the perspective of observability [141, 53, 54, 30, 60]. Using the interpretation of observability as a signal-to-noise ratio of an observer, less observability potentially corresponds to less noise needed for obtaining some fixed level of privacy. In particular, our method provides tools for analyzing this signal-to-noise ratio with sensing delay and locality of sensors.*

7.4 Results

This section presents our main three technical results. Theorem 7.6 states that the controllability volume given by (7.6) is equivalent to classical controllability (7.2) when the constraint set is empty.

Algorithm 1 offers a provably correct closed-form solution for controllability, and Theorem 7.11 provides conditions under which adding constraints to the closed-loop map ϕ does not change the set of feasible states and inputs. As a consequence, the controllability does not change with the addition of the constraints (Corollary 7.12).

7.4.1 Connection to classical methods

Our first result shows that, in the setting without support constraints, the generalized definition of controllability (Definition 7.1) reduces to the classical definition (7.2). This shows that the formulation is consistent with classical controllability.

Theorem 7.6. *When $S_u = 0$ and $S_x = 0$, $Y_c^2 = V_c$ for all $A \in \mathbb{R}^{N_x \times N_x}$, $B \in \mathbb{R}^{N_x \times N_u}$.*

Proof. When the constraint set is empty, that is, when S_x and S_u are zero, the remaining constraint in S_c and C_c is the terminal constraint $x_T = 0$. Using $u = \phi_u x_0$, we can rewrite the objective of (7.6) as $\|u\|^2$. The constraint $S_c \vec{\phi}_u + C_c = 0$ can be written as $A^T x_0 + Du = 0$. Here, D takes the form $D = [A^{T-1}B \ \dots \ B]$ and $u = [u(0), \dots, u(T-1)]$. The optimal u is given by

$$u = -D^\dagger A^T x_0.$$

Plugging this into $h_c(x_0)$ results in

$$\begin{aligned} h_c(x_0) &= x_0^\top \left(A^{T^\top} D^{\dagger^\top} D^\dagger A^T \right) x_0 \\ &= \|D^\dagger A^T x_0\|_2^2. \end{aligned}$$

To compute the volume of all x_0 such that $h_c(x_0) \leq 1$, note that the volume in terms of a new variable $y = A^T x_0$ is the volume of an ellipse defined by $D^\dagger$:

$$\begin{aligned} \text{volume} \left\{ A^T x_0 \mid h_c(x_0) \leq 1 \right\} &= \text{volume} \left\{ y \mid \|D^\dagger y\|^2 \leq 1 \right\} \\ &= \sqrt{\det DD^\top}, \end{aligned}$$

and using the definition of D results in $Y_c^2 = \sum_{i=0}^{T-1} A^i B B^\top (A^\top)^i = V_c$. $\square$

7.4.2 Computing Constrained Controllability

Our second contribution is a novel algorithm (Algorithm 1) to calculate controllability with constraints in closed form. Algorithm 1 starts by eliminating the linear constraints in (7.6) to reformulate it as an unconstrained optimization problem. We first decompose $S_c \vec{\phi}_u + C_c = 0$ as n independent sets of linear constraints $S_i \Phi_i + C_i = 0$, for $i \in [N_x]$ in Step 1. Here, $\Phi_i \in \mathbb{R}^{TN_u}$ contains the entries in $\vec{\phi}_u$ that correspond to $\phi_u(t)[i, :]$ for $t < T$. Such decomposition can always be implemented due to the special structure of the support constraints induced by actuation delay, communication speed, and locality constraints. We state this result in Lemma 7.7.

Algorithm 1 Constrained Controllability

Require: Horizon T ; Constraint matrices S_c and C_c .

- 1: Decompose $S_c \vec{\phi}_u + C_c = 0$ as $S_i \Phi_i + C_i = 0$, $i \in [N_x]$. $\triangleright \Phi_i$ contains the entries in $\vec{\phi}_u$ that correspond to $x_0[i]$.
- 2: Compute the general solution $\Phi_i = \omega_i + V_i \lambda_i$, $i \in [N_x]$. $\triangleright \lambda_i \in \mathbb{R}^{r_i}$ and $V_i \in \mathbb{R}^{(TN_u) \times r_i}$ has full column rank.
- 3: Define $M = [V_1, V_2, \dots, V_n] \in \mathbb{R}^{(TN_u) \times (\sum_{i=1}^{N_x} r_i)}$.
- 4: Define $N = [\omega_1, \omega_2, \dots, \omega_n] \in \mathbb{R}^{(TN_u) \times N_x}$.
- 5: Return the volume

$$\det(A^T) / \sqrt{\det((-MM^\dagger N + N)^\top (-MM^\dagger N + N))}.$$

Lemma 7.7. *For any k , the constraint $(S_c \vec{\phi}_u + C_c)[k] = 0$ cannot depend on entries from both Φ_i and Φ_j ($i \neq j$).*

Proof. Under the definitions of constraints on actuation delay, communication speed, and locality in Section 7.3, the constraint $(S_c \vec{\phi}_u + C_c)[k] = 0$ corresponds to either $\vec{\phi}_u[\tau] = 0$ or $\vec{\phi}_x[\tau] = 0$ for some index $\tau \in [TN_u N_x]$. In the first case, it satisfies the statement of Lemma 7.7.

Otherwise, $(S_c \vec{\phi}_u + C_c)[k] = 0$ must correspond to $\vec{\phi}_x[\tau] = 0$ for some index $\tau \in [TN_u N_x]$. Suppose it is converted from $\phi_x(t)[j, i] = c$, where $c \in \{0, 1\}$. By (7.4), we see that $\phi_x(\tau + 1)[:, i] = A\phi_x(\tau)[:, i] + B\phi_u(\tau)[:, i]$ for all $\tau < t$. Therefore, $\phi_x(t)[j, i]$ can be expressed as a linear combination of $\{\phi_u(\tau)[:, i] \mid 0 \leq \tau < t\}$. Since all entries in $\phi_u(\tau)[:, i]$ are in Φ_i , $(S_c \vec{\phi}_u + C_c)[k] = 0$ only involves entries from Φ_i . $\square$

After the decomposition, Step 2 solves the general solution $\Phi_i = \omega_i + V_i \cdot \lambda_i$ of each constraint set $S_i \Phi_i + C_i = 0$, so we can change the variables from $\vec{\phi}_u$ to $\{\lambda_i\}_{i \in [N_x]}$ in the optimization problem (7.6). With the matrices M and N defined in Steps 3-4, we see $h_c(x_0)$ can be expressed as

$$\min_{\lambda_1, \dots, \lambda_{N_x}} \|Nx_0 + M(x_0[1]\lambda_1, \dots, x_0[N_x]\lambda_{N_x})^\top\|_2^2. \quad (7.8)$$

Finally, we conclude the discussion with a proposition that shows the output of Algorithm 1 is equal to the constrained controllability Y_c defined in (7.6).

Proposition 7.8. *The constrained controllability Y_c satisfies*

$$Y_c = \text{volume} \left\{ A^T x_0 : \|(-MM^\dagger N + N)x_0\|_2^2 \leq 1 \right\},$$

which is the output of Algorithm 1.

Proof. By (7.8), we see $h_c(x_0) = \min_y \|Nx_0 + My\|_2^2$ holds for all x_0 that satisfies $x_0[i] \neq 0$ for all $i \in [N_x]$, because we can do the change of variables

$$y := (x_0[1]\lambda_1, \dots, x_0[N_x]\lambda_{N_x})^\top \in \mathbb{R}^{\sum_{i=1}^{N_x} r_i}.$$

Since this equation holds almost everywhere, we obtain that

$$\begin{aligned} Y_c &= \text{volume} \left\{ A^T x_0 : \min_y \|Nx_0 + My\|_2^2 \leq 1 \right\} \\ &= \text{volume} \left\{ A^T x_0 : \|(-MM^\dagger N + N)x_0\|_2^2 \leq 1 \right\}, \end{aligned}$$

which is equal to the return value of Algorithm 1. $\square$

It is worth noting that the time complexity of Algorithm 1 can grow polynomially in time horizon T and system dimensions (N_x, N_u) . In Section 7.5.2, we show that the rank condition is faster than the volume computation, with the gap increasing as the state size increases.

7.4.3 Maintaining Performance Under Constraints

An important question for system design is whether the performance of the unconstrained system is preserved after imposing additional constraints on actuation delay, communication speed, and locality. While a straightforward answer is to compute the constrained controllability Y_c analytically with Algorithm 1, it can be inefficient for large-scale systems with high-dimensional state spaces. To address this challenge, we provide sufficient conditions under which

1. adding constraints does not decrease controllability;
2. constrained systems achieve the optimal unconstrained convex cost.

More importantly, they can be verified efficiently. Our result is inspired by a recent work [62] on comparing localized/global model predictive control (MPC), with the critical difference that we require $x_T = 0$ in (7.6). While [62] requires that one must check the conditions hold for each x_0 , we prove that if our conditions are satisfied, the performance is maintained for x_0 almost everywhere (a.e.).

First, we define sets $\mathcal{S}(x_0)$ and $\mathcal{T}(x_0)$ as the set of all feasible states and control inputs under the dynamics (7.1) and initial condition x_0 , without and with constraints. Let

$$Z_{AB0} := \left[\begin{array}{cc|cc} I & & 0 & \\ -A & I & -B & 0 \\ & \ddots & & \ddots \\ & -A & I & -B \\ & & -A & -B \end{array} \right],$$

such that $Z_{AB0} \in \mathbb{R}^{N_x T \times (N_x + N_u) T}$. Let $Z_p := Z_{AB0}^\dagger \begin{bmatrix} I \\ 0 \end{bmatrix}$, $Z_h := I - Z_{AB0}^\dagger Z_{AB0}$. Let $Z_h^{blk} = \text{blkdiag}(Z_h, \dots, Z_h)$, repeated N_x times. The indices R are defined as the indices of $\vec{\phi}$ where the entry is constrained to be zero. Denote the matrix $F = Z_h^{blk}[R]$; due to the structure of Z_h^{blk} , F has the form $F = \text{blkdiag}(\{Z_h[R_i]\}_{i \in [N_x]})$, where each R_i is the subset of indices in R that corresponds to the i^{th} block of Z_h^{blk} . Define $x_0 := [x_0[1]I \ x_0[2]I \ \dots \ x_0[N_x]I]$. We use the results of the following lemmas to prove the main results. Finally, let $H := [H_1, \dots, H_{N_x}]$ where $H_i = I - F[R_i]^\dagger F[R_i]$.

Lemma 7.9. *The dimension of $\mathcal{S}(x_0)$ is given by $\dim \mathcal{S}(x_0) = \text{rank}(Z_h)$.*

Lemma 7.10. *The dimension of $\mathcal{T}(x_0)$ is given by $\dim \mathcal{T}(x_0) = \text{rank}(Z_h x_0 (I - F^\dagger F))$.*

See the appendix for proofs, which follow closely from [62].

Theorem 7.11 (Optimality of Constrained Systems). *For any constraint set S_x, S_u , if there exists a feasible solution $\phi \in \mathcal{S}$ satisfying (7.4), $x_T = 0$, and $\text{rank}(Z_h) = \text{rank}(Z_h H)$, then $\mathcal{S}(x_0) = \mathcal{T}(x_0)$ for all x_0 almost everywhere.*

Proof. Note that F is a block-diagonal matrix; the product $Z_h x_0 (I - F^\dagger F)$ is therefore

$$\begin{aligned} Z_h x_0 (I - F^\dagger F) &= [Z_h x_0[1](I - Z_h[R_1]^\dagger Z_h[R_1]) \ \dots \\ &\quad Z_h x_0[N_x](I - Z_h[R_{N_x}]^\dagger Z_h[R_{N_x}])] . \end{aligned}$$

Let $a = \text{rank}(Z_h H)$. The matrix $Z_h H$ can be row-reduced into a matrix $\hat{H}$, where the first a rows of $\hat{H}$ are linearly independent and the last $N_x - a$ rows are linearly dependent on the first a rows. Select $\lambda \in \mathbb{R}^a$ such that

$$\sum_{j=1}^a \lambda_j [\hat{H}_1[j], \dots, \hat{H}_{N_x}[j]] = 0 .$$

Then because the first a rows of $\hat{H}$ are linearly independent, $\lambda_j = 0$ for $j \in [1, \dots, a]$. For any x_0 , define $\tilde{x}_0$ as x_0 re-weighted by the row-reduction used to define $\hat{H}$ from H . If $\tilde{x}_0$ satisfies $\tilde{x}_0[i] \neq 0 \ \forall i \in [N_x]$, then, for any $\tilde{\lambda} \in \mathbb{R}^a$ such that

$$\sum_{j=0}^a \tilde{\lambda}_j [\tilde{x}_0[1]\hat{H}_1[j], \dots, \tilde{x}_0[N_x]\hat{H}_{N_x}[j]] = 0 ,$$

we have that $\tilde{\lambda}_j \tilde{x}_0[i] = 0$ for all $i \in [N_x]$ and $j \in [a]$. Because $\tilde{x}_0[i] \neq 0$ for all $i \in [N_x]$, so $\tilde{\lambda}_j = 0$ for all $j \in [a]$. This results in the desired rank condition $\text{rank}(Z_h x_0 H) = \text{rank}(Z_h)$ for all

x_0 such that $\tilde{x}_0[i] \neq 0$ for all $i \in [N_x]$. Note that we do not analyze the linearly dependent rows of $\hat{H}$ because they will remain linearly dependent when rescaled by $\tilde{x}_0$. By Lemmas 7.9 and 7.10, $\dim \mathcal{S}(x_0) = \dim \mathcal{T}(x_0)$ for x_0 almost everywhere. Since $\mathcal{T}(x_0) \subseteq \mathcal{S}(x_0)$, the same dimensionality implies that the sets are equal. $\square$

Finally, we present a corollary of Theorem 7.11 that gives criteria for optimality conditions.

Corollary 7.12. *If there exists a feasible solution $\phi \in \mathcal{S}$ and $\text{rank}(Z_h) = \text{rank}(Z_h H)$, then for any convex cost function $f(x, u)$, the unconstrained cost is achieved under the constraints $\phi \in \mathcal{S}$ for x_0 almost everywhere, with $V_c = Y_c^2$.*

Proof. By Theorem 7.11, the theorem assumptions imply that $\mathcal{S}(x_0) = \mathcal{T}(x_0)$ for x_0 almost everywhere. Therefore any optimization problem over the constrained set obtains the same value as the unconstrained problem for x_0 almost everywhere.

Now select $f(x, u) = \|u\|_2^2$. The above result implies that the optimization $h_c(x_0)$ in (7.6) is over the same set with and without constraints, so the solution is the same for x_0 almost everywhere. Because the volume of an ellipse does not change over a measure-zero set, the volume over all x_0 is also the same. The result $V_c = Y_c^2$ then holds by Theorem 7.6. $\square$

From this result, we conclude that for commonly-used cost functions such as quadratic cost, the addition of constraints does not necessarily hurt the performance of the controller. This contrasts with some recent results on control [107, 70] and reinforcement learning [71, 140], where localized policies are less optimal than global policies. While this may seem surprising, it is worth noticing that we allow the controller to decide the (constrained) mapping $\vec{\phi}_u$ based on the initial state x_0 , which is easier than committing to a policy before x_0 is sampled or picked by an adversary. Further, the SLS formulation allows us to consider a more general policy class $u(t) = K(t)x(t)$ than the static policy $u(t) = Kx(t)$ in [107].

We show in the following numerical results how locality constraints satisfy this rank requirement and the controllability correspondingly does not change within this region. In contrast, actuation delay does not satisfy this condition and we observe a decrease in the controllability. In addition to explaining the contrasting behavior of different types of constraints, computing the rank condition is faster and simpler to implement compared with computing the volume. We show a comparison in computation times in the numerical results.

Remark 7.13. *While the formulation of $h_c(x_0)$ might suggest one needs global knowledge of x_0 to solve ϕ_u , one can show when locality constraints are enforced, the value of each entry of ϕ_u depends only on the local entries of x_0 . For example: let A be a tri-diagonal matrix (i.e., $A[i, j] = 0$ if*




Figure 7.1: When actuation delay is zero, the addition of locality constraints on the system does not change the controllability volume for locality greater than five. The controllability decreases with increasing actuation delay.

$|i - j| > 1$), $B = I$, $T = 1$ and the locality constraint include only the immediate neighbors. Then $\text{supp } \phi_u(t)$ is also tri-diagonal. Let ϕ_u denote $\phi_u(0)$ in this setting with only one time step. The objective $\|\phi_u x_0\|^2$ is

$$\begin{aligned} \|\phi_u x_0\|^2 &= (\phi_u[0, 0]x_0[0] + \phi_u[0, 1]x_0[1])^2 + \\ &\sum_{j=1}^{N_x-2} (\phi_u[j, j-1]x_0[j-1] + \phi_u[j, j]x_0[j] + \phi_u[j, j+1]x_0[j+1])^2 \\ &+ (\phi_u[N_x-1, N_x-2]x_0[N_x-2] + \phi_u[N_x-1, N_x-1]x_0[N_x-1])^2. \end{aligned}$$

Note that each $\phi_u[j]$ depends only on $x_0[j-1]$ to $x_0[j+1]$. As this objective is row-separable and the dynamical constraints are column separable, ϕ_u can be computed locally using the alternating direction method of multipliers (ADMM), where at each iteration, a node only needs information from its direct neighbor. See [7, Section 5.5] for details.

7.5 Numerical Results

7.5.1 Distributed Control of Linearized Swing Equations

To illustrate our results concretely, we consider the setting of the linearized swing equations, as used in [62], which models power generation at multiple substations connected together in a power grid. This example allows us to illustrate how controllability changes based on actuation delay and locality constraints. The model is given by

$$\begin{aligned} x(t+1)[i] &= \sum_{j: \mathcal{G}[i,j]=1} A_{ij}x(t)[j] + B_i u(t)[i] \\ A_{ii} &= \begin{bmatrix} 1 & \Delta t \\ -\frac{k_i}{m_i}\Delta t & 1 - \frac{d_i}{m_i}\Delta t \end{bmatrix}, \quad A_{ij} = \begin{bmatrix} 0 & 0 \\ \frac{k_{ij}}{m_i}\Delta t & 0 \end{bmatrix}, \quad B_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \end{aligned}$$




Figure 7.2: As the state dimension grows, the amount by which the rank condition computation time is faster than the controllability volume time increases.

The state $x(t)[i] = [\theta_i(t); \omega_i(t)]$ is the phase angle and frequency deviations at node i . The parameters k_{ij} , m_i^{-1} , and d_i are coupling, inertia, and damping parameters, with $k_i = \sum_{j: \mathcal{G}[i,j]=1} k_{ij}$. In the numerical implementation, the discretization step is $\Delta t = 0.1$, $T = 12$. There are nine two-state subsystems with one actuator per subsystem, resulting in $N_x = 18$ and $N_u = 9$. The subsystems are arranged in a 3×3 grid, with possible neighbors of a particular subsystem being the other subsystems to the left, right, top or bottom. The coefficients were sampled uniformly $k_{ij} \sim [0.5, 1]$, $m_i^{-1} \sim [0, 2]$, $d_i \sim [1, 1.5]$. In Figure 7.1, we observe that the settings where rank condition and feasibility hold coincides with the maximum controllability; importantly, we see no decrease in controllability with locality constraints until the system is no longer controllable. The controllability decreases with increasing actuation delay, and as expected, the rank condition does not hold for any non-zero actuation delay. Large delay renders the system not controllable.

7.5.2 Computation of Rank Condition Versus Volume

The rank condition, which rigorously explains why communication speed and locality constraints may not result in decreased controllability, is easier to implement and faster to compute for large systems. We recorded the total time required to compute the controllability volume for all possible combinations of communication and locality constraints for a system where the dynamics matrix A is given by a tri-diagonal stochastic matrix, B is filled with one and zero entries with actuation density 0.75 relative to the state dimension, the time horizon is set to $T = 10$, and 40 trials were recorded for each combination of constraints. As seen in Figure 7.2, the time is comparable for N_x ranging from 5 – 10, but the volume computation becomes increasingly slower with larger N_x , with up to $6\times$ speedup using the rank condition instead of the total volume. Users can compute the controllability volume for the unconstrained system and then compute the rank condition and feasibility of the locality- and communication-constrained system to determine the controllability, which will either be equal to the unconstrained system or possibly zero if the condition is not satisfied.

7.6 Concluding Remarks

This work proposes notions of constrained controllability Y_c and constrained observability Y_o using SLS, and shows they reduce to classic definitions for unconstrained systems. We provide closed-form solutions for computing Y_c, Y_o and an efficient criteria to check whether a certain set of constraints can cause performance loss. In future work, we aim to apply this framework to solve real-world sensor and actuator problems, including air and water quality sensor placement for environmental monitoring, and investigate methods for computing the controllability in a distributed fashion for large-scale systems, which would greatly improve the range of settings in which this method can be used.

Appendix

7.1 Proof of Lemma 7.9

To prove Lemma 7.9, we first show that the set $\mathcal{S}(x_0)$ can be written as

$$\mathcal{S}(x_0) = \{\mathbf{s} : \mathbf{s} = Z_p x_0 + Z_h x_0 \lambda, \lambda \in \mathbb{R}^{(N_u + N_x)N_x T}\},$$

where $\mathbf{s} = [\mathbf{x}; \mathbf{u}]$. Using the dynamics and terminal constraint $x_T = 0$, the closed-loop map ϕ can be written as

$$Z_{AB0}\phi = \begin{bmatrix} I \\ 0 \end{bmatrix}. \quad (7.9)$$

This results in $\phi = Z_{AB0}^\dagger \begin{bmatrix} I \\ 0 \end{bmatrix} + (I - Z_{AB0}^\dagger Z_{AB0})\Lambda$, where Λ is a free variable. Using the definitions of Z_p and Z_h results in $\phi = Z_p + Z_h \Lambda$. Then the state and input signals $\mathbf{x}$ and $\mathbf{u}$ result are given by $\mathbf{s} = \phi x_0$, and we use the matrix x_0 to switch the multiplication of Λx_0 to $x_0 \lambda$. The dimension of $\mathcal{S}(x_0)$ is given by the degrees of freedom of $\mathbf{s}$, which is $\text{rank}(Z_h x_0) = \text{rank}(Z_h)$ because x_0 is full row-rank.

7.2 Proof of Lemma 7.10

This proof proceeds similarly to that of Lemma 7.9, except now the constraint on ϕ is defined in terms of $\vec{\phi}$. From (7.9) and the additional requirements that $S_u \vec{\phi}_u = 0$ and $S_x \vec{\phi}_x = 0$,

$$\begin{bmatrix} S_x \\ S_u \end{bmatrix} (\hat{Z}_p + \hat{Z}_h \vec{\Lambda}) = 0 \Rightarrow \hat{Z}_p[R] + \hat{Z}_h[R] \vec{\Lambda},$$

resulting in Λ being parameterized as

$$\vec{\Lambda} = -\hat{Z}_h[R]^\dagger \hat{Z}_p[R] + (I - Z_h[R]^\dagger Z_h[R])\mu,$$

for a free variable μ . Using the definition of F and noting that $\hat{Z}_h[R] = Z_h^{blk}[R]$, the dimension of $\mathcal{T}(x_0)$ is given by $\text{rank}(Z_h x_0 (I - F^\dagger F))$.

*Chapter 8***LEARNING THE UNCERTAINTY SETS FOR CONTROL DYNAMICS VIA
SET MEMBERSHIP: A NON-ASYMPTOTIC ANALYSIS**

Y. Li, J. Yu, L. Conger, T. Kargin, and A. Wierman. Learning the Uncertainty Sets of Linear Control Systems via Set Membership: A Non-asymptotic Analysis. en. In *Proceedings of the 41st International Conference on Machine Learning*, pages 29234–29265. PMLR, July 2024

ABSTRACT

This chapter studies uncertainty set estimation for unknown linear systems. Uncertainty sets are crucial for the quality of robust control since they directly influence the conservativeness of the control design. Departing from the confidence region analysis of least squares estimation, this chapter focuses on set membership estimation (SME). Though good numerical performances have attracted applications of SME in the control literature, the non-asymptotic convergence rate of SME for linear systems remains an open question. This chapter provides the first convergence rate bounds for SME and discusses variations of SME under relaxed assumptions. We also provide numerical results demonstrating SME's practical promise.

8.1 Introduction

The problem of estimating unknown linear dynamical systems of the form $x_{t+1} = A^*x_t + B^*u_t + w_t$ with unknown parameters (A^*, B^*) has seen considerable progress recently [103, 23, 110, 127, 111, 27, 2]. Most literature focuses on the analysis of the least squares estimator (LSE) and its variants, where sharp bounds on the convergence rates for subGaussian disturbances w_t have been obtained [110, 111]. Building on this, there is a rapidly growing body of literature on “learning to control” unknown linear systems that leverages LSE to achieve various control objectives, such as stability and regret [21, 58, 110, 52, 81, 29].

However, for successful application of learning-based control methods to safety-critical applications, it is crucial to quantify the uncertainties of the estimated system and to robustly satisfy safety constraints and stability despite these uncertainties [126, 17]. A promising framework for achieving this is to estimate the uncertainty set of the unknown system parameters and to utilize robust controllers to satisfy the robust constraints under any parameters in the uncertainty set [17, 43]. Uncertainty set estimation is crucial for the success of robust control: on the one hand, too large of an uncertainty set gives rise to over-conservative control actions, resulting in degraded performance; on the other hand, if the uncertainty set is underestimated and fails to contain the true system, the resulting controller may lead to unsafe behaviors [17, 96].

To estimate uncertainty sets, a popular method is to construct LSE’s confidence regions [29, 110]. However, this approach yields a confidence region for a point estimate rather than directly estimating the uncertainty set of the model. Further, the confidence regions are usually derived from concentration inequalities, which allows convergence rate analysis but may suffer conservative constant factors [96, 110].

In this paper, we instead focus on a direct uncertainty set estimation method: set membership estimation (SME), which estimates the uncertainty set without relying on the concentration inequalities underlying the approaches based on LSE. SME has a long history in the control community [135, 59, 78, 75, 36, 14]. SME has primarily been proposed for scenarios with bounded disturbances, which is common in safety-critical systems, e.g. power systems [97], unmanned aerial vehicles (UAV) [13, 89], and building control [139]. Further, the bounded disturbance is a standard assumption in control when certain safety requirements are desired, such as robust (adaptive) constrained control [79, 78, 29], online (constrained) control [64, 65, 73, 66].

Consequently, SME has been widely adopted in the robust (adaptive) constrained control literature [78, 18, 137, 94, 93, 104] and the online control literature [45, 135, 133, 134]. Figure 8.1 provides a toy example illustrating SME’s promising performance under bounded disturbances.

On the theory side, the convergence analysis of SME generally considers a simple regression




Figure 8.1: A visualized toy example of uncertainty set comparison between SME in (8.2) and LSE confidence regions in [110, 2] for a one-dimensional system $x_{t+1} = A^*x_t + B^*u_t + w_t$, with $w_t, u_t \in [-1, 1]$ generated i.i.d. from a truncated Gaussian distribution. Detailed experiment settings are in Section 8.I. **Figure (a)** compares the diameters of the uncertainty sets from SME and LSE 90% confidence bounds. **Figure (b) and (c)** visualize the uncertainty sets after $T = 5$ and $T = 250$ data points.

problem: $y_t = \theta^*x_t + w_t$ with a deterministic sequence of x_t and bounded i.i.d. disturbances w_t [10, 5, 56, 11, 33]. This regression problem does not capture the correlation between x_t and the history $w_{t-1}, \dots, w_0$ in the dynamical systems. This issue was largely overlooked in the vast literature of empirical algorithm design related to SME (for example, see [78, 57], etc.). It is not until recently that [80] provide the first *asymptotic convergence* guarantees for SME in linear systems. However, the *non-asymptotic convergence rate* still remains open for SME in linear dynamical systems.

Contributions. This paper tackles the open question above by providing non-asymptotic bounds on the convergence rates of SME for linear systems. To the best of our knowledge, this is the first convergence rate analysis of SME for dynamical systems in the literature.

We consider two scenarios in our analysis. Firstly, when a *tight* bound $\mathcal{W}$ on the support of w_t

is known, we provide an instance-dependent convergence rate for SME. Interestingly, for several common distributions of w_t , SME enjoys a convergence rate $\tilde{O}(n_x^{1.5}(n_x + n_u)^2/T)$, which is faster than the LSE's error bound $O(\frac{\sqrt{n_x+n_u}}{\sqrt{T}})$ in terms of the number of samples T but is worse in terms of the dependence on state and control dimensions n_x, n_u . The improved convergence rate of SME with respect to T is enabled by leveraging the additional boundedness property of w_t , which is a common assumption in robust constrained control but is not utilized in LSE's analysis. Secondly, when a tight bound of w_t is *unknown*, we introduce a UCB-SME algorithm that learns conservative upper bounds of w_t from data and constructs uncertainty sets based on the conservative upper bounds. We also provide a convergence rate of UCB-SME, which has the same dependence on T but has worse dependence on n_x by a factor of $\sqrt{n_x}$ compared with the convergence rate with a known tight bound.

Our estimation error bound relies on a novel construction of an event sequence based on designing a sequence of stopping times. This construction, together with the BMSB condition in [111], addresses the challenge caused by the correlation between x_t, u_t , and the history disturbances (see the proof of Theorem 8.8 for more details).

Moreover, our results lay a foundation for future non-asymptotic analysis of control designs based on SME. To illustrate this, we apply our results to robust-adaptive model predictive control and robust system-level-synthesis (SLS) and discuss the novel non-asymptotic guarantees enabled by our convergence rates of SME.

Finally, we conduct extensive simulations to compare the numerical behaviors of SME, UCB-SME, and LSE's confidence regions, which demonstrates the promising performance of SME and UCB-SME.

8.2 Problem Formulation and Preliminaries

8.2.1 Problem Formulation

This paper focuses on the identification of uncertainty sets of unknown system parameters in the linear dynamical system:

$$x_{t+1} = A^* x_t + B^* u_t + w_t \quad (8.1)$$

where A^*, B^* are the unknown system parameters, $x_t \in \mathbb{R}^{n_x}, u_t \in \mathbb{R}^{n_u}$. For notational simplicity, we define $\theta^* = (A^*, B^*)$ by matrix concatenation and $z_t = (x_t^\top, u_t^\top)^\top \in \mathbb{R}^{n_z}$ by vector concatenation, where $n_z = n_x + n_u$. Accordingly, the system (8.1) can be written as $x_{t+1} = \theta^* z_t + w_t$.

The goal of the uncertainty set identification problem is to determine a set Θ_T that contains the true parameters $\theta^* = (A^*, B^*)$ based on a sequence of data $\{x_t, u_t, x_{t+1}\}_{t=0}^{T-1}$. Set Θ_T is called an uncertainty set since it captures the remaining uncertainty on the system model after the revelation of the data sequence $\{x_t, u_t, x_{t+1}\}_{t=0}^{T-1}$.

Uncertainty sets play an important role in robust control, where one aims to achieve robust constraint satisfaction [78, 79], robust objective optimization [131], and/or robust stability [119] for any model in the uncertainty set.¹ Therefore, the diameter of the uncertainty sets heavily influences the conservativeness of robust controllers and thus the control performance. Formally, we define the diameter as follows.

Definition 8.1 (Diameter of a set of matrices). *Consider a set $\mathbb{S}$ of matrices $\theta \in \mathbb{R}^{n_x \times n_z}$. We define the diameter of $\mathbb{S}$ in Frobenius norm as $\text{diam}(\mathbb{S}) = \sup_{\theta, \theta' \in \mathbb{S}} \|\theta - \theta'\|_F$.*

8.2.2 Set Membership Estimation (SME)

In this section, we review set membership estimation (SME), which is an uncertainty set identification method that has been studied in the control literature for decades [79, 14]. SME primarily focuses on systems with *bounded* disturbances, i.e. $w_t \in \mathcal{W}$ for some bounded $\mathcal{W}$ for all $t \geq 0$. When $\mathcal{W}$ is known, SME computes an uncertainty/membership set by

$$\Theta_T = \bigcap_{t=0}^{T-1} \{\hat{\theta} : x_{t+1} - \hat{\theta} z_t \in \mathcal{W}\}. \quad (8.2)$$

It is straightforward to see that $\theta^* \in \Theta_T$ when $w_t \in \mathcal{W}$.

The bounded disturbance assumption may seem restrictive, considering that the uncertainty set identification based on the confidence region of LSE only requires subGaussian disturbances [111]. However, in many control applications, it is reasonable and common to assume bounded w_t . For example, bounded disturbances is a standard assumption in the robust constrained control literature, such as robust constrained LQR [79, 78, 80, 29], and online constrained control of linear systems [73, 64]. This is different from unconstrained control, where unbounded subGaussian disturbances are usually considered [119]. The difference in the disturbance formulation is largely motivated by the applications: constrained control is mostly applied to safety-critical applications, where the disturbances are usually bounded. For example, in UAV and flight control, the disturbances are mostly caused by wind gusts, and wind disturbances are bounded in practice [13, 89]. Similarly, in building thermal control, the disturbances are caused by external heat exchanges, which are also bounded [139].

Ideally, one hopes that Θ_T converges to the singleton of the true model $\{\theta^*\}$ or at least a small neighborhood of θ^* . This usually calls for additional assumptions, such as the persistent excitation property on the observed data and additional stochastic properties on w_t . In this paper, we consider the following assumptions to establish convergence rate bounds on the diameter of Θ_T , which, to the best of our knowledge, is the first non-asymptotic guarantee of SME for linear dynamical systems.

¹In addition to model uncertainties, robust control may also consider other system uncertainties, e.g., disturbances, measurement noises, etc.

The first assumption formalizes the bounded disturbance assumption discussed above and introduces stochastic properties of w_t for analytical purposes.

Assumption 8.2 (Bounded i.i.d. disturbances). *The disturbances are box-constrained, $w_t \in \mathcal{W} := \{w \in \mathbb{R}^{n_x} : \|w\|_\infty \leq w_{\max}\}$ for all $t \geq 0$. Further, w_t is i.i.d., has zero mean and positive definite covariance matrix Σ_w .*

Assumption 8.2 is common in SME literature, e.g. [5, 80, 33]. In terms of generality, boundedness is essential for SME. The stochastic properties, such as i.i.d., zero mean, positive definite covariance, are standard in the recent learning-based control literature and allow the use of statistical tools utilized and developed in the recent literature for non-asymptotic analysis [111, 67]. Besides, it is worth mentioning that SME still works in non-stochastic settings. In particular, as long as $w_t \in \mathcal{W}$, even without the stochastic properties in Assumption 8.2, the SME algorithm (8.2) still generates a valid uncertainty set that contains θ^* . It is an interesting future direction to study the convergence rate of SME without assuming stochastic disturbances.

Next, we introduce the assumptions on u_t , which relies on the block-martingale small-ball (BMSB) condition proposed in [111]. It can be shown that the BMSB guarantees persistent excitation (PE) with high probability under proper conditions (see Proposition 2.5 in [111] and Lemma 8.16). The PE condition requires that z_t explores all directions, which is essential for system identification [90].

Definition 8.3 (Persistent excitation). *There exists $\alpha > 0$ and $m \in \mathcal{N}+$, such that for any $t_0 \geq 0$,*

$$\frac{1}{m} \sum_{t=t_0}^{t_0+m-1} \begin{pmatrix} x_t \\ u_t \end{pmatrix} (x_t^\top, u_t^\top) \succeq \alpha^2 I_{n_x+n_u}.$$

Definition 8.4 (BMSB [111]). *Consider a filtration $\{\mathcal{F}_t\}_{t \geq 1}$ and an $\{\mathcal{F}_t\}_{t \geq 1}$ -adapted random process $\{Z_t\}_{t \geq 1}$ in $\mathbb{R}^d$. $\{Z_t\}_{t \geq 1}$ satisfies the (k, Γ_{sb}, p) -block martingale small-ball (BMSB) condition for $k > 0$, a positive definite Γ_{sb} , and $0 \leq p \leq 1$, if the following holds: for any fixed $\lambda \in \mathbb{R}^d$ with $\|\lambda\|_2 = 1$, we have $\frac{1}{k} \sum_{i=1}^k \mathbb{P}(|\lambda^\top Z_{t+i}| \geq \sqrt{\lambda^\top \Gamma_{sb} \lambda} \mid \mathcal{F}_t) \geq p$ for all $t \geq 1$.*

The following is the assumption on u_t .

Assumption 8.5 (BMSB and boundedness). *With filtration $\mathcal{F}_t = \mathcal{F}(w_0, \dots, w_{t-1}, z_0, \dots, z_t)$, the $\mathcal{F}_t$ -adapted stochastic process $\{z_t\}_{t \geq 0}$ satisfies $(1, \sigma_z^2 I_{n_z}, p_z)$ -BMSB for some $\sigma_z, p_z > 0$. Besides, there exists $b_z \geq 0$ such that $\|z_t\|_2 \leq b_z$ almost surely for all $t \geq 0$.*

Assumption 8.5 requires u_t to guarantee both BMSB and bounded z_t . This can be satisfied by several robust (adaptive) constrained control policies, such as robust (adaptive) model predictive

control (MPC) [79, 78, 80], system level synthesis (SLS) [29], and control barrier functions (CBF) [77]. In the following, we briefly discuss robust (adaptive) MPC as an example. SLS and CBF can be similarly shown to satisfy Assumption 8.5.

Example 8.6 (Robust (adaptive) MPC). *Robust MPC is a popular method for the robust constrained control [99], which aims to optimize the control objective while satisfying robust safety constraints,*

$$z_t \in \mathbb{Z}_{\text{safe}}, \text{ where } x_{t+1} = \theta z_t + w_t, \forall \theta \in \Theta_0, w_t \in \mathcal{W}, \quad (8.3)$$

where Θ_0 is an initial uncertainty set known a priori, and the safety constraint $\mathbb{Z}_{\text{safe}}$ is usually bounded. The robust MPC policy, denoted by $u_t = \pi_{\text{RMPC}}(x_t; \Theta_0, \mathcal{W})$, satisfies the constraints (8.3) for any $\theta \in \Theta_0$. Therefore, it naturally guarantees bounded z_t under the true θ^* . Further, as shown in [69], BMSB can be achieved by adding a random disturbance, i.e. $u_t = \pi_{\text{RMPC}}(x_t; \Theta_0, \mathcal{W}) + \eta_t$, where η_t is i.i.d., bounded, and has positive definite covariance. Therefore, the randomly perturbed robust MPC can satisfy Assumption 8.5. Robust adaptive MPC is based on the same control design, $u_t = \pi_{\text{RMPC}}(x_t; \Theta_t, \mathcal{W})$, but utilizes adaptively updated uncertainty sets Θ_t . Notice that Θ_t is usually updated by SME in the literature of robust adaptive MPC [78, 79, 57].

We also note that BMSB and bounded z_t with high probability are assumed in LSE literature (Theorem 2.4 [111]), and bounded z_t with high probability under subGaussian disturbances corresponds to bounded z_t under bounded disturbances for linear systems (see bounded-input-bounded-output stability in Sec. 9 of [42]).

Finally, we assume that the bound $w_{\max}$ on w_t is tight in all directions, which is common in the literature on SME analysis [10, 5, 80].

Assumption 8.7 (Tight bound on w_t). *For any $\epsilon > 0$, there exists $q_w(\epsilon) > 0$, such that for any $1 \leq j \leq n$, we have*

$$\min(\mathbb{P}(w_t^j \leq \epsilon - w_{\max}), \mathbb{P}(w_t^j \geq w_{\max} - \epsilon)) \geq q_w(\epsilon),$$

where w_t^j denotes the j th entry of vector w_t . Without loss of generality, we can further assume $q_w(\epsilon)$ to be non-decreasing with ϵ and $q_w(2w_{\max}) = 1$.²

In essence, Assumption 8.7 requires that a hyper-cubic $\mathcal{W} = \{w : \|w\|_{\infty} \leq w_{\max}\}$ should be *tight* on the support of w_t in all coordinate directions, that is, there exists a positive probability $q_w(\epsilon)$ such that w_t visits an ϵ -neighborhood of $w_{\max}$ and $-w_{\max}$, respectively, on all coordinates.

²This is because $\mathbb{P}(w_t^j \leq \epsilon - w_{\max})$ and $\mathbb{P}(w_t^j \geq w_{\max} - \epsilon)$ are non-decreasing with ϵ , and $\mathbb{P}(w_t^j \geq -w_{\max}) = \mathbb{P}(w_t^j \leq w_{\max}) = 1$ by Assumption 8.2.

When the support of w_t is indeed $\mathcal{W} = \{w : \|w\|_\infty \leq w_{\max}\}$, many common distributions enjoy $q_w(\epsilon) \geq \Omega(\epsilon)$.³ For example, for the uniform distribution on $\mathcal{W}$, we have $q_w(\epsilon) = \frac{\epsilon}{2w_{\max}}$; for the truncated Gaussian distribution with zero mean, $\sigma_w^2 I_n$ covariance, and truncated region $\mathcal{W}$, we have $q_w(\epsilon) = \frac{\epsilon}{2w_{\max}\sigma_w} \exp(\frac{-w_{\max}^2}{2\sigma_w^2})$; and for the uniform distribution on the boundary of $\mathcal{W}$ (a generalization of Rademacher distribution), we have $q_w(\epsilon) \geq \frac{1}{2n_x} \geq \Omega(\epsilon)$ (see Section 8.C.2 for more details).

However, knowing a tight bound on the support of w_t can be challenging in practice. Therefore, we will discuss how to relax this assumption and learn a tight bound from data in Section 8.3.2.

Further, the requirement of a hyper-cubic $\mathcal{W}$ can be restrictive because different entries of disturbances may have different magnitudes, resulting in a hyper-rectangular support that violates Assumption 8.7. Our follow-up work [132] relaxes this assumption and generalizes the results in this paper.

8.3 Set Membership Convergence Analysis

8.3.1 Convergence Rate of SME with Known $w_{\max}$

We now present the main result (Theorem 8.8) of this paper, which is a non-asymptotic bound on the estimation error of SME given bounded i.i.d. stochastic disturbances.

Theorem 8.8 (Convergence rate of SME). *For any $m > 0$ any $\delta > 0$, when $T > m$, we have*

$$\begin{aligned} \mathbb{P}(\text{diam}(\Theta_T) > \delta) &\leq \underbrace{\frac{T}{m} \tilde{O}(n_z^{2.5}) a_2^{n_z} \exp(-a_3 m)}_{\mathcal{T}_1} \\ &\quad + \underbrace{\tilde{O}((n_x n_z)^{2.5}) a_4^{n_x n_z} \left(1 - q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{\lceil T/m \rceil}}_{\mathcal{T}_2(\delta)} \end{aligned}$$

where $a_1 = \frac{\sigma_z p_z}{4}$, $a_2 = \frac{64b_z^2}{\sigma_z^2 p_z^2}$, $a_3 = \frac{p_z^2}{8}$, $a_4 = \frac{4b_z \sqrt{n_x}}{a_1}$, p_z, σ_z, b_z are defined in Assumption 8.5, $\lceil \cdot \rceil$ denotes the ceiling function, and $\text{diam}(\cdot)$ is defined in Definition 8.1, the factors hidden in $\tilde{O}(\cdot)$ are provided in Appendix 8.D.4.

Theorem 8.8 provides an upper bound on the “failure” probability of SME, i.e., the probability that the diameter of the uncertainty set is larger than δ . In this bound, $\mathcal{T}_1$ decays exponentially with m , so for any small $\epsilon > 0$, m can be chosen such that $\mathcal{T}_1 \leq \epsilon$, which indicates $m \geq O(n_z + \log T + \log(1/\epsilon))$. For any $\delta > 0$, $\mathcal{T}_2(\delta)$ decays exponentially with the number of data points T and involves a distribution-dependent function $q_w(\cdot)$, which characterizes how likely it is for w_t to visit the boundary of $\mathcal{W}$ as

³The $\Omega(\cdot)$ notation is the lower bound version of $O(\cdot)$.

defined in Assumption 8.7. To ensure the probability upper bound in Theorem 8.8 to be less than 1, one can choose $m = O(\log T)$ and a large enough T such that $T \geq O(m) = O(\log(T))$. If w_t is more likely to visit the boundary, (a larger $q_w(\cdot)$), then SME is less likely to generate an uncertainty set with a diameter bigger than δ .

Estimation error bounds when $q_w(\epsilon) = \Omega(\epsilon)$. To provide intuition for $\mathcal{T}_2(\delta)$ and discuss the estimation error bound in Theorem 8.8 more explicitly, we consider distributions satisfying $q_w(\epsilon) = \Omega(\epsilon)$ for all $\epsilon > 0$. Notice that several common distributions satisfy this additional requirement, such as uniform distribution and truncated Gaussian distribution as discussed after Assumption 8.7.

Corollary 8.9 (Estimation error bound when $q_w(\epsilon) = \Omega(\epsilon)$). *For any $\epsilon > 0$, let*

$$m \geq O(n_z + \log T + \log(1/\epsilon))$$

in the following.⁴ If w_t is generated i.i.d. by a distribution satisfying $q_w(\epsilon) = \Omega(\epsilon)$ for all $\epsilon > 0$, then with probability at least $1 - 2\epsilon$, for any $\hat{\theta}_T \in \Theta_T$, we have

$$\|\hat{\theta}_T - \theta^*\|_F \leq \text{diam}(\Theta_T) \leq \tilde{O}\left(\frac{n_x^{1.5}(n_x + n_u)^2}{T}\right).$$

Corollary 8.9 indicates that the estimation error of any point in the uncertainty set Θ_T can be bounded by $\tilde{O}\left(\frac{n_x^{1.5}(n_x + n_u)^2}{T}\right)$ when $q_w(\epsilon) \geq \Omega(\epsilon)$.

Dynamical systems without control inputs. SME also applies to dynamical systems with no control inputs, i.e., $x_{t+1} = A^*x_t + w_t$, where the uncertainty set of A^* can be computed by $\mathbb{A}_T = \bigcap_{t=0}^{T-1} \{\hat{A} : \|x_{t+1} - \hat{A}x_t\|_\infty \leq w_{\max}\}$. Its convergence rate can be similarly derived via the proof of Theorem 8.8.

Corollary 8.10 (Convergence rate with $B^* = 0$ (informal)). *For stable A^* , for any $m > 0, \delta > 0, T > m$, we have*

$$\begin{aligned} \mathbb{P}(\text{diam}(\mathbb{A}_T) > \delta) &\leq \frac{T}{m} \tilde{O}(n_x^{2.5}) a_2^{n_x} \exp(-a_3 m) \\ &\quad + \tilde{O}(n_x^5) a_4^{n_x^2} (1 - q_w(\frac{a_1 \delta}{4\sqrt{n_x}}))^{\lceil T/m \rceil} \end{aligned}$$

Consequently, when $q_w(\epsilon) = \Omega(\epsilon)$, e.g. uniform or truncated Gaussian, we have $\text{diam}(\mathbb{A}_T) \leq \tilde{O}(n_x^{3.5}/T)$.

⁴A detailed formula is provided in Appendix 8.E.

Note that [111] have shown a lower bound $\Omega(\sqrt{n_x}/\sqrt{T})$ for the estimation of linear systems with no control inputs when w_t follows an (unbounded) Gaussian distribution. Interestingly, Corollary 8.10 reveals that, for some bounded-support distributions of w_t , e.g. Uniform and truncated Gaussian, SME is able to converge at a faster rate $\tilde{O}(1/T)$ in terms of the sample size T . This does not conflict with the lower bound in [111] because SME's rate only holds for bounded disturbances. In fact, from (8.2), it is straightforward to see that SME does not even converge under Gaussian disturbances. Therefore, SME is mostly useful in applications with bounded disturbances, e.g. robust constrained control, safety-critical systems, etc., while LSE's confidence regions are preferred for unbounded disturbances.

Lastly, Corollary 8.10 shows that SME's convergence rate has a poor dependence with respect to n_x : $\tilde{O}(n_x^{3.5})$. This is likely a proof artifact because we do not observe such poor dimension scaling in simulation (see Figure 8.3). It is left as future work to refine the dimension dependence.

8.3.2 SME with Unknown $w_{\max}$

Next, we discuss the convergence rates of SME without knowing a tight bound $w_{\max}$ in three steps: 1) only knowing a conservative upper bound of $w_{\max}$, 2) learning $w_{\max}$ from data, and 3) a variant of SME that converges without prior knowledge of $w_{\max}$.

1) SME with a conservative upper bound for $w_{\max}$. In many practical scenarios, it is easier to obtain an over-estimation of the range of the disturbances instead of a tight upper bound, i.e., $\hat{w}_{\max} \geq w_{\max}$. In this case, we can show that the uncertainty set converges to a small neighborhood around θ^* of size $O(\sqrt{n_x}(\hat{w}_{\max} - w_{\max}))$ at the same convergence rate as Theorem 8.8.

Theorem 8.11 (Conservative bound on $w_{\max}$). *When $w_{\max}$ in Assumption 8.7 is unknown but an upper bound $\hat{w}_{\max} \geq w_{\max}$ is known, consider the following SME algorithm:*

$$\hat{\Theta}_T(\hat{w}_{\max}) = \bigcap_{t=0}^{T-1} \{\hat{\theta} : \|x_{t+1} - \hat{\theta}z_t\|_{\infty} \leq \hat{w}_{\max}\},$$

For any $m > 0$, $\delta > 0$, $T > m$, we have

$$\mathbb{P}(\text{diam}(\hat{\Theta}_T) > \delta + a_5 \sqrt{n_x}(\hat{w}_{\max} - w_{\max})) \leq \mathcal{T}_1 + \mathcal{T}_2(\delta)$$

where $a_5 = \frac{4}{a_1}$, $\mathcal{T}_1, \mathcal{T}_2(\delta)$ are defined in Theorem 8.8.

2) Learning $w_{\max}$. When $w_{\max}$ is not accurately known, we can try to learn it from the data. Let's first consider the learning algorithm studied in [10].

$$\bar{w}_{\max}^{(T)} = \min_{\theta} \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta z_t\|_{\infty}. \quad (8.4)$$

Though algorithm (8.4) cannot provide an upper bound on $w_{\max}$ under finite samples because $\bar{w}_{\max}^{(T)} \leq w_{\max}$ for finite T ,⁵ it can be shown that $\bar{w}_{\max}^{(T)}$ converges to $w_{\max}$ as $T \rightarrow +\infty$. The convergence for linear regression has been established in [10]. The following theorem establishes the convergence and convergence rate of algorithm (8.4) for linear dynamical systems. Based on this convergence rate, we will design an online learning algorithm (8.5) that generates converging upper bounds of $w_{\max}$.

Theorem 8.12. *The estimation $\bar{w}_{\max}^{(T)}$ of $w_{\max}$ satisfies:*

$$0 \leq w_{\max} - \bar{w}_{\max}^{(T)} \leq \underbrace{b_z \text{diam}(\Theta_T)}_{\mathcal{T}_3} + \underbrace{w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty}}_{\mathcal{T}_4}$$

Therefore, for any $\delta > 0$,

$$\mathbb{P}(w_{\max} - \bar{w}_{\max}^{(T)} > \delta) \leq \mathcal{T}_1 + \mathcal{T}_2 \left(\frac{\delta}{2b_z} \right) + \mathcal{T}_5 \left(\frac{\delta}{2} \right),$$

where $\mathcal{T}_5(\delta) = (1 - q_w(\delta))^T$.

Notice that $\mathcal{T}_4$ is the smallest possible learning error of $w_{\max}$ from history w_t , which can be achieved if one can directly measure w_t . However, with unknown θ^* , it is challenging to measure/compute w_t exactly, then Theorem 8.12 shows that the learning error of $w_{\max}$ has an additional term $\mathcal{T}_3$ that depends on the uncertainty around θ^* . Therefore, the convergence rate of $\bar{w}_{\max}^{(T)}$ can be obtained by our non-asymptotic analysis of SME in Theorem 8.8.

Further, when $q_w(\epsilon) = \Omega(\epsilon)$, the convergence rate of $\bar{w}_{\max}^{(T)}$ can be explicitly bounded by $\tilde{O}(n_x^{1.5} n_z^2 / T)$, which is of the same order as the convergence rate of the diameter of Θ_T .

Corollary 8.13. *For any $0 < \epsilon < 1/3$ and any $T \geq 1$, there exists $\delta_T > 0$ satisfying $\lim_{T \rightarrow \infty} \delta_T = 0$ such that*

$$0 \leq w_{\max} - \bar{w}_{\max}^{(T)} \leq \delta_T$$

with probability at least $1 - 3\epsilon$.

In particular, when $q_w(\delta) = O(\delta)$, with probability $1 - 3\epsilon$,

$$0 \leq w_{\max} - \bar{w}_{\max}^{(T)} \leq \delta_T = \tilde{O}(n_x^{1.5} n_z^2 / T)$$

3) SME with unknown $w_{\max}$. Unfortunately, $\bar{w}_{\max}^{(T)}$ cannot be directly applied to SME because $\bar{w}_{\max}^{(T)} \leq w_{\max}$, which may cause $\theta^* \notin \hat{\Theta}_T(\bar{w}_{\max}^{(T)})$. However, by leveraging our convergence rate bound

⁵If SME does not use an upper bound on $w_{\max}$, the generated uncertainty set may not contain the true parameter θ^* .

in Theorem 8.12, we can construct an upper confidence bound (UCB) of $w_{\max}$ and a corresponding UCB-SME algorithm:

$$\hat{w}_{\max}^{(T)} = \bar{w}_{\max}^{(T)} + \delta_T, \quad \hat{\Theta}_T^{\text{ucb}} = \hat{\Theta}_T(\hat{w}_{\max}^{(T)}), \quad (8.5)$$

where δ_T is defined in Corollary 8.13.

Then, by combining Theorem 8.11 and Corollary 8.13, we can verify the well-definedness of UCB-SME and obtain its convergence rate.

Theorem 8.14. *For any $0 < \epsilon < 1/3$, any $T \geq 1$, with probability at least $1 - 3\epsilon$, we have*

$$\theta^* \in \hat{\Theta}_T^{\text{ucb}}, \quad \text{diam}(\hat{\Theta}_T^{\text{ucb}}) \leq O(\sqrt{n_x} \delta_T).$$

In particular, if $q_w(\epsilon) = \Omega(\epsilon)$, then $\text{diam}(\hat{\Theta}_T^{\text{ucb}}) \leq O(n_x^2 n_z^2 / T)$ with probability at least $1 - 3\epsilon$.

Notice that UCB-SME converges at the same rate in terms of T but $\sqrt{n_x}$ -worse in terms of dimensionality when compared with SME knowing a tight bound $w_{\max}$.

Remark 8.15 (Computation complexity). *SME can be computed by linear programming since all constraints are linear in (8.2). Further, UCB-SME can also be computed by linear programming because (8.4) can be reformulated as a linear program. However, the number of constraints for SME and UCB-SME increases linearly with T . To address the computation issue of SME, many computationally efficient algorithms have been proposed based on approximations of (8.2), e.g. [80, 133, 11]. The convergence rates of these approximate algorithms are unknown and how to design computationally efficient UCB-SME remains open.*

8.4 Proof Sketch of Theorem 8.8

The major technical novelty of this paper is the proof of Theorem 8.8, thus we describe the key ideas here. The complete proof is provided in Appendix 8.D. For ease of notation and without loss of generality, we assume T/m is an integer in the following.

Specifically, we first define a set Γ_T on the model estimation error $\gamma = \hat{\theta} - \theta^*$ by leveraging the observation that $x_{s+1} - \hat{\theta} z_s = w_s - (\hat{\theta} - \theta^*) z_s$,

$$\Gamma_t = \bigcap_{s=0}^{t-1} \{\gamma : \|w_s - \gamma z_s\|_{\infty} \leq w_{\max}\}, \quad \forall t \geq 0. \quad (8.6)$$

Notice that $\Theta_t = \theta_* + \Gamma_t$, so $\text{diam}(\Theta_t) = \text{diam}(\Gamma_t)$, and

$$\text{diam}(\Gamma_t) = \sup_{\gamma, \gamma' \in \Gamma_t} \|\gamma - \gamma'\|_F \leq 2 \sup_{\gamma \in \Gamma_t} \|\gamma\|_F.$$




Figure 8.2: **Figures (a)-(b)** compares the diameters of SME, UCB-SME, and SME with loose disturbance upper bounds that are 2, 3, 5, and 10 times larger than the true disturbance bound $w_{\max}$, as well as the baseline uncertainty set from the 90% confidence region of LSE. **Figure (c)** shows the convergence to the true bound $w_{\max}$ of the lower estimation $\hat{w}_{\max}$ in (8.4) and the UCB $\hat{w}_{\max}$ generated by the UCB-SME algorithm in **Figures (a)-(b)**.

Thus, we can define $\mathcal{E}_1 := \{\exists \gamma \in \Gamma_T, \text{ s.t. } \|\gamma\|_F \geq \frac{\delta}{2}\}$ such that $\mathbb{P}(\text{diam}(\Theta_T) > \delta) \leq \mathbb{P}(\mathcal{E}_1)$. Next, we define an event $\mathcal{E}_2$ below, which is essentially PE on every time segments $km + 1 \leq t \leq km + m$ for $k \geq 0$, where the choice of m will be specified later.

$$\mathcal{E}_2 = \left\{ \frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z}, \forall 0 \leq k \leq \left\lfloor \frac{T}{m} \right\rfloor - 1 \right\}$$

where $a_1 = \frac{\sigma_z p_z}{4}$. Now, by dividing the event $\mathcal{E}_1$ based on $\mathcal{E}_2$, we obtain

$$\mathbb{P}(\text{diam}(\Theta_T) > \delta) \leq \mathbb{P}(\mathcal{E}_1) \leq \mathbb{P}(\mathcal{E}_2^c) + \mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2).$$

The proof can be completed by establishing the following bounds on $\mathbb{P}(\mathcal{E}_2^c)$ and $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2)$.

Lemma 8.16 (Bound on $\mathbb{P}(\mathcal{E}_2^c)$). $\mathbb{P}(\mathcal{E}_2^c) \leq \mathcal{T}_1$, where $a_2 = \frac{64b_z^2}{\sigma_z^2 p_z^2}$ and $a_3 = \frac{p_z^2}{8}$.

Lemma 8.17 (Bound on $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2)$). $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \leq \mathcal{T}_2(\delta)$, where $a_4 = \max(1, 4b_z\sqrt{n_x}/a_1)$.

Roughly, Lemma 8.16 indicates that PE holds with high probability, which is proved by leveraging the BMSB assumption and set discretization. The proof of Lemma 8.17 is more involved and is our major technical contribution. On a high level, the proof relies on two technical lemmas below.

Lemma 8.18 (Discretization of $\mathcal{E}_1 \cap \mathcal{E}_2$ (informal)). *Let $\mathcal{M} = \{\gamma_1, \dots, \gamma_{v_\gamma}\}$ denote an ϵ_γ -net of $\{\gamma : \|\gamma\|_F = 1\}$. Under a proper choice of ϵ_γ , we have $v_\gamma = \tilde{O}(n_x^{2.5} n_z^{2.5}) a_4^{n_x n_z}$.⁶ We can construct $\tilde{\Gamma}_T$ such that*

$$\begin{aligned} \mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) &\leq \mathbb{P}(\{\exists 1 \leq i \leq v_\gamma, d \geq 0, \text{ s.t. } d\gamma_i \in \tilde{\Gamma}_T\} \cap \mathcal{E}_2) \\ &\leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2) \end{aligned}$$

where $\mathcal{E}_{1,i} = \{\exists d \geq 0, \text{ s.t. } d\gamma_i \in \tilde{\Gamma}_T\}$.

Lemma 8.18 leverages finite set discretization to bound the existence of a feasible element in an infinite continuous set. The formal version of Lemma 8.18 is provided as Lemma 8.33 in the appendix.

Lemma 8.19 (Construction of event $G_{i,k}$ via stopping times (informal)). *Consider $\mathcal{F}_t$ as defined in Assumption 8.5. Under the conditions in Lemma 8.18, we construct $G_{i,k}$ for all i and all $0 \leq k \leq T/m - 1$ by*

$$G_{i,k} = \left\{ b_{i,km+L_{i,k}} w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} \geq \frac{a_1 \delta}{4\sqrt{n_x}} - w_{\max}, \text{ and } \frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z} \right\}.$$

where $b_{i,t}, j_{i,t}$ are measurable in $\mathcal{F}_t$, and $L_{i,k}$ is constructed as a stopping time with respect to $\{\mathcal{F}_{km+l}\}_{l \geq 0}$. The formal definitions of $b_{i,t}, j_{i,t}, L_{i,k}$ are provided in Appendix 8.D.3.1.

Then, we have

$$\mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2) \leq \mathbb{P}\left(\bigcap_{k=0}^{T/m-1} G_{i,k}\right) \leq \left(1 - q_w \left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{\frac{T}{m}}$$

⁶The exact formulas of v_γ and ϵ_γ are in Lemma 8.28.

The constructions of $G_{i,k}$ and $L_{i,k}$ in Lemma 8.19 are our major technical contribution. With the constructions above, the proof can be completed by leveraging the conditional independence property of stopping times, which is briefly discussed below. Notice that by conditioning on the event $\{L_{i,k} = l\}$, we have $w_{km+L_{i,k}} = w_{km+l}$ and w_{km+l} is independent of $\mathcal{F}_{km+l}$. Consequently, w_{km+l} is also independent of $b_{i,km+L_{i,k}}, j_{i,km+L_{i,k}}$ conditioning on $\{L_{i,k} = l\}$ since $b_{i,km+l}, j_{i,km+l}$ are measurable in $\mathcal{F}_{km+l}$. Therefore, the probability of $G_{i,k}$ conditioning on $\{L_{i,k} = l\}$ can be bounded by the probability distribution of w_t , which enjoys good properties such as Assumption 8.7. More details of the proof are in Appendix 8.D.3.3.

In conclusion, Lemma 8.17 follows directly from Lemma 8.18 and Lemma 8.19. Combining Lemma 8.17 and Lemma 8.16 completes the proof of Theorem 8.8.

Remark 8.20 (Convergence rate of SME for general time series). *Similar to Theorem 2.4 in [111], our results for linear dynamical systems can also be generalized to general time series with linear responses:*

$$y_t = \theta^* z_t + w_t, \quad t \geq 0,$$

where $\mathcal{F}_t^y = \mathcal{F}(w_0, \dots, w_t, z_0, \dots, z_t)$, $y_t \in \mathbb{R}^{n_y}$ is measurable in $\mathcal{F}_t^y$ but not in $\mathcal{F}_{t-1}^y$. The SME algorithm is

$$\Theta_T^y = \bigcap_{t=0}^{T-1} \{\hat{\theta} : y_t - \hat{\theta} z_t \in \mathcal{W}\}.$$

Under Assumptions 8.2, 8.5, and 8.7, we have

$$\begin{aligned} \mathbb{P}(\text{diam}(\Theta_T^y) > \delta) &\leq \frac{T}{m} \tilde{O}(n_z^{2.5}) a_2^{n_z} \exp(-a_3 m) \\ &\quad + \tilde{O}((n_y n_z)^{2.5}) a_4^{n_y n_z} \left(1 - q_w \left(\frac{a_1 \delta}{4\sqrt{n_y}}\right)\right)^{\lceil T/m \rceil}, \end{aligned}$$

where a_1, a_2, a_3 are defined in Theorem 8.8 and $a_4 = \frac{4b_z \sqrt{n_y}}{a_1}$.

8.5 Applications to Robust Adaptive Control

Robust adaptive control usually involves two steps: updating the uncertainty set estimation, and designing robust controllers based on the updated uncertainty set. SME can be naturally applied to robust adaptive control as the updating rule of the uncertainty set estimation. To illustrate this, we discuss the applications of SME to two popular controllers, robust adaptive MPC and robust SLS. We focus on the implications of our convergence rates.

Application of SME to robust adaptive MPC. SME has long been adopted in the robust adaptive MPC design (see e.g., [57, 78, 79]). Despite the regret analysis for unconstrained MPC and its variants (e.g. [138, 63]), the non-asymptotic analysis for robust adaptive MPC remains unsolved.

Applying Theorem 8.8 straightforwardly, we can obtain a non-asymptotic estimation error bound for robust adaptive MPC below, which lays a foundation for future regret analysis. For simplicity, we consider a tight bound $\mathcal{W}$ is known below, but our results for unknown $\mathcal{W}$ can also be applied similarly.

Corollary 8.21. *Consider the robust adaptive MPC controller introduced in Example 8.6, where Θ_t is updated by SME and $\mathcal{W}$ is known.⁷ Under the conditions of Corollary 8.9, the estimation error for any $\hat{\theta}_T \in \Theta_T$ can be bounded by $\|\hat{\theta}_T - \theta^*\|_F \leq \tilde{O}(\frac{n_x^{1.5}n_z^2}{T})$ with high probability.*

Application of SME to robust SLS. Robust SLS has been proposed in [29] for robust constrained control under system uncertainties [29]. Since [29] assumes bounded disturbances, one can apply SME for the uncertainty set estimation in place of the LSE's confidence regions in [29]. Then, by leveraging Theorems 3.1, 4.1 in [29] and our Theorem 8.8, we can directly obtain a non-asymptotic sub-optimality gap for learning-based robust SLS with SME as the uncertainty set estimation. For simplicity, we consider a known tight bound $\mathcal{W}$, but our results for unknown $\mathcal{W}$ can also be similarly applied here.

Corollary 8.22. *Under the conditions in Theorem 3.1 in [29] and Corollary 8.9, for large enough T , we have $\frac{J(A^*, B^*, \hat{\mathbf{K}}) - J^*}{J^*} \leq \tilde{O}(n_x^{1.5}n_z^2/T)$, where $\hat{\mathbf{K}}$ denotes the robust SLS controller in [29] under the uncertainty set Θ_T constructed by SME, $J(A^*, B^*, \hat{\mathbf{K}}) = \lim_{T \rightarrow +\infty} \mathbb{E} \frac{1}{T} \sum_{t=0}^{T-1} (x_t^\top Q x_t + u_t^\top R u_t)$ denotes the infinite-horizon averaged total cost by implementing the robust SLS controller $\hat{\mathbf{K}}$, and J^* denotes the optimal infinite-horizon averaged total cost.*

8.6 Numerical Experiments

We evaluate the empirical performance of SME on various systems and applications. For all experiments, we use the 90% confidence regions of LSE computed by Lemma E.3 in [110] and Theorem 1 in [2] as the baseline. The details of the simulation settings are provided in Appendix 8.I.⁸

Comparison of SME, SME with loose bound, UCB-SME, and LSE. This experiment is based on the linearized longitudinal flight control dynamics of Boeing 747 as studied in recent literature on learning-based control of linear systems [58, 85].

In Figure 8.2, we show the diameters of SME, SME with loose disturbance bounds, and UCB-SME on the identification problem of the Boeing 747 dynamics with i.i.d. truncated Gaussian (Figure 8.2a)

⁷When $\mathcal{W}$ is unknown, Theorems 8.11-8.14 all apply.

⁸The code to reproduce all the experimental results can be found at <https://github.com/jy-cds/non-asymptotic-set-membership>.




Figure 8.3: Diameters of the uncertainty sets constructed by SME, UCB-SME, and LSE for systems with different dimensions.

and uniform (Figure 8.2b) disturbances. We use control actions sampled from a uniform distribution in both cases. In Figure 8.2c, we show that both the upper bound $\hat{w}_{\max}$ used for UCB-SME and the lower bound $\bar{w}_{\max}$ in (8.4) converge to the true bound $w_{\max}$ as T increases. The quantitative behaviors of SME and its variants are consistent with those predicted by our theoretical results. In particular, in Figure 8.2a and Figure 8.2b, SME and UCB-SME outperform the 90% confidence regions of LSE in both the magnitude and the convergence rate. In Figure 8.2c, we verify that the UCB estimation $\hat{w}_{\max}^{(T)}$ converges to the true disturbance bound $w_{\max}$ from above, while the estimation $\bar{w}_{\max}^{(T)}$ converges from below. It is worth noting that $\bar{w}_{\max}^{(T)}$ converges to $w_{\max}$ very quickly in the simulations, allowing $\bar{w}_{\max}^{(T)}$ to be another potential approximation of $w_{\max}$ for SME when T is very large.

Scaling with dimension. We compare the scaling of SME, SME-UCB, and LSE with respect to the system dimensions in Figure 8.3. We use an autonomous system $x_{t+1} = A^*x_t + w_t$, where $A \in \mathbb{R}^{n_x \times n_x}$ has varying n_x . Disturbances w_t are sampled from a truncated Gaussian distribution and uniform distribution with $w_{\max} = 2$. Surprisingly, the scaling of SME with respect to the dimension of the system is not significantly worse than that of LSE in the simulation. This suggests that the convergence rate in Corollary 8.9 can potentially be improved in terms of the dimension dependence, which is left for future investigation.

Application to robust adaptive MPC. We provide an example of the quantitative impact of using SME for adaptive robust MPC in Figure 8.4. We consider the task of constrained linear quadratic tracking problem as in [69]. The model uncertainty set is estimated online with SME and LSE's 90% confidence region. Control actions are computed using the tube-based robust MPC [100, 84] with the uncertainty sets. We also plot the optimal MPC controller with accurate model information. Thanks to the fast convergence of SME, the tracking performance of the tube-based robust MPC with SME estimation quickly coincides with OPT, while the same controller based on LSE's confidence region estimation converges more slowly.




Figure 8.4: Linear quadratic tracking of robust adaptive MPC based on SME, LSE's confidence regions, and the accurate model (OPT).

8.7 Concluding Remarks

This work provides the first convergence rates for SME in linear dynamical systems with bounded disturbances and discusses variants of SME with unknown bound on w_t . Numerical experiments demonstrate SME's promising performance under bounded disturbances.

Regarding future directions, this work only considers box constraints on w_t , so it is worth extending the analysis to more general constraints. This paper only measures the size of the uncertainty sets by their diameters. We leave for future work to consider other metrics, such as volume. Further, our bounds suffer poor dependence on the system dimension, which is not reflected in simulations. Hence, it is important to refine the bounds and discuss the fundamental limits. Another exciting direction is to speed up the computation of SME since the current computation complexity increases linearly with the sample size. The convergence rate of such algorithms is an important *open question*. Other interesting directions include the extensions of the SME analysis to nonlinear systems, where recent nonlinear system identification literature [105, 37] may provide insights; and analyzing SME in the presence of other uncertainties, e.g. measurement noises [102].

SME is a valid estimation for bounded non-stochastic disturbances [36, 86, 59, 75]. Thus, a fruitful direction is to study SME's convergence rates under non-stochastic w_t . Another method for uncertainty set estimation is the credible regions of Bayesian approaches, e.g. Thompson sampling for linear systems [52, 3] and Gaussian processes for nonlinear systems [35]. A future direction is to study the convergence rates of credible regions.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

Roadmap for the appendices

- Section 8.A introduces additional notation used throughout the Appendix.
- Section 8.B provides more literature review on LSE and SM, and a more detailed discussion on the technical contributions of this paper.
- Section 8.C provides more discussions on examples that satisfy Assumptions 8.5 and 8.7.
- Section 8.D presents the proof of Theorem 8.8. In particular, we provide helper lemmas in Section 8.D.1 and prove Lemma 8.16, Lemma 8.17 in Section 8.D.2 and Section 8.D.3 respectively. A more precise upper bound for Theorem 8.8 (without the $\tilde{O}(\cdot)$ notation) is provided in Appendix 8.D.4.
- Section 8.E presents a proof of Corollary 8.9
- Section 8.F provides a proof of Corollary 8.10.
- Section 8.G presents a proof of Theorem 8.11.
- Appendix 8.H provides proofs of Theorem 3, Corollary 3, and Theorem 4.
- Section 8.I provides details of the simulation.

8.A Additional notations

Let $\mathbb{S}_n(0, 1)$ denote the unit sphere in $\mathbb{R}^n$ in l_2 norm, i.e., $\mathbb{S}_n(0, 1) = \{x \in \mathbb{R}^n : \|x\|_2 = 1\}$. Let $\mathbb{S}_{n \times m}(0, 1)$ denote the unit sphere in $\mathbb{R}^{n \times m}$ with respect to the Frobenius norm, i.e., $\mathbb{S}_{n \times m}(0, 1) = \{M \in \mathbb{R}^{n \times m} : \|M\|_F = 1\}$. Let $\bar{\mathbb{B}}_n(0, 1)$ denote the closed unit ball in $\mathbb{R}^n$ in l_2 norm, i.e., $\bar{\mathbb{B}}_n(0, 1) = \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}$. Let $\bar{\mathbb{B}}_{n \times m}(0, 1)$ denote the closed unit ball in $\mathbb{R}^{n \times m}$ in Frobenius norm, i.e., $\bar{\mathbb{B}}_{n \times m}(0, 1) = \{M \in \mathbb{R}^{n \times m} : \|M\|_F \leq 1\}$. For a matrix $M \in \mathbb{R}^{n \times m}$, $\text{vec}(M)$ is the vectorization of M . Moreover, we define the inverse mapping of $\text{vec}(\cdot)$ as $\text{mat}(\cdot)$, i.e., for a vector $d \in \mathbb{R}^{nm}$, $\text{mat}(d) \in \mathbb{R}^{n \times m}$. Consider a σ -algebra $\mathcal{F}$ and a random variable X , we write $X \in \mathcal{F}$ if X is measurable with respect to $\mathcal{F}$, i.e., for all Borel measurable sets $B \subseteq \mathbb{R}$, we have $X^{-1}(B) \in \mathcal{F}$. We can similarly define $\mathcal{F}$ -measurable random matrices and random vectors. Further, consider a polyhedral $\mathcal{D} = \{x : Ax \leq b\}$, we write $\mathcal{D} \in \mathcal{F}$ if matrix A and vector b are measurable with respect to $\mathcal{F}$. Consider two symmetric matrices $A, B \in \mathbb{R}^{n \times n}$, we write $A \succeq B$ if $A - B$ is a positive definite matrix. We define $\min \emptyset = +\infty$. For a set $\mathcal{E}$, let $\mathbb{1}_{\mathcal{E}}$ denote the indicator function on $\mathcal{E}$. For a vector $x \in \mathbb{R}^n$, we use x^j to denote the j th coordinate of x . Throughout the paper, we use $\text{TrunGauss}(0, \sigma_w, [-w_{\max}, w_{\max}])$ to refer to the truncated Gaussian distribution generated by Gaussian distribution with zero mean and σ_w^2 variance with truncated range $[-w_{\max}, w_{\max}]$. The same applies to multi-variate truncated Gaussian distributions.

8.B More discussions on least square and set membership

System identification studies the problem of estimating the parameters of an unknown dynamical systems from trajectory data. There are two main classes of estimation methods: point estimator such as least square estimation (LSE), and set estimator such as set membership estimation (SME). In the following, we provide more discussions and literature review on LSE and SME. We will also discuss the major technical novelties of this work.

8.B.1 Least square estimation

For linear dynamical systems $x_{t+1} = A^*x_t + B^*u_t + w_t = \theta^*z_t + w_t$, given a trajectory of data $\{x_t, u_t\}_{t \geq 0}$, least square estimation generates a point estimator that minimizes the following quadratic error [120, 76]:

$$\hat{\theta}_{\text{LSE}} = \min_{\hat{\theta}} \sum_{t=0}^{T-1} \|x_{t+1} - \hat{\theta}z_t\|_2^2.$$

Least-square estimation is widely used and its convergence (rate) guarantees have been investigated for a long time. In particular, non-asymptotic convergence rate guarantees of LSE has become increasingly important as these guarantees are the foundations for non-asymptotic performance analysis of learning-based/adaptive control algorithms. Earlier non-asymptotic analysis of LSE focused on the simpler regression model $y_t = \theta^*x_t + w_t$, where x_t and y_t are independent [19, 125, 46].

Recently, there is one major breakthrough in [111] that provides LSE's convergence rate analysis for linear dynamical system $x_{t+1} = \theta^*z_t + w_t$, where x_{t+1} and $z_t = [x_t^\top, u_t^\top]^\top$ are correlated. More specifically, [111] establishes a fundamental property, *block-martingale small-ball (BMSB)*, to analyze LSE under correlated data. BMSB enables a long list of subsequent literature on LSE's non-asymptotic analysis for different types of dynamical systems, e.g., [92, 26, 143, 98, 34, 127, 118, 142, 65].

Though LSE is a point estimator, one can establish confidence region of LSE based on proper statistical assumptions on w_t . The pioneer works on the confidence region of LSE for linear dynamical systems are [1, 2], which construct ellipsoid confidence regions for LSE. Moreover, the non-asymptotic bounds on estimation errors established in [111, 29] can also be viewed as confidence bounds. Further, the estimation error $\tilde{O}(\frac{\sqrt{n_x + n_z}}{\sqrt{T}})$ has been shown to match the fundamental lower bound for any estimation methods for unbounded disturbances in [111]. However, these confidence bounds all rely on statistical inequalities, which may result in loose constant factors despite an optimal convergence rate. When applying these confidence bounds to robust control, where the controller is required to satisfy certain stability and constraint satisfaction properties for every possible system

in the confidence region, a loose constant factor will result in a larger confidence region and a more conservative control design. Finally, in robust control and many practical applications, the disturbances are usually bounded, and it will be interesting to see how the knowledge of the boundedness will improve the uncertainty set estimation.

On a side note, this paper is also related with the ambiguity set estimation for the transition probabilities in robust Markov decision processes [96]. There are attempts on improving the ambiguity set estimation based on LSE for less conservative robust MDP [96].

8.B.2 Set membership

Set membership is commonly used in robust control for uncertainty set estimation [87, 4, 115, 18, 137, 94, 93, 104]. There is a long history of research on SME for both deterministic disturbances, such as [11, 36, 55, 86, 59, 75], and stochastic disturbances, such as [10, 11, 56, 5, 80]. For the stochastic disturbances, both convergence and convergence rate analysis have been investigated under the persistent excitation (PE) condition. However, the existing convergence rates are only established for simpler regression problems, $y_t = \theta^* x_t + w_t$, where y_t and x_t are independent [5, 11, 10, 56].

Recently, [80] provided an initial attempt to establish the convergence guarantee of SME for linear dynamical systems $x_{t+1} = \theta^* z_t + w_t$ for correlated data x_{t+1} and z_t . However, [80] assumes that PE holds deterministically, and designs a special control design based on constrained optimization to satisfy PE deterministically. Therefore, the convergence for general control design and the convergence rate analysis remain open questions for correlated data arising from dynamical systems.

In this paper, we establish the convergence rate guarantees of SME on linear dynamical systems under the BMSB conditions in [111]. Compared with [80], BMSB condition can be satisfied by adding an i.i.d. random noise to a general class of control designs [65].

Technically, one major challenge of SME analysis compared with the LSE analysis is that the diameter of the membership set does not have an explicit formula, which is in stark contrast with LSE, where the point estimator is the solution to a quadratic program and has explicit form. A common trick to address this issue in the analysis of SME is to connect the diameter bound with the values of disturbances subsequences $\{w_{s_k}\}_{k \geq 0}$: it can be generally shown that a large diameter indicates that a long subsequence of disturbances are far away from the boundary of $\mathcal{W}$. However, existing construction methods of $\{w_{s_k}\}_{k \geq 0}$ will cause the time indices $\{s_k\}_{k \geq 0}$ to *correlate* with the realization of the sequences $\{x_t, u_t, w_t\}_{t \geq 0}$ [5, 80, 11].⁹ Consequently, in the correlated-data scenario and when PE does not hold deterministically, under the existing construction methods

⁹In [80], the correlation between $\{s_k\}_{k \geq 0}$ and $\{x_t, u_t, w_t\}_{t \geq 0}$ is via the PE condition, but [80] assume deterministic PE to avoid this correlation issue.

in [5, 80, 11], the probability of $\{w_{s_k}\}_{k \geq 0}$ with correlated time indices *cannot* be bounded by the probability of the *independent* sequence $\{w_t\}_{t \geq 0}$. One major **technical contribution** of this paper is to provide a novel construction of $\{w_{s_k}\}_{k \geq 0}$ based on a sequence of stopping times and establish conditional independence properties despite correlated data and stochastic PE condition (BMSB). More details can be found in Lemma 8.19 and the proof of Lemma 8.17.

Though we only consider box constraints for w_t , it is worth mentioning that SME can be applied to much more general forms of disturbances. For example, a common alternative is the ellipsoidal-bounded disturbance where $\mathcal{W} := \{w \in \mathcal{R}^{n_x} : w^\top P w \leq 1\}$ with positive definite $P \in \mathcal{R}^{n_x \times n_x}$ [11, 121, 32, 74] and polytopic-bounded disturbance $\mathcal{W} := \{w \in \mathcal{R}^{n_x} : G w \leq h\}$ for positive definite $G \in \mathcal{R}^{n_x \times n_x}$ and $h \in \mathcal{R}^{n_x}$ [36, 80, 79]. There are also SME literature assuming bounded energy of the disturbance sequences [11]. It is an interesting future direction to extend the analysis in this paper to more general disturbance constraints.

Further, exact SME involves the intersection of an increasing number of sets, thus causing the computation complexity increases with time t , which can become prohibitive when t is large. There are many methods trying to reduce the computation complexity by approximating the membership sets (see e.g., [75, 80], etc.). It is an exciting future direction to study the diameter bounds of the approximated SME methods.

Lastly, it is worth mentioning that SME can also be applied to the uncertainty set estimation in perception-based control [38, 116].

8.C More discussions on Assumptions 8.5 and 8.7

8.C.1 More discussions on Assumption 8.5

The BMSB condition has been widely used in learning-based control. It has been shown that BMSB can be satisfied in many scenarios. For example, [111, 119] showed that linear systems with i.i.d. perturbed linear control policies, i.e., $x_{t+1} = Ax_t + B(Kx_t + \eta_t) + w_t$,¹⁰ satisfy BMSB if the disturbances w_t and η_t are i.i.d. and follow Gaussian distributions with positive definite covariance matrices. Later, [29] showed that $x_{t+1} = Ax_t + B(Kx_t + \eta_t) + w_t$ can still satisfy BMSB even for non-Gaussian distributions of w_t, η_t , as long as w_t and η_t have independent coordinates and finite fourth moments. Recently, [65] extended the results to linear systems with nonlinear policies, i.e., $x_{t+1} = Ax_t + B(\pi_t(x_t) + \eta_t) + w_t$, and showed that BMSB still holds as long as the nonlinear policies π_t generate bounded trajectories of states and control inputs, and w_t, η_t are bounded and follow distributions with certain anti-concentrated properties (a special case is positive definite covariance matrix).

¹⁰Though we only describe a static linear policy $u_t = Kx_t$ here, the results in [111, 119, 29] hold for dynamic linear policies.

8.C.2 More discussions on Assumption 8.7

In this subsection, we provide two example distributions, truncated Gaussian and uniform distributions, and discuss their corresponding $q_w(\epsilon)$ functions. It will be shown that for both distributions below, $q_w(\epsilon) = O(\epsilon)$.

Lemma 8.23 (Example of uniform distribution). *Consider w_t that follows a uniform distribution on $[-w_{\max}, w_{\max}]^{n_x}$. Then, $q_w(\epsilon) = \frac{\epsilon}{2w_{\max}}$.*

Proof. Since $\text{Unif}(\mathcal{W})$ is symmetric, we only need to consider one direction $j = 1$.

$$\begin{aligned} \mathbb{P}(w^j + w_{\max} \leq \epsilon) &= \int_{w^1 + w_{\max} \leq \epsilon} \int_{w^2, \dots, w^{n_x} \in [-w_{\max}, w_{\max}]} \frac{1}{(2w_{\max})^{n_x}} \mathbb{1}_{(w \in \mathcal{W})} dw \\ &= \int_{w^1 \leq \epsilon - w_{\max}} \frac{1}{2w_{\max}} \mathbb{1}_{(w \in \mathcal{W})} dw^1 = \frac{\epsilon}{2w_{\max}} \end{aligned}$$

Similarly, $\mathbb{P}(w_{\max} - w^1 \leq \epsilon) = \int_{w^1 \geq w_{\max} - \epsilon} \frac{1}{2w_{\max}} \mathbb{1}_{(w \in \mathcal{W})} dw^1 = \frac{\epsilon}{2w_{\max}}$. $\square$

Lemma 8.24 (Example of truncated Gaussian distribution). *Consider w_t follows a truncated Gaussian distribution on $[-w_{\max}, w_{\max}]^{n_x}$ generated by a Gaussian distribution with zero mean and $\sigma_w I_{n_x}$ covariance matrix. Then, $q_w(\epsilon) = \frac{1}{\min(\sqrt{2\pi}\sigma_w, 2w_{\max})} \exp(-\frac{w_{\max}^2}{2\sigma_w^2})\epsilon$.*

Proof. Since this distribution is symmetric and each coordinate is independent, we only need to consider one direction j . Let X denote a Gaussian distribution with zero mean and σ_w^2 variance. By the definition of truncated Gaussian distributions, we have

$$\mathbb{P}(w^j + w_{\max} \leq \epsilon) = \frac{\mathbb{P}(-w_{\max} \leq X \leq -w_{\max} + \epsilon)}{\mathbb{P}(-w_{\max} \leq X \leq w_{\max})}$$

Notice that X/σ_w follows the standard Gaussian distribution, so we can obtain the following bounds.

$$\begin{aligned} \mathbb{P}(-w_{\max} \leq X \leq -w_{\max} + \epsilon) &= \int_{-w_{\max}/\sigma_w}^{(-w_{\max} + \epsilon)/\sigma_w} \frac{1}{\sqrt{2\pi}} \exp(-\frac{z^2}{2}) dz \\ &\geq \frac{1}{\sqrt{2\pi}} \exp(-w_{\max}^2/(2\sigma_w^2)) \frac{\epsilon}{\sigma_w} \end{aligned}$$

and

$$\begin{aligned} \mathbb{P}(-w_{\max} \leq X \leq w_{\max}) &= \int_{-w_{\max}/\sigma_w}^{w_{\max}/\sigma_w} \frac{1}{\sqrt{2\pi}} \exp(-\frac{z^2}{2}) dz \\ &\leq \min(1, \frac{1}{\sqrt{2\pi}} \frac{2w_{\max}}{\sigma_w}) \end{aligned}$$

Therefore, we obtain

$$\begin{aligned}
\mathbb{P}(w^j + w_{\max} \leq \epsilon) &= \frac{\mathbb{P}(-w_{\max} \leq X \leq -w_{\max} + \epsilon)}{\mathbb{P}(-w_{\max} \leq X \leq w_{\max})} \\
&\geq \max\left(\frac{1}{\sqrt{2\pi}} \exp(-w_{\max}^2/\sigma_w^2) \frac{\epsilon}{\sigma_w}, \frac{\epsilon}{2w_{\max}} \exp\left(\frac{-w_{\max}^2}{2\sigma_w^2}\right)\right) \\
&= \frac{1}{\min(\sqrt{2\pi}\sigma_w, 2w_{\max})} \exp\left(\frac{-w_{\max}^2}{2\sigma_w^2}\right) \epsilon
\end{aligned}$$

Finally, $\mathbb{P}(w_{\max} - w^1 \leq \epsilon)$ can be bounded similarly. $\square$

Lemma 8.25 (Example of uniform distribution on the boundary of $\mathcal{W}$ (a generalization of Rademacher distribution)). *Consider w_t follows a uniform distribution on $\{w : \|w\|_\infty = w_{\max}\}$. Then $q_w(\epsilon) = \frac{1}{2n_x}$.*

Proof. Since the hyper-cube $\{w : \|w\|_\infty = w_{\max}\}$ has $2n_x$ facets, the probability on each facet is $\frac{1}{2n_x}$. Therefore, $\mathbb{P}(w^j \leq \epsilon - w_{\max}) \geq \mathbb{P}(w^j = -w_{\max}) = \frac{1}{2n_x}$ for all j . The same applies to $\mathbb{P}(w^j \geq -\epsilon + w_{\max})$. $\square$

8.D Proof of Theorem 8.8

The section provides more details for the proof of Theorem 8.8. In particular, we first provide technical lemmas for set discretization, then prove Lemma 8.16 and Lemma 8.17 respectively. The proof of Theorem 8.8 follows naturally by combining Lemma 8.16 and Lemma 8.17.

8.D.1 Technical lemmas: set discretization

This subsection provide useful technical lemmas for the proofs of Lemma 8.16 and Lemma 8.17. The results are based on a finite-ball covering result that is classical in the literature [101] [124].

Theorem 8.26 (Theorem 1.1 and 1.2 in [124] and Theorem 2 in [101] (revised to match the setting of this paper)). *Consider a closed ball $\bar{\mathcal{B}}_n(0, 1) = \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}$ in l_2 norm. Considering covering this ball $\bar{\mathcal{B}}_n(0, 1)$ with smaller closed balls $\bar{\mathcal{B}}_n(z, \epsilon)$ for $z \in \mathbb{R}^n$. Let $v_{\epsilon, n}$ denote the minimal number of smaller balls needed to cover $\bar{\mathcal{B}}_n(0, 1)$. For $n \geq 1$ and $0 < \epsilon < 1/2$, we have*

$$v_{\epsilon, n} \leq 544n^{2.5} \log(n/\epsilon) \left(\frac{1}{\epsilon}\right)^n$$

Proof. Theorem 1.1 and 1.2 in [124] and Theorem 2 in [101] discuss the upper bounds of $v_{\epsilon, n}$ in several different cases. These upper bounds in these different cases are unified by the upper bound in the theorem above by algebraic manipulations. $\square$

We apply Theorem 8.26 to obtain the number of covering balls in the two settings below. These two settings will be considered in the proofs of Lemma 8.16 and 8.17 respectively.

Corollary 8.27. *There exists a finite set $\mathcal{M}' = \{\lambda_1, \dots, \lambda_{v_\lambda}\} \subseteq \mathbb{S}_{n_z}(0, 1)$ such that for any $\lambda \in \mathbb{R}^{n_z}$ with $\|\lambda\|_2 = 1$, there exists $\lambda_i \in \mathcal{M}'$ such that $\|\lambda - \lambda_i\|_2 \leq 2\epsilon_\lambda$.*

In the following, we consider $\epsilon_\lambda = \sigma_z^2 p_z^2 / (64b_z^2) = 1/a_2$. Notice that $\epsilon_\lambda < 1/2$. Accordingly,

$$v_\lambda \leq 544n_z^{2.5} \log(a_2 n_z) a_2^{n_z}. \quad (8.7)$$

Proof. $\epsilon_\lambda \leq 1/64 < 1/2$ because $p_z \leq 1$ and $\sigma_z \leq b_z$ by the definitions of BMSB and b_z . Then, the bound on v_λ follows from Theorem 8.26. $\square$

Lemma 8.28. *There exists a finite set $\mathcal{M} = \{\gamma_1, \dots, \gamma_{v_\gamma}\} \subseteq \mathbb{S}_{n_x \times n_z}(0, 1)$ such that for any $\gamma \in \mathbb{R}^{n_x \times n_z}$ and $\|\gamma\|_F = 1$, there exists $\gamma_i \in \mathcal{M}$ such that $\|\gamma - \gamma_i\|_F \leq 2\epsilon_\gamma$. Consider $\epsilon_\gamma = \frac{a_1}{4b_z \sqrt{n_x}} = 1/a_4$. Notice that $\epsilon_\gamma < 1/2$. Accordingly,*

$$v_\gamma \leq 544n_x^{2.5} n_z^{2.5} \log(a_4 n_x n_z) a_4^{n_z n_x}.$$

Proof. The proof is basically by mapping the matrices to vectors based on matrix vectorization, then mapping the vectors back to matrices. These two mappings are isomorphism.

Specifically, consider a closed unit ball in $\mathbb{R}^{n_x n_z}$. There exist $v_{\epsilon, n_x n_z}$ smaller closed balls to cover it, denoted by $\mathcal{B}_1, \dots, \mathcal{B}_{v_{\epsilon, n_x n_z}}$. Consider the non-empty sets from $\mathcal{B}_1 \cap \mathbb{S}_{n_x n_z}(0, 1), \dots, \mathcal{B}_{v_{\epsilon, n_x n_z}} \cap \mathbb{S}_{n_x n_z}(0, 1)$. For any $1 \leq i \leq v_{\epsilon, n_x n_z}$, if $\mathcal{B}_i \cap \mathbb{S}_{n_x n_z}(0, 1) \neq \emptyset$, select a point $\text{vec}(\gamma) \in \mathcal{B}_i \cap \mathbb{S}_{n_x n_z}(0, 1)$. Notice that $\|\text{vec}(\gamma)\|_2 = 1$. In this way, we construct a finite sequence $\{\text{vec}(\gamma_1), \dots, \text{vec}(\gamma_{v_\gamma})\}$ where $v_\gamma \leq v_{\epsilon_\gamma, n_x n_z}$.¹¹

For any $\gamma \in \mathbb{R}^{n_x \times n_z}$, we have $\text{vec}(\gamma) \in \mathbb{R}^{n_x n_z}$ and $\|\text{vec}(\gamma)\|_2 = 1$. Hence, there exists $1 \leq i \leq v_\gamma$ such that $\text{vec}(\gamma) \in \mathcal{B}_i \cap \mathbb{S}_{n_x n_z}(0, 1)$. Hence, $\|\text{vec}(\gamma) - \text{vec}(\gamma_i)\|_2 \leq 2\epsilon_\gamma$. Moreover, $\|\gamma_i\|_F = \|\text{vec}(\gamma_i)\|_2 = 1$. Therefore, $\|\gamma_i - \gamma\|_F \leq 2\epsilon_\gamma$. So the set $\mathcal{M} = \{\gamma_1, \dots, \gamma_{v_\gamma}\}$ satisfies our requirement. $\square$

8.D.2 Proof of Lemma 8.16

Essentially, Lemma 8.16 shows that PE holds with high probability under the BMSB condition. This result has been established in Proposition 2.5 in [111], though in a different form. The rest of this subsection will prove the PE condition needed in this paper based on Proposition 2.5 in [111].

Firstly, we review Proposition 2.5 in [111] for the convenience of the reader.

Theorem 8.29 (Proposition 2.5 in [111] when $k = 1$). *Let $\{Z_t\}_{t \geq 1}$ be an $\{\mathcal{F}_t^Z\}_{t \geq 1}$ -adapted random process taking values in $\mathbb{R}$. Z_0 is given. If $\{Z_t\}_{t \geq 0}$ is $(1, v, p)$ -BMSB, then*

$$\mathbb{P}\left(\sum_{t=1}^T Z_t^2 \leq v^2 p^2 T / 8\right) \leq \exp(-T p^2 / 8)$$

¹¹Here, without loss of generality, we consider $\mathcal{B}_1 \cap \mathbb{S}_{n_x n_z}(0, 1), \dots, \mathcal{B}_{v_\gamma} \cap \mathbb{S}_{n_x n_z}(0, 1)$ are not empty.

Next, we prove the PE in one segment of data sequence.

Lemma 8.30 (Probability of PE in one segment). *For any $m \geq 1$, for any $k \geq 0$, we have*

$$\mathbb{P}\left(\sum_{t=km+1}^{km+m} z_t z_t^\top \succ (\sigma_z^2 p_z^2 m / 16) I_{n_z} \mid \mathcal{F}_{km}\right) \geq 1 - v_\lambda \exp(-m p_z^2 / 8)$$

Proof. Consider $\mathcal{M}' = \{\lambda_1, \dots, \lambda_{v_\lambda}\}$ defined in Corollary 8.27. For any $\lambda_i \in \mathcal{M}'$, $\lambda_i^\top z_t$ satisfies the $(1, \sigma_z, p_z)$ -BMSB condition. Therefore, by Theorem 8.29, we have

$$\mathbb{P}\left(\sum_{t=1}^T \lambda_i^\top z_t z_t^\top \lambda_i \leq \sigma_z^2 p_z^2 T / 8\right) \leq \exp(-T p_z^2 / 8).$$

Notice that the horizon length T is arbitrary and the starting stage $t = 1$ can also be different because we consider a time-invariant dynamical system in this paper. Therefore, for any $m \geq 1, k \geq 0$, for any $\lambda_i \in \mathcal{M}'$, we have

$$\mathbb{P}\left(\sum_{i=1}^m \lambda_i^\top z_{km+i} z_{km+i}^\top \lambda_i \leq \sigma_z^2 p_z^2 m / 8 \mid \mathcal{F}_{km}\right) \leq \exp(-m p_z^2 / 8),$$

where we condition on $\mathcal{F}_{km}$ to make sure z_{km} is known under $\mathcal{F}_{km}$, which is required by Theorem 8.29.

For arbitrary λ such that $\|\lambda\|_2 = 1$, there exists $\lambda_i \in \mathcal{M}'$ such that $\|\lambda - \lambda_i\|_2 \leq 2\epsilon_\lambda$. Therefore, we can bound $\sum_{t=km+1}^{km+m} \lambda^\top z_t z_t^\top \lambda$ by $\sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i$.

$$\begin{aligned} \sum_{t=km+1}^{km+m} \lambda^\top z_t z_t^\top \lambda &= \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i + \sum_{t=km+1}^{km+m} (\lambda + \lambda_i)^\top z_t z_t^\top (\lambda - \lambda_i) \\ &\geq \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i - \sum_{t=km+1}^{km+m} \|\lambda + \lambda_i\|_2 \|z_t\|_2^2 \|\lambda_i - \lambda\|_2 \\ &\stackrel{(a)}{\geq} \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i - \sum_{t=km+1}^{km+m} 4b_z^2 \epsilon_\lambda \\ &= \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i - 4b_z^2 \epsilon_\lambda m \stackrel{(b)}{\geq} \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i - \sigma_z^2 p_z^2 m / 16 \end{aligned}$$

where (a) is by Assumption 8.5, $\|\lambda - \lambda_i\|_2 \leq 2\epsilon_\lambda$, and $\|\lambda\|_2 = \|\lambda_i\|_2 = 1$; and (b) is by choosing $\epsilon_\lambda \leq \sigma_z^2 p_z^2 / (64b_z^2)$.

Therefore, by the definition of positive definiteness and the inequalities above, we can complete the proof by the following.

$$\begin{aligned}
\mathbb{P}\left(\sum_{t=km+1}^{km+m} z_t z_t^\top \succ (\sigma_z^2 p_z^2 m/16) I_{n_z} \mid \mathcal{F}_{km}\right) &= \mathbb{P}(\forall \|\lambda\|_2 = 1, \sum_{t=km+1}^{km+m} \lambda^\top z_t z_t^\top \lambda > \sigma_z^2 p_z^2 m/16 \mid \mathcal{F}_{km}) \\
&\geq \mathbb{P}(\forall 1 \leq i \leq v_\lambda, \sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i > \sigma_z^2 p_z^2 m/8 \mid \mathcal{F}_{km}) \\
&\geq 1 - \sum_{i=1}^{v_\lambda} \mathbb{P}\left(\sum_{t=km+1}^{km+m} \lambda_i^\top z_t z_t^\top \lambda_i \leq \sigma_z^2 p_z^2 m/8 \mid \mathcal{F}_{km}\right) \\
&\geq 1 - v_\lambda \exp(-mp_z^2/8),
\end{aligned}$$

which completes the proof. $\square$

Now, we are ready for the proof of Lemma 8.16.

Proof of Lemma 8.16. Recall that $\mathcal{E}_2 = \{\frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z}, \forall 0 \leq k \leq \lceil T/m \rceil - 1\}$, where $a_1 = \sigma_z p_z/4$. Hence

$$\mathcal{E}_2 = \bigcap_{k=0}^{T/m-1} \left\{ \sum_{t=km+1}^{km+m} z_t z_t^\top \succ (\sigma_z^2 p_z^2 m/16) I_{n_z} \right\}.$$

Therefore,

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_2) &\geq 1 - \sum_{k=0}^{T/m-1} \mathbb{P}\left(\sum_{t=km+1}^{km+m} z_t z_t^\top \preceq (\sigma_z^2 p_z^2 m/16) I_{n_z}\right) \\
&\geq 1 - \frac{T}{m} v_\lambda \exp(-mp_z^2/8) \\
&= 1 - \frac{T}{m} (544 n_z^{2.5} \log(a_2 n_z) a_2^{n_z}) \exp(-mp_z^2/8),
\end{aligned}$$

where we use Lemma 8.30 and the fact that if $\mathbb{P}(\sum_{t=km+1}^{km+m} z_t z_t^\top \preceq (\sigma_z^2 p_z^2 m/16) I_{n_z} \mid \mathcal{F}_{km}) \leq v_\lambda \exp(-mp_z^2/8)$, then $\mathbb{P}(\sum_{t=km+1}^{km+m} z_t z_t^\top \preceq (\sigma_z^2 p_z^2 m/16) I_{n_z}) \leq v_\lambda \exp(-mp_z^2/8)$. $\square$

8.D.3 Proof of Lemma 8.17

This proof takes four major steps:

- (i) Define $b_{i,t}, j_{i,t}, L_{i,k}$.
- (ii) Provide a formal definition of $\mathcal{E}_{1,k}$ based on $b_{i,t}, j_{i,t}, L_{i,k}$ and prove a formal version of Lemma 8.18.

(iii) Prove Lemma 8.19.

(iv) Prove Lemma 8.17 by the formal version of Lemma 8.18 and Lemma 8.19.

It is worth mentioning that the formal definition of $\mathcal{E}_{1,k}$ is slightly different from the definition in Lemma 8.18, but we still have $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{1,k} \cap \mathcal{E}_2)$, which is the key property that will be used in the proof of Lemma 8.17.

8.D.3.1 Step (i): Definitions of $b_{i,t}, j_{i,t}, L_{i,k}$.

Recall the discretization of $\mathbb{S}_{n_x \times n_z}(0, 1)$ in Lemma 8.28, which generates the set $\mathcal{M} = \{\gamma_1, \dots, \gamma_{v_\gamma}\}$. We are going to define $b_{i,t}, j_{i,t}, L_{i,k}$ for $\gamma_i \in \mathcal{M}$ for each $1 \leq i \leq v_\gamma$. Notice that $\mathcal{M}$ is a deterministic set of matrices.

Lemma 8.31 (Definition of $b_{i,t}, j_{i,t}$). *For any $\gamma_i \in \mathcal{M}$, any $0 \leq t \leq T$, there exist $b_{i,t} \in \{-1, 1\}$ and $1 \leq j_{i,t} \leq n_x$ such that $b_{i,t}, j_{i,t} \in \mathcal{F}(z_t) \subseteq \mathcal{F}_t$ and*

$$\|\gamma_i z_t\|_\infty = b_{i,t}(\gamma_i z_t)^{j_{i,t}}.$$

Note that one way to determine $b_{i,t}, j_{i,t}$ from z_t is by the following: first pick the smallest j such that $|(\gamma_i z_t)^j| = \|\gamma_i z_t\|_\infty$, then let $b_{i,t} = \text{sgn}((\gamma_i z_t)^j)$, where $\text{sgn}(\cdot)$ denotes the sign of a scalar argument.

Proof. For any $\gamma_i \in \mathcal{M}$, any $0 \leq t \leq T$, we have

$$\|\gamma_i z_t\|_\infty = \max_{1 \leq j \leq n_x} \max_{b \in \{-1, 1\}} b(\gamma_i z_t)^j$$

Hence, there exist $b_{i,t}, j_{i,t}$ such that $\|\gamma_i z_t\|_\infty = b_{i,t}(\gamma_i z_t)^{j_{i,t}}$. Further, $b_{i,t}, j_{i,t}$ only depend on γ_i and z_t , so they are $\mathcal{F}(z_t)$ -measurable, and $\mathcal{F}(z_t) \subseteq \mathcal{F}_t$. $\square$

Lemma 8.32 (Definition of stopping times $L_{i,k}$). *Let $\eta = \frac{a_1}{\sqrt{n_x}}$. For any $\gamma_i \in \mathcal{M}$, any $0 \leq k \leq T/m - 1$, we can define a random time index $1 \leq L_{i,k} \leq m + 1$ by*

$$L_{i,k} = \min(m + 1, \min\{l \geq 1 : \|\gamma_i z_{km+l}\|_\infty \geq \eta\}).$$

Then, we have $1 \leq L_{i,k} \leq m + 1$. Further, for any $1 \leq l \leq m$, $\{L_{i,k} = l\} \in \mathcal{F}_{km+l}$, and $\{L_{i,k} = m + 1\} \in \mathcal{F}_{km+m} \subseteq \mathcal{F}_{km+m+1}$. In other words, $L_{i,k}$ is a stopping time with respect to filtration $\{\mathcal{F}_{km+l}\}_{l \geq 1}$.

Proof. For any i and any k , it is straightforward to see that $L_{i,k}$ is well-defined and $1 \leq L_{i,k} \leq m + 1$.

When $L_{i,k} = l \leq m$, this is equivalent with $\|\gamma_i z_{km+l}\|_\infty \geq \eta$ but $\|\gamma_i z_{km+s}\| < \eta$ for $1 \leq s < l$. Notice that this event is only determined by $z_{km+l}, \dots, z_{km+1}$, so $\{L_{i,k} = l\} \in \mathcal{F}_{km+l}$.

When $L_{i,k} = m + 1$, this is equivalent with $\|\gamma_i z_{km+s}\| < \eta$ for $1 \leq s \leq m$. Notice that this event is only determined by $z_{km+m}, \dots, z_{km+1}$, so $\{L_{i,k} = m + 1\} \in \mathcal{F}_{km+m}$.

Therefore, by definition, $L_{i,k}$ is a stopping time with respect to filtration $\{\mathcal{F}_{km+l}\}_{l \geq 1}$. $\square$

8.D.3.2 Step (ii): a formal version of Lemma 8.18 and its proof

Lemma 8.33 (Discretization of $\mathcal{E}_1 \cap \mathcal{E}_2$ (Formal version of Lemma 8.18)). *Let $\mathcal{M} = \{\gamma_1, \dots, \gamma_{v_\gamma}\}$ be an ϵ_γ -net of $\{\gamma : \|\gamma\|_F = 1\}$ as defined in Lemma 8.28, where $\epsilon_\gamma = \min(\frac{a_1}{4b_z\sqrt{n_x}}, 1)$, $v_\gamma = \tilde{O}(n_x^{2.5} n_z^{2.5}) a_4^{n_x n_z}$, and $a_4 = \frac{4b_z\sqrt{n_x}}{a_1}$. Define*

$$\mathcal{E}_{1,i} = \{\exists \gamma \in \Gamma_T, \text{ s.t. } b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1 \delta}{4\sqrt{n_x}}, \forall k \geq 0\}.$$

Then, we have

$$\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2).$$

The rest of this subsection is dedicated to the proof of Lemma 8.33. As an overview: firstly, we will discuss the implications of $\mathcal{E}_2$ on $\gamma_i \in \mathcal{M}$. Then, we discuss the implications of $\mathcal{E}_2$ on any γ . Lastly, we prove Lemma 8.33 by combining the implications of $\mathcal{E}_2$ on any γ and $\|\gamma\|_F \geq \delta/2$.

Lemma 8.34 (The implication of $\mathcal{E}_2$ on γ_i). *If $\mathcal{E}_2$ happens, then for any $\gamma_i \in \mathcal{M}$, any $0 \leq k \leq T/m - 1$, we have*

$$\max_{1 \leq s \leq m} \|\gamma_i z_{km+s}\|_\infty \geq \frac{a_1}{\sqrt{n_x}}.$$

Therefore, almost surely, we have $1 \leq L_{i,k} \leq m$ and

$$b_{i,km+L_{i,k}}(\gamma_i z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{\sqrt{n_x}}.$$

Proof. If $\mathcal{E}_2$ happens, then by definition, we have

$$\frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z},$$

for all $0 \leq k \leq T/m - 1$.

Now, for any $\gamma_i \in \mathcal{M}$, we have that

$$\frac{1}{m} \sum_{s=1}^m \gamma_i z_{km+s} z_{km+s}^\top \gamma_i^\top \succeq a_1^2 \gamma_i \gamma_i^\top. \quad (8.8)$$

Therefore, by taking trace at each side of (8.8), we obtain

$$\frac{1}{m} \sum_{s=1}^m \text{Tr}(\gamma_i z_{km+s} z_{km+s}^\top \gamma_i^\top) \geq a_1^2 \text{Tr}(\gamma_i \gamma_i^\top) \quad (8.9)$$

Since $\gamma_i \in \mathbb{S}_{n_x \times n_z}(0, 1)$, we have $\|\gamma_i\|_F = 1$, so $\text{Tr}(\gamma_i \gamma_i^\top) = \text{Tr}(\gamma_i^\top \gamma_i) = \|\gamma_i\|_F^2 = 1$. Further, we have

$$\text{Tr}(\gamma_i z_{km+s} z_{km+s}^\top \gamma_i^\top) = \text{Tr}(z_{km+s}^\top \gamma_i^\top \gamma_i z_{km+s}) = z_{km+s}^\top \gamma_i^\top \gamma_i z_{km+s} = \|\gamma_i z_{km+s}\|_2^2.$$

Consequently, we have

$$\frac{1}{m} \sum_{s=1}^m \|\gamma_i z_{km+s}\|_2^2 \geq a_1^2$$

for all k .

By the pigeonhole principle, we have that

$$\max_{1 \leq s \leq m} \|\gamma_i z_{km+s}\|_2^2 \geq a_1^2.$$

This is equivalent with $\max_{1 \leq s \leq m} \|\gamma_i z_{km+s}\|_2 \geq a_1$.

Notice that $\|\gamma_i z_{km+s}\|_2 \leq \sqrt{n_x} \|\gamma_i z_{km+s}\|_\infty$, so $\max_{1 \leq s \leq m} \sqrt{n_x} \|\gamma_i z_{km+s}\|_\infty \geq a_1$, which completes the proof of the first inequality in the lemma statement.

Next, we prove the second inequality in the lemma statement. Notice that by the definition of $L_{i,k}$ in Lemma 8.32 and by $\eta = \frac{a_1}{\sqrt{n_x}}$, we have $1 \leq L_{i,k} \leq m$ and $\|\gamma_i z_{km+L_{i,k}}\|_\infty \geq \frac{a_1}{\sqrt{n_x}}$ for all k . Further, by Lemma 8.31, we have $\|\gamma_i z_{km+L_{i,k}}\|_\infty = b_{i,km+L_{i,k}}(\gamma_i z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}}$ almost surely. Hence, we have $b_{i,km+L_{i,k}}(\gamma_i z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{\sqrt{n_x}}$, which completes the proof. $\square$

Lemma 8.35 (The implication of $\mathcal{E}_2$ on γ_{z_t}). *If $\mathcal{E}_2$ happens, then for any $\gamma \in \mathbb{R}^{n_x \times n_z}$, there exists $1 \leq i \leq v_\gamma$, such that*

$$b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{2\sqrt{n_x}} \|\gamma\|_F,$$

for all $0 \leq k \leq T/m - 1$.

Proof. Firstly, when $\gamma = 0$, the inequality holds because both sides are 0.

Next, when $\gamma \neq 0$, it suffices to prove $b_{i,km+L_{i,k}}(\frac{\gamma}{\|\gamma\|_F} z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{2\sqrt{n_x}}$. Therefore, we will only consider $\gamma \in \mathbb{S}_{n_x \times n_z}(0, 1)$. By Lemma 8.28, there exists $\gamma_i \in \mathcal{M}$ such that $\|\gamma - \gamma_i\|_F \leq 2\epsilon_\gamma = \min(\frac{a_1}{2b_z \sqrt{n_x}}, 2)$. Notice that by Lemma 8.34, if $\mathcal{E}_2$ happens, for all k , we have

$$b_{i,km+L_{i,k}}(\gamma_i z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{\sqrt{n_x}}.$$

Therefore,

$$\begin{aligned}
b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} &= b_{i,km+L_{i,k}}(\gamma_i z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \\
&\quad - b_{i,km+L_{i,k}}((\gamma_i - \gamma) z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \\
&\geq \frac{a_1}{\sqrt{n_x}} - |b_{i,km+L_{i,k}}((\gamma_i - \gamma) z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}}| \\
&\geq \frac{a_1}{\sqrt{n_x}} - \|(\gamma_i - \gamma) z_{km+L_{i,k}}\|_2 \\
&\geq \frac{a_1}{\sqrt{n_x}} - \|\gamma_i - \gamma\|_2 \|z_{km+L_{i,k}}\|_2 \\
&\geq \frac{a_1}{\sqrt{n_x}} - 2\epsilon_\gamma b_z \geq \frac{a_1}{2\sqrt{n_x}}
\end{aligned}$$

□

Proof of Lemma 8.33. By Lemma 8.35, under $\mathcal{E}_2$, for any $\gamma \in \mathbb{R}^{n_x \times n_z}$, there exists $1 \leq i \leq v_\gamma$, such that

$$b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{2\sqrt{n_x}} \|\gamma\|_F,$$

for all $0 \leq k \leq T/m - 1$. Therefore, if $\mathcal{E}_1 \cap \mathcal{E}_2$ happens, there exists $\gamma \in \Gamma_T$ and a corresponding i , such that

$$b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1}{2\sqrt{n_x}} \|\gamma\|_F \geq \frac{a_1 \delta}{4\sqrt{n_x}}.$$

Therefore,

$$\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \leq \mathbb{P}\left(\bigcup_{i=1}^{v_\gamma} \mathcal{E}_{1,i} \cap \mathcal{E}_2\right) \leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2),$$

which completes the proof. □

8.D.3.3 Proof of Lemma 8.19

Notice that Lemma 8.19 states two inequalities: in the following, we will first prove the first inequality $\mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2) \leq \mathbb{P}(\cap_{k=0}^{T/m-1} G_{i,k})$, then prove the second inequality on $\mathbb{P}(G_{i,k} \mid \cap_{k'=0}^{k-1} G_{i,k'})$.

Lemma 8.36 (Bound $\mathcal{E}_{1,i} \cap \mathcal{E}_2$ by $G_{i,k}$). *Under the conditions in Lemma 8.19, for any $1 \leq i \leq v_\gamma$, we have*

$$\mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2) \leq \mathbb{P}\left(\bigcap_{k=0}^{T/m-1} G_{i,k}\right).$$

Proof. Firstly, for any $\gamma \in \Gamma_T$, we have $\|w_t - \gamma z_t\|_\infty \leq w_{\max}$ for all $t \geq 0$. This suggests that, for any $1 \leq j \leq n_x$, we have

$$-w_{\max} \leq w_t^j - (\gamma z_t)^j \leq w_{\max}.$$

Hence, we have $b(\gamma z_t)^j \leq bw_t^j + w_{\max}$ for any $b \in \{-1, 1\}$, $1 \leq j \leq n_x$, and $t \geq 0$.

Next, by $\mathcal{E}_{1,i}$, there exists $\gamma \in \Gamma_T$ such that $b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1\delta}{4\sqrt{n_x}}$ for all $k \geq 0$.

Therefore, $b_{i,km+L_{i,k}}w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}$ for all k .

Finally, $\mathcal{E}_{1,i} \cap \mathcal{E}_2$ implies that $b_{i,km+L_{i,k}}w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}$ and $\frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z}$ for all k , which is $\cap_k G_{i,k}$ by the definition of $G_{i,k}$.

□

Lemma 8.37 (Bound on $\mathbb{P}(G_{i,k} \mid \cap_{k'=0}^{k-1} G_{i,k'})$). *Under the conditions in Lemma 8.19, for any $1 \leq i \leq v_\gamma$ and any $k \geq 0$, we have*

$$\mathbb{P}(G_{i,k} \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) \leq 1 - q_w\left(\frac{a_1\delta}{4\sqrt{n_x}}\right).$$

Proof. Firstly, notice that when $\frac{1}{m} \sum_{s=1}^m z_{km+s} z_{km+s}^\top \succeq a_1^2 I_{n_z}$, we have $1 \leq L_{i,k} \leq m$ by the proof of Lemma 8.34. Therefore, we have

$$\begin{aligned} \mathbb{P}(G_{i,k} \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) &\leq \mathbb{P}(b_{i,km+L_{i,k}}w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}, 1 \leq L_{i,k} \leq m \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) \\ &\leq \sum_{l=1}^m \mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}, L_{i,k} = l \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) \\ &\leq \sum_{l=1}^m \mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}} \mid L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}) \mathbb{P}(L_{i,k} = l \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) \\ &\stackrel{(a)}{\leq} (1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}})) \sum_{l=1}^m \mathbb{P}(L_{i,k} = l \mid \bigcap_{k'=0}^{k-1} G_{i,k'}) \\ &\leq 1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}}) \end{aligned}$$

The inequality (c) is proved in the following:

$$\begin{aligned}
& \mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}} \mid L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}) \\
&= \int_{v_{0:km+l}} \mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}, w_{0:km+l} = v_{0:km+l} \mid L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}) dv_{0:km+l} \\
&= \int_{v_{0:km+l} \in S_{km+l}} \mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}} \mid w_{0:km+l} = v_{0:km+l}) \\
&\quad \times \mathbb{P}(w_{0:km+l} = v_{0:km+l} \mid L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}) dv_{0:km+l} \\
&\stackrel{(b)}{\leq} (1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}})) \int_{v_{0:km+l} \in S_{km+l}} \mathbb{P}(w_{0:km+l} = v_{0:km+l} \mid L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}) dv_{0:km+l} \\
&= 1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}}),
\end{aligned}$$

where we define a shorthand notation $w_{0:km+l} = (w_0, \dots, w_{km+l-1})$, and we use $v_{0:km+l}$ to denote a realization of $w_{0:km+l}$, then we define the set of values of $w_{0:km+l}$ as S_{km+l} such that $L_{i,k} = l, \bigcap_{k'=0}^{k-1} G_{i,k'}$ holds. Notice that $L_{i,k} = l$ can be determined by a set of values of $w_{0:km+l}$ because $L_{i,k}$ is a stopping time of $\{F_{km+l}\}_{l \geq 1}$ and thus $\{L_{i,k} = l\} \in \mathcal{F}_{km+l}$. The inequality (b) above is because of the following: firstly, notice that $b_{i,km+l}, j_{i,km+l} \in \mathcal{F}_{km+l}$, so $b_{i,km+l}, j_{i,km+l}$ are deterministic values when $w_{0:km+l} = v_{0:km+l}$. Further, since w_{km+l} is independent of $w_{0:km+l}$, we have $\mathbb{P}(w_{\max} + bw_{km+l}^j \geq \epsilon \mid w_{0:km+l} = v_{0:km+l}) \leq 1 - q_w(\epsilon)$ for any deterministic b, j and any $\epsilon > 0$ by Assumption 8.7. Hence, we have $\mathbb{P}(b_{i,km+l}w_{km+l}^{j_{i,km+l}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}} \mid w_{0:km+l} = v_{0:km+l}) \leq 1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}})$. $\square$

8.D.3.4 Proof of Lemma 8.17

The proof is by leveraging Lemma 8.33 and Lemma 8.19.

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) &\leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{1,i} \cap \mathcal{E}_2) \\
&\leq \sum_{i=1}^{v_\gamma} \mathbb{P}\left(\bigcap_{k=0}^{T/m-1} G_{i,k}\right) \\
&= \sum_{i=1}^{v_\gamma} \mathbb{P}(G_{i,0}) \mathbb{P}(G_{i,1} \mid G_{i,0}) \cdots \mathbb{P}(G_{i,T/m-1} \mid \bigcap_{k=0}^{T/m-2} G_{i,k}) \\
&\leq \sum_{i=1}^{v_\gamma} \left(1 - q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{T/m} \\
&\leq 544 n_x^{2.5} n_z^{2.5} \log(a_4 n_x n_z) a_4^{n_z n_x} \left(1 - q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{T/m}.
\end{aligned}$$

8.D.4 A more precise upper bound for Theorem 8.8

By the proof of Lemma 8.16 and Lemma 8.17 above, we have

$$\begin{aligned}
\mathbb{P}(\text{diam}(\Theta_T) > \delta) &\leq 544 \frac{T}{m} n_z^{2.5} \log(a_2 n_z) a_2^{n_z} \exp(-a_3 m) \\
&\quad + 544 n_x^{2.5} n_z^{2.5} \log(a_4 n_x n_z) a_4^{n_z n_x} \left(1 - q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{T/m}
\end{aligned} \tag{8.10}$$

8.E Proof of Corollary 8.9

The proof involves two parts. Firstly, we will show that Term 1 $\leq \epsilon$ under our choice of m . Secondly, we will let Term 2 = ϵ , then we will show $\delta \leq \tilde{O}(n_x^{1.5} n_z^2 / T)$, which completes the proof.

Step 1: show Term 1 $\leq \epsilon$. Notice that when $m \geq \frac{1}{a_3} (\log(\frac{T}{\epsilon}) + n_z \log(a_2) + 2.5 \log(n_z) + \log \log(a_2 n_z) + 7) = O(n_z + \log T + \log(1/\epsilon))$, we have $T \tilde{O}(n_z^{2.5}) a_2^{n_z} \exp(-a_3 m) \leq \epsilon$. Since $m \geq 1$, we obtain Term 1 $\leq \epsilon$.

Step 2: let Term 2 = ϵ and show $\delta \leq \tilde{O}(n_x^{1.5} n_z^2 / T)$. Let Term 2 = ϵ , then we have $(1 - q_w(\frac{a_1 \delta}{4\sqrt{n_x}}))^{T/m} = \frac{\epsilon}{\tilde{O}(n_x^{2.5} n_z^{2.5}) a_4^{n_x n_z}}$. Then, we obtain $(1 - q_w(\frac{a_1 \delta}{4\sqrt{n_x}})) = \left(\frac{\epsilon}{\tilde{O}(n_x^{2.5} n_z^{2.5}) a_4^{n_x n_z}} \right)^{m/T}$, which is equivalent with

$$q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right) = 1 - \left(\frac{\epsilon}{\tilde{O}(n_x^{2.5} n_z^{2.5}) a_4^{n_x n_z}} \right)^{m/T}.$$

When $q_w(\frac{a_1\delta}{4\sqrt{n_x}}) = O(\frac{a_1\delta}{4\sqrt{n_x}})$, we obtain

$$\begin{aligned}\delta &= O\left(\frac{4\sqrt{n_x}}{a_1}\right) \left(1 - \left(\frac{\epsilon}{\tilde{O}(n_x^{2.5}n_z^{2.5})a_4^{n_x n_z}}\right)^{m/T}\right) \\ &\leq O\left(\frac{4\sqrt{n_x}}{a_1}\right) \log\left(\left(\frac{\epsilon}{\tilde{O}(n_x^{2.5}n_z^{2.5})a_4^{n_x n_z}}\right)^{m/T}\right) \\ &= O\left(\frac{4\sqrt{n_x}m}{a_1 T}\right) (\log(1/\epsilon) + n_x n_z + \log(n_x n_z)) \\ &= \tilde{O}\left(\frac{n_x^{1.5}n_z^2}{T}\right).\end{aligned}$$

Step 3: prove Corollary 8.9. By leveraging the bounds above and Theorem 8.8, we have $\mathbb{P}(\text{diam}(\Theta_T) \leq \tilde{O}\left(\frac{n_x^{1.5}n_z^2}{T}\right)) \geq \mathbb{P}(\text{diam}(\Theta_T) \leq \delta) \geq 1 - 2\epsilon$.

Since $\theta^* \in \Theta_T$ by definition, for any $\hat{\theta}_T \in \Theta_T$, we have $\|\hat{\theta}_T - \theta^*\|_F \leq \text{diam}(\Theta_T) \leq \tilde{O}\left(\frac{n_x^{1.5}n_z^2}{T}\right)$ with probability at least $1 - 2\epsilon$.

8.F Proof of Corollary 8.10

We provide a formal version of Corollary 8.10 and its proof below.

Corollary 8.38 (Convergence rate when $B^* = 0$ (formal version)). *When A^* is (κ, ρ) -stable, i.e., $\|(A^*)^t\|_2 \leq \kappa(1 - \rho)^t$ for all t with $\rho < 1$, for any $m > 0$ and any $\delta > 0$, when $T > m$, we have*

$$\mathbb{P}(\text{diam}(\mathbb{A}_T) > \delta) \leq \frac{T}{m} \tilde{O}(n_x^{2.5})a_2^{n_x} \exp(-a_3 m) + \tilde{O}(n_x^5)a_4^{n_x^2} (1 - q_w(\frac{a_1\delta}{4\sqrt{n_x}}))^{[T/m]}$$

where $b_x = \kappa\|x_0\|_2 + \kappa\sqrt{n_x}/\rho$, $p_x = 1/192$, $\sigma_x = \sqrt{\lambda_{\min}(\Sigma_w)/2}$, $a_1 = \frac{\sigma_x p_x}{4}$, $a_2 = \frac{64w_{\max}}{\sigma_x^2 p_x^2}$, $a_3 = \frac{p_x^2}{8}$, $a_4 = \frac{4b_x\sqrt{n_x}}{a_1}$.

Consequently, when the distribution of w_t satisfies $q_w(\epsilon) = O(\epsilon)$, e.g. uniform or truncated Gaussian, we have $\|\hat{\theta} - \theta^*\| \leq \tilde{O}(n_x^{3.5}/T)$.

The proof of Corollary 8.10 is exactly the same as the proofs of Theorem 8.8 and Corollary 8.9. When A^* is stable, we can show that $\|x_t\|_2 \leq b_x$ for all t . Further, by [29], the sequence $\{x_t\}_{t \geq 0}$ satisfies the $(1, \sigma_x, p_x)$ -BMSB condition. Therefore, we complete the proof.

8.G Proof of Theorem 8.11

Specifically, we define $\epsilon_0 = \frac{4\sqrt{n_x}}{a_1}(\hat{w}_{\max} - w_{\max})$.

The proof is similar to the proof of Theorem 8.8. Firstly, we define $\hat{\Gamma}_T$ as a translation of the set $\hat{\Theta}_T$:

$$\hat{\Gamma}_t = \bigcap_{s=0}^{t-1} \{\gamma : \|w_s - \gamma z_s\|_\infty \leq \hat{w}_{\max}\}, \quad \forall t \geq 0. \quad (8.11)$$

Notice that

$$\hat{\Theta}_T = \theta^* + \hat{\Gamma}_T$$

by considering $\gamma = \hat{\theta} - \theta^*$. Therefore, we can upper bound our goal event $\{\text{diam}(\hat{\Theta}_T) > \delta + \epsilon_0\}$ by the event $\mathcal{E}_3$ defined below.

$$\mathbb{P}(\text{diam}(\hat{\Theta}_T) > \delta + \epsilon_0) \leq \mathbb{P}(\mathcal{E}_3), \text{ where } \mathcal{E}_3 := \{\exists \gamma \in \hat{\Gamma}_T, \text{ s.t. } \|\gamma\|_F \geq \frac{\delta + \epsilon_0}{2}\}. \quad (8.12)$$

Next, notice that

$$\mathbb{P}(\text{diam}(\hat{\Theta}_T) > \delta + \epsilon_0) \leq \mathbb{P}(\mathcal{E}_3) \leq \mathbb{P}(\mathcal{E}_3 \cap \mathcal{E}_2) + \mathbb{P}(\mathcal{E}_2^c)$$

By Lemma 8.16, we have already shown $\mathbb{P}(\mathcal{E}_2^c) \leq \text{Term 1}$. So we only need to discuss $\mathbb{P}(\mathcal{E}_3 \cap \mathcal{E}_2)$.

Lemma 8.39.

$$\mathbb{P}(\mathcal{E}_3 \cap \mathcal{E}_2) \leq \text{Term 2}$$

Proof. Firstly, define

$$\mathcal{E}_{3,i} = \{\exists \gamma \in \hat{\Gamma}_T, \text{ s.t. } b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1(\delta + \epsilon_0)}{4\sqrt{n_x}}, \forall k \geq 0\}.$$

We have $\mathbb{P}(\mathcal{E}_3 \cap \mathcal{E}_2) \leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{3,i} \cap \mathcal{E}_2)$ based on the same proof ideas of Lemma 8.33.

Next, we will show that

$$\Pr(\mathcal{E}_{3,k} \cap \mathcal{E}_2) \leq \mathbb{P}\left(\bigcap_{k=0}^{T/m-1} G_{i,k}\right) \quad (8.13)$$

This is because for any $\gamma \in \hat{\Gamma}_T$, we have $b(\gamma z_t)^j \leq bw_t^j + \hat{w}_{\max}$ for any $b \in \{-1, 1\}$, $1 \leq j \leq n_x$, and $t \geq 0$. By $\mathcal{E}_{3,i}$, there exists $\gamma \in \hat{\Gamma}_T$ such that $b_{i,km+L_{i,k}}(\gamma z_{km+L_{i,k}})^{j_{i,km+L_{i,k}}} \geq \frac{a_1(\delta + \epsilon_0)}{4\sqrt{n_x}}$ for all $k \geq 0$. Thus, $b_{i,km+L_{i,k}}w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} + \hat{w}_{\max} \geq \frac{a_1(\delta + \epsilon_0)}{4\sqrt{n_x}}$ for all k . Notice that this is equivalent with $b_{i,km+L_{i,k}}w_{km+L_{i,k}}^{j_{i,km+L_{i,k}}} + w_{\max} \geq \frac{a_1\delta}{4\sqrt{n_x}}$ for all k because $\epsilon_0 = \frac{4\sqrt{n_x}}{a_1}(\hat{w}_{\max} - w_{\max})$. In this way, we can prove (8.13).

Finally, we can complete the proof by the following.

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_3 \cap \mathcal{E}_2) &\leq \sum_{i=1}^{v_\gamma} \mathbb{P}(\mathcal{E}_{3,i} \cap \mathcal{E}_2) \leq \sum_{i=1}^{v_\gamma} \mathbb{P}\left(\bigcap_{k=0}^{T/m-1} G_{i,k}\right) \\
&= \sum_{i=1}^{v_\gamma} \mathbb{P}(G_{i,0}) \mathbb{P}(G_{i,1} \mid G_{i,0}) \cdots \mathbb{P}(G_{i,T/m-1} \mid \bigcap_{k=0}^{T/m-2} G_{i,k}) \\
&\leq \sum_{i=1}^{v_\gamma} \left(1 - q_w\left(\frac{a_1 \delta}{4\sqrt{n_x}}\right)\right)^{T/m} \leq \text{Term 1}
\end{aligned}$$

where the second last inequality is by Lemma 8.37 and the last inequality uses the definition of v_γ in Lemma 8.28. $\square$

8.H Proofs of Theorem 8.12, Corollary 8.13, and Theorem 8.14

This section provides proofs of the main results related to the SME with unknown $w_{\max}$ as discussed in section 8.3.2. Namely, Theorem 8.12 and Corollary 8.13 provide the rate of convergence of the estimator $\bar{w}_{\max}^{(T)}$ defined in (8.4) to $w_{\max}$, and Theorem 8.14 states the rate of convergence of UCB-SME algorithm introduced in (8.5).

For ease of notation, we introduce the following function indexed by the time horizon $T > 0$,

$$W_T : \theta \mapsto \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta z_t\|_\infty. \quad (8.14)$$

The estimator $\bar{w}_{\max}^{(T)}$ is simply the infimum of this function, i.e., $\bar{w}_{\max}^{(T)} = \inf_\theta W_T(\theta)$.

8.H.1 Proof of Theorem 8.12

The proof of Theorem 8.12 involves two steps:

- *Step 1:* We demonstrate that the learning error of $w_{\max}$ incurred by the estimator $\bar{w}_{\max}^{(T)}$ is governed by the diameter of the uncertainty set Θ_T and the minimum learning error achievable if θ^* were known.
- *Step 2:* We then provide an upper bound the probability of learning error exceeding a fixed threshold.

Before we proceed with the proof of Theorem 8.12, we present the the following technical lemma.

Lemma 8.40. *Consider the sequence of functions $\{W_T\}_{T>0}$ defined in (8.14). The following holds:*

- i. W_T is convex in $\mathbb{R}^{n_x \times n_z}$,

ii. The sequence $\{\inf_{\theta} W_T(\theta)\}_{T>0}$ is bounded and monotonically non-decreasing, i.e.,

$$0 \leq \inf_{\theta} W_T(\theta) \leq \inf_{\theta} W_{T+1}(\theta) \leq w_{\max},$$

for all $T > 0$,

iii. W_T attains its minimum in Θ_T , i.e., $\arg \min_{\theta} W_T(\theta) \subset \Theta_T$.

Proof. (i.) For $0 \leq t \leq T-1$, the function $\theta \mapsto \|x_{t+1} - \theta z_t\|_{\infty}$ is convex due to convexity of norms. Since the maximum of convex functions is convex [16], convexity of W_T follows.

(ii.) Notice that W_{T+1} can be defined in terms of W_T recursively as $W_{T+1}(\theta) = \max(W_T(\theta), \|x_{T+1} - \theta z_T\|_{\infty})$. Thus, $W_T(\theta) \leq W_{T+1}(\theta)$ for all $\theta \in \mathbb{R}^{n_x \times n_z}$, implying monotonicity of $\{\inf_{\theta} W_T(\theta)\}_{T>0}$. To see boundedness, first notice that

$$W_T(\theta^*) = \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta^* z_t\|_{\infty} = \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} \leq w_{\max},$$

since $x_{t+1} = \theta^* z_t + w_t$. Therefore, for any $T > 0$, we have that

$$\inf_{\theta} W_T(\theta) = \inf_{\theta} \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta z_t\|_{\infty} \leq \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta^* z_t\|_{\infty} \leq w_{\max}$$

.

(iii.) First, we show that W_T attains its minimum on $\mathbb{R}^{n_x \times n_z}$. If $z_t = 0$ for $t \in [T]$, then W_T is a constant function and any $\theta \in \mathbb{R}^{n_x \times n_z}$ is a minimum of W_T . Now, suppose $z_t \neq 0$ for some $t \in [T]$. Then, W_T diverges at the infinity, i.e., $\lim_{k \rightarrow \infty} W_T(\theta_k) = \infty$ for any sequence $\{\theta_k\}_{k \in \mathcal{N}}$ such that $\|\theta_k\| \rightarrow \infty$ as $k \rightarrow \infty$. Since W_T is convex and bounded below with finite infimum, there exists a global minimizer $\bar{\theta}_T \in \mathbb{R}^{n_x \times n_z}$ such that $W_T(\bar{\theta}_T) = \inf_{\theta} W_T(\theta) = \bar{w}_{\max}^{(T)}$. Furthermore, by (ii), we have that $\|x_{t+1} - \bar{\theta}_T z_t\|_{\infty} \leq w_{\max}$ for all $t \in [T]$ and any global minimizer $\bar{\theta}_T \in \arg \min_{\theta} W_T(\theta)$, hence $\bar{\theta}_T \in \Theta_T$ by definition. $\square$

Step 1 of the Proof of Theorem 8.12: We first show that the error margin of the estimate $\bar{w}_{\max}^{(T)}$ from $w_{\max}$ is governed by the sum of two factors: (i) the diameter of Θ_T , which arises due to the lack of knowledge of θ^* , and (ii) the minimum learning error achievable if θ^* were known, namely

$$0 \leq w_{\max} - \bar{w}_{\max}^{(T)} \leq b_z \text{diam}(\Theta_T) + w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty}. \quad (8.15)$$

First, $0 \leq w_{\max} - \bar{w}_{\max}^{(T)}$ is simply due to Lemma 8.40. Next, we prove the second inequality $w_{\max} - \bar{w}_{\max}^{(T)} \leq b_z \text{diam}(\Theta_T) + w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty}$. By Lemma 8.40, there exists $\bar{\theta}_T \in \Theta_T$

such that $W_T(\bar{\theta}_T) = w_{\max}$ and

$$\begin{aligned} w_{\max} &= \max_{0 \leq t \leq T-1} \|x_{t+1} - \bar{\theta}_T z_t\|_{\infty}, \\ &= \max_{0 \leq t \leq T-1} \|x_{t+1} - \theta^* z_t + (\theta^* - \bar{\theta}_T) z_t\|_{\infty}, \\ &\geq \max_{0 \leq t \leq T-1} (\|x_{t+1} - \theta^* z_t\|_{\infty} - \|(\theta^* - \bar{\theta}_T) z_t\|_{\infty}), \end{aligned}$$

where the inequality is due to reverse triangle inequality. Furthermore, by using the equivalence of ℓ_2 and ℓ_{∞} norms, i.e., $\|x\|_2 \leq \|x\|_{\infty}$ for $x \in \mathbb{R}^{n_x}$, we bound $w_{\max}$ further below by

$$\begin{aligned} w_{\max} &\geq \max_{0 \leq t \leq T-1} (\|x_{t+1} - \theta^* z_t\|_{\infty} - \|(\theta^* - \bar{\theta}_T) z_t\|_2), \\ &\geq \max_{0 \leq t \leq T-1} (\|x_{t+1} - \theta^* z_t\|_{\infty} - \|\theta^* - \bar{\theta}_T\|_2 \|z_t\|_2), \\ &\geq \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} - b_z \text{diam}(\Theta_T), \end{aligned}$$

where the second inequality is due to $\|\theta^* - \bar{\theta}_T\|_2 := \sup_{z \neq 0} \frac{\|(\theta^* - \bar{\theta}_T)z\|_2}{\|z\|_2} \leq \frac{\|(\theta^* - \bar{\theta}_T)z_t\|_2}{\|z_t\|_2}$ and the third inequality follows from the assumption $\|z_t\|_2 \leq b_z$, the equivalence of Frobenius and spectral norms $\|\theta^* - \bar{\theta}_T\|_2 \leq \|\theta^* - \bar{\theta}_T\|_F$, and $\theta^*, \bar{\theta}_T \in \Theta_T$. Consequently,

$$w_{\max} - \bar{w}_{\max}^{(T)} \leq w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} + b_z \text{diam}(\Theta_T).$$

This completes the proof of the first step. $\square$

Step 2 of the Proof of Theorem 8.12: Using the learning error bound in (8.15), we obtain an upper bound on the probability of learning error exceeding a fixed $\delta > 0$ as shown below

$$\mathbb{P}(w_{\max} - \bar{w}_{\max}^{(T)} > \delta) \leq \mathcal{T}_1 + \mathcal{T}_2 \left(\frac{\delta}{2b_z} \right) + \mathcal{T}_5 \left(\frac{\delta}{2} \right), \quad (8.16)$$

where $\mathcal{T}_5(\delta) := (1 - q_w(\delta))^T$.

First, using the fact that $\{w_t\}_{t=0}^{T-1}$ are iid, we show that

$$\begin{aligned} \mathbb{P} \left(w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} > \delta \right) &= \mathbb{P}(w_{\max} - \delta > \|w_t\|_{\infty}, \forall 0 \leq t \leq T-1), \\ &= \prod_{t=0}^{T-1} \mathbb{P}(w_{\max} - \delta > \|w_t\|_{\infty}), \\ &\leq \prod_{t=0}^{T-1} \mathbb{P}(w_{\max} - \delta > w_t^1), \\ &\leq (1 - q_w(\delta))^T, \end{aligned}$$

where the first inequality is due to $w_t^1 \leq \|w_t\|_\infty$ and the second inequality is from Assumption 8.7. Finally, we obtain the desired convergence rate using the error bound in (8.15) as follows

$$\begin{aligned}
\mathbb{P}(w_{\max} - \bar{w}_{\max}^{(T)} > \delta) &\leq \mathbb{P}\left(b_z \text{diam}(\Theta_T) + w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_\infty > \delta\right) \\
&\leq \mathbb{P}\left(b_z \text{diam}(\Theta_T) > \delta/2 \text{ or } w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_\infty > \delta/2\right) \\
&\leq \mathbb{P}\left(\text{diam}(\Theta_T) > \frac{\delta}{2b_z}\right) + \mathbb{P}\left(w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_\infty > \delta/2\right) \\
&\leq \mathcal{T}_1 + \mathcal{T}_2\left(\frac{\delta}{2b_z}\right) + \mathcal{T}_5(\delta/2).
\end{aligned}$$

where the last inequality is by Theorem 8.8.

This completes the second and the last step of the proof. $\square$

8.H.2 Proof of Corollary 8.13

First, by the proof of Corollary 8.9 in section 8.E, we have that $\mathcal{T}_1 = \frac{T}{m} \tilde{O}(n_z^{2.5}) a_2^{n_z} \exp(-a_3 m) \leq \epsilon$ whenever $m \geq O(n_z + \log T + \log \frac{1}{\epsilon})$.

Next, we show $\mathcal{T}_5(\delta_T/2) \leq \mathcal{T}_2(\frac{\delta_T}{2b_z})$. Since $b_z \geq \sigma_z$ by the definition of BMSB, we have $\frac{a_1 \delta_T}{8\sqrt{n_x} b_z} \leq \frac{\delta_T}{2}$. Since $q_w(\cdot)$ is a non-decreasing function, we have $1 - q_w(\frac{a_1 \delta_T}{8\sqrt{n_x} b_z}) \geq 1 - q_w(\frac{\delta_T}{2})$. Notice that $m \geq 1$, and the constant factors in front of the $(1 - q_w(\cdot))^{\lceil T/m \rceil}$ in $\mathcal{T}_2$ is also larger than 1. Consequently, $\mathcal{T}_2(\frac{\delta_T}{2b_z}) \geq \mathcal{T}_5(\delta_T/2)$. Therefore, the choice of δ_T for the second term $\mathcal{T}_2$ also guarantees $\mathcal{T}_5(\delta_T/2) \leq \epsilon$.

Therefore, it suffices to ensure $\mathcal{T}_2(\frac{\delta_T}{2b_z}) \leq \epsilon$. Notice that, when $\frac{\delta_T}{2b_z} = 2w_{\max}$, then $\mathcal{T}_2(\frac{\delta_T}{2b_z}) = 0 \leq \epsilon$, so there exists δ_T such that $\mathcal{T}_2(\frac{\delta_T}{2b_z}) \leq \epsilon$.

Next, we will show that there exists such δ_T that diminishes to zero as T goes to infinity. Notice that we need

$$1 - q_w\left(\frac{a_1 \delta_T}{8b_z \sqrt{n_x}}\right) \leq \left(\frac{\epsilon}{\tilde{O}((n_x n_z)^{2.5} a_4^{n_x n_z})}\right)^{1/\lceil T/m \rceil},$$

so that

$$q_w\left(\frac{a_1 \delta_T}{8b_z \sqrt{n_x}}\right) \geq 1 - \left(\frac{\epsilon}{\tilde{O}((n_x n_z)^{2.5} a_4^{n_x n_z})}\right)^{1/\lceil T/m \rceil},$$

where the right hand side converges to zero as $T \rightarrow \infty$.

Now, consider $\delta(k) = 1/k$. Since $q_w\left(\frac{a_1 \delta(k)}{8b_z \sqrt{n_x}}\right) > 0$, there exists a large enough T_k for any $k > 0$ such that for any $T \geq T_k$, we have that

$$q_w\left(\frac{a_1 \delta(k)}{8b_z \sqrt{n_x}}\right) \geq 1 - \left(\frac{\epsilon}{\tilde{O}((n_x n_z)^{2.5} a_4^{n_x n_z})}\right)^{1/\lceil T_k/m \rceil}.$$

Furthermore, for any $T > 0$, we can define

$$\delta_T = \begin{cases} \delta(k), & \text{if } T_k \leq T < T_{k+1}, \text{ for } k > 0, \\ 2w_{\max}, & \text{if } T < T_1. \end{cases}$$

In this way, δ_T satisfies $\mathcal{T}_2(\frac{\delta_T}{2b_z}) \leq \epsilon$ and $\delta_T \rightarrow 0$ as $T \rightarrow +\infty$.

Finally, using the proof of Corollary 8.9, we can show that there exists $\frac{\delta_T}{2b_z} = \tilde{O}(n_x^{1.5}n_z^2/T)$ such that $\mathcal{T}_2(\frac{\delta_T}{2b_z}) \leq \epsilon$ whenever $q_w(\delta) = O(\delta)$. This implies $\delta_T = \tilde{O}(n_x^{1.5}n_z^2/T)$ and completes the proof. $\square$

8.H.3 Proof of Theorem 8.14

We first show that the unknown θ^* is a member of USC-SME uncertainty set $\hat{\Theta}_T^{\text{ucb}}$ with high probability. By Theorem 8.12, Corollary 8.13, and the definition in (8.5), we have

$$\mathbb{P}(w_{\max} > \hat{w}_{\max}^{(T)}) = \mathbb{P}(w_{\max} - \bar{w}_{\max}^{(T)} > \delta_T) \leq 3\epsilon,$$

which implies $1 - 3\epsilon \leq \mathbb{P}(w_{\max} \leq \hat{w}_{\max}^{(T)}) \leq \mathbb{P}(\theta^* \in \hat{\Theta}_T^{\text{ucb}})$.

Next, we show that the diameter of the UCB-SME uncertainty set is controlled by δ_T with high probability. Notice that $\hat{\Theta}_T^{\text{ucb}} \subseteq \Theta_T(w_{\max} + \delta_T)$ because $\bar{w}_{\max}^{(T)} \leq w_{\max}$. Therefore, by Theorem 8.11, the following holds for any constant $r > 0$:

$$\begin{aligned} \mathbb{P}(\text{diam}(\hat{\Theta}_T^{\text{ucb}}) > r + a_5\sqrt{n_x}\delta_T) &\leq \mathbb{P}(\text{diam}(\Theta_T(w_{\max} + \delta_T)) > r + a_5\sqrt{n_x}\delta_T), \\ &\leq \mathcal{T}_1 + \mathcal{T}_2(r). \end{aligned}$$

Let $r = \delta_T$, then, using the inequality $\mathcal{T}_2(\delta_T) \leq \mathcal{T}_2(\delta_T/2b_z)$, we have that

$$\begin{aligned} \mathbb{P}(\text{diam}(\hat{\Theta}_T^{\text{ucb}}) > \delta_T + a_5\sqrt{n_x}\delta_T) &\leq \mathbb{P}(\text{diam}(\Theta_T(w_{\max} + \delta_T)) > \delta_T + a_5\sqrt{n_x}\delta_T), \\ &\leq 2\epsilon. \end{aligned}$$

Therefore, with probability $1 - 2\epsilon$, the diameter of $\hat{\Theta}_T^{\text{ucb}}$ is bounded above by

$$\text{diam}(\hat{\Theta}_T^{\text{ucb}}) \leq \delta_T + a_5\sqrt{n_x}\delta_T = O(\sqrt{n_x}\delta_T).$$

Finally, we can verify that the event $\{\text{diam}(\hat{\Theta}_T^{\text{ucb}}) \leq \delta_T + a_5\sqrt{n_x}\delta_T = O(\sqrt{n_x}\delta_T)\}$ and the event

$\{\theta^* \in \hat{\Theta}_T^{\text{ucb}}\}$ simultaneously happen with probability at least $1 - 3\epsilon$ as follows:

$$\begin{aligned}
& \mathbb{P}\left(\theta^* \notin \hat{\Theta}_T(\hat{w}_{\max}^{(T)}), \text{ or } \text{diam}(\hat{\Theta}_T(w_{\max} + \delta_T)) > \delta_T + a_5\sqrt{n_x}\delta_T\right) \\
& \leq \mathbb{P}\left(w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} \geq \delta_T/2, \text{ or } \text{diam}(\Theta_T) > \delta_T/2b_z, \right. \\
& \quad \left. \text{or } \text{diam}(\hat{\Theta}_T(w_{\max} + \delta_T)) > \delta_T + a_5\sqrt{n_x}\delta_T\right) \\
& \leq \mathbb{P}\left(w_{\max} - \max_{0 \leq t \leq T-1} \|w_t\|_{\infty} \geq \delta_T/2\right) + \mathbb{P}\left(\text{diam}(\Theta_T) > \delta_T/2b_z, \right. \\
& \quad \left. \text{or } \text{diam}(\hat{\Theta}_T(w_{\max} + \delta_T)) > \delta_T + a_5\sqrt{n_x}\delta_T\right) \\
& \leq \epsilon + \mathbb{P}(\mathcal{E}_2) + \sum_{i=1}^{v_\gamma} \mathbb{P}\left(\bigcap_k G_{i,k}(\min(\delta_T/2b_z, \delta_T))\right) \leq 3\epsilon.
\end{aligned}$$

The third inequality follows from

- the proof of Theorem 8.12 in section 8.H.1,
- Theorem 8.11,
- the fact that the probabilities $\mathbb{P}(\text{diam}(\Theta_T) > \delta_T/2b_z)$ and $\mathbb{P}(\text{diam}(\hat{\Theta}_T(w_{\max} + \delta_T)) > \delta_T + a_5\sqrt{n_x}\delta_T)$ are bounded by the same events $\mathcal{E}_2$,
- and $G_{i,k}(\delta_T), G_{i,k}(\delta_T/2b_z) \subseteq G_{i,k}(\min(\delta_T/2b_z, \delta_T))$, where $G_{i,k}(\delta)$ is defined in Lemma 8.19 as a function of δ .

This completes the proof. $\square$

8.I Simulation details and additional experiments

This section provides the details on the simulation experiments, along with some additional results. The code for replicating the presented results can be found in the github repository: <https://github.com/jy-cds/non-asymptotic-set-membership>.

8.I.1 Baseline: LSE's confidence regions

In all our experiments, we use the 90% confidence region of the LSE as the baseline uncertainty set. The diameters of LSE's confidence regions are computed by taking minimum of the formulas provided in the following two papers: Lemma E.3 in [110] and Theorem 1 in [2]. To apply Theorem 1 in [2], we used regularization parameter $\lambda = 0.1$, $\delta = 0.1$ for 90% confidence, $S = \sqrt{\text{Tr}(\theta^{*,\top}\theta^*)}$, variance proxy $L = 1$ for truncated Gaussian distribution and $L = 4/3$ for uniform distribution.

To determine the parameters in Lemma E.3 of [110], we approximately optimize the projection matrix P in Lemma E.3 as follows. First, we consider an orthogonal transformation of the empirical covariance matrix $\Lambda = \sum_{t=1}^T z_t z_t^\top$ with $\Lambda = G M G^\top$ where G is unitary. This transforms the event $\mathcal{E}$ in Lemma E.3 to $M \geq \lambda_1 P_0 + \lambda_2 (I - P_0)$, where $G P_0 G^\top = P$. We select P_0 as a block matrix $[[I_p, 0], [0, 0]]$, then optimize over the block size p in search of the tightest LSE confidence bound.

8.I.2 Figure 8.1: SME and LSE uncertainty set visualization

In this experiment, we consider $x_{t+1} = A^* x_t + B^* u_t + w_t$, where $A^* = 0.8$ and $B^* = 1$ are unknown. $w_t \sim \text{TrunGauss}(0, \sigma_w, [-w_{\max}, w_{\max}])$ is i.i.d. and $u_t \sim \text{TrunGauss}(0, \sigma_u, [-u_{\max}, u_{\max}])$ are also i.i.d generated, where $\sigma_w = \sigma_u = 0.5$, and $w_{\max} = u_{\max} = 1$. We compare SME that knows $w_{\max} = 1$ and LSE's 90% confidence region computed based on Section 8.I.1.

8.I.3 Figure 8.2

In this experiment, we consider the the linearized longitudinal flight control dynamics of Boeing 747 [58, 85] with i.i.d. bounded inputs and disturbances sampled from truncated Gaussian and uniform distribution. The dynamics is $x_{t+1} = A^* x_t + B^* u_t + w_t$ with

$$A = \begin{bmatrix} 0.99 & 0.03 & -0.02 & -0.32 \\ 0.01 & 0.47 & 4.7 & 0 \\ 0.02 & -0.06 & 0.4 & 0 \\ 0.01 & -0.04 & 0.72 & 0.99 \end{bmatrix} \quad B = \begin{bmatrix} 0.01 & 0.99 \\ -3.44 & 1.66 \\ -0.83 & 0.44 \\ -0.47 & 0.25 \end{bmatrix}.$$

Disturbances are sampled from $\text{TrunGauss}(0, I, [-w_{\max}, w_{\max}]^4)$ as well as $\text{Unif}([-w_{\max}, w_{\max}]^4)$, while control inputs are samples from $\text{TrunGauss}(0, I, [-w_{\max}, w_{\max}]^2)$ in both disturbance settings, with $w_{\max} = 2$. To compute the UCB for SME using (8.5), we heuristically define $\delta_T = \beta \frac{n_x^{1.5} \cdot n_z^2 \cdot (\max_t \|x_t\|)}{T}$, where $n_x = 4$ and $n_z = 6$ are the system dimension, while β is a tunable parameter. This definition matches the dimension and time order of the theoretical analysis in Corollary 8.13. In both experiments of Figure 8.2, we fix $\beta = 0.01$.

In Figure 2(a)-(b), we plot SME with accurate and conservative bounds of $w_{\max}$, UCB-SME, and LSE's 90% confidence regions computed by Section 8.I.1. We use 10 different seeds to generate the disturbance sequences for each plot, and use the shaded region to denote 1 standard deviation from the mean (colored lines).

8.I.4 Figure 8.3

In this experiment, we consider autonomous systems of the form $x_{t+1} = A^* x_t + w_t$, where $A^* \in \mathbb{R}^{n_x}$ is randomly sampled and its spectral radius is normalized to be 0.9. We simulate SME and LSE for $n_x = 5, 10, 15, 20, 25$. The disturbances are sampled from $\text{TrunGauss}(0, I, [-w_{\max}, w_{\max}]^{n_x})$ as

well as $\text{Unif}([-w_{\max}, w_{\max}]^{n_x})$ with $w_{\max} = 2$. This simulation is run on 10 random seeds and the total length of the simulation is set to be $T = 1000$ across all n_x experiments. The mean is plotted as solid lines and the shaded regions denote 1 standard deviation from the mean.

Though SME's theoretical bound with respect to the dimension is $\tilde{O}(n_x^{1.5}n_z^2)$ from Corollary 8.10, which is much worse than LSE's bound, it is not reflected in Figure 8.3. Therefore, it is promising that the dimension scaling in the analysis in Section 8.3 can be further tightened. We leave this for future work.

8.1.5 Figure 8.4

To illustrate the quantitative impact of using SME for adaptive tube-based robust MPC, we study tube-based robust MPC for a system $x_{t+1} = A^*x_t + B^*u_t + w_t$ with nominal system $A^* = 1.2$, $B^* = 0.9$ with an initial model uncertainty set $\Theta_0 := [1, 1.2] \times [0.9, 1.1]$. We use the basic tube-based robust MPC method [100, 84] and parameterize the control policy as $u_k = Kx_k + v_k + \eta_k$, where $K = -1$, v_k is determined by the tube-based robust MPC algorithm, and η_k is a bounded exploration injection with $\eta_k \sim \text{Unif}([-0.01, 0.01])$. The disturbance w_k has a known bound of $w_{\max} = 0.1$ and is generated to be i.i.d. $\text{Unif}([-0.1, 0.1])$. The horizon of the tube-based robust MPC is set to be 5. The state and input constraints are such that $x_k \in [-10, 10]$ and $u_k \in [-10, 10]$ for all $k \geq 0$. We consider the task of constrained LQ tracking problem with a time-varying cost function $c_t := (x_t - g_t)^\top Q(x_t - g_t) + u_t^\top Ru_t$ where the target trajectory is generated as $g_t = 8 \sin(t/20)$.

We compare the performance of an adaptive tube-based robust MPC controller that uses the SME for uncertainty set estimation against one that uses the LSE 90% confidence region (LSE). For better visualization of the trajectory difference as a result of different estimation methods, we used the minimum of the the dominant factors in Dean, Mania, Matni, Recht, and Tu [27, equation C.12] and the LSE 90% confidence region for the LSE uncertainty set. We also plot the offline optimal RMPC controller, i.e., the controller that has knowledge of the true underlying system parameters (OPT).

Since the controller has to robustly satisfy constraints against the worst-case model in the uncertainty set, smaller uncertainty set for the tube-based robust MPC means more optimal trajectories can be computed. This observation is consistent with the extensive empirical results in the control literature [78, 80, 57].

*Chapter 9***CONVEX CONSTRAINED CONTROLLER SYNTHESIS FOR EVOLUTION
EQUATIONS**

L. Conger, A. P. Leeman, and F. Hoffmann. Convex Constrained Controller Synthesis for Evolution Equations. In *2025 American Control Conference (ACC)*, pages 1169–1176, July 2025. DOI: 10.23919/ACC63710.2025.11108072. ISSN: 2378-5861

ABSTRACT

We propose a convex controller synthesis framework for a large class of constrained linear systems, including those described by (deterministic and stochastic) partial differential equations and integral equations, commonly used in fluid dynamics, thermo-mechanical systems, quantum control, or transportation networks. Most existing control techniques rely on a (finite-dimensional) discrete description of the system, via ordinary differential equations. Here, we work instead with more general (infinite-dimensional) Hilbert spaces. This enables the discretization to be applied after the optimization (optimize-then-discretize). Using output-feedback System Level Synthesis (SLS), we formulate the controller synthesis as a convex optimization problem. Structural constraints like sensor and communication delays, and locality constraints, are incorporated while preserving convexity, allowing parallel implementation and extending key SLS properties to infinite dimensions. The proposed approach and its benefits are demonstrated on a linear Boltzmann equation.

9.1 Introduction

Many dynamical systems are described by evolution equations and have inherently continuous state spaces, such as fluid dynamics and plasma behavior in fusion reactors [20], heat transfer [82], aerospace applications [49], and bacterial movement [95]. For control of wind farms [106], considering the wind velocity as the state, one can maximize the power. In these setting, computing controllers is particularly challenging due to the infinite-dimensional nature of the state. One set of current control techniques largely relies on a discretization of the system, resulting in an ordinary differential equation (ODE) [44]. Although this traditional technique has proven useful for many evolution systems, the discretization generates a large number of states, and leads to an optimization problem often too costly to solve. A second set of methods to solve evolution equation control




Figure 9.1: Traditional control approaches discretize PDE dynamics and then solve a finite-dimensional optimal control problem (discretize-then-optimize). We propose optimizing in the infinite-dimensional space, and then discretizing the solution (optimize-then-discretize). This requires an infinite-dimensional SLS framework.

problems includes techniques used for boundary control [113]. Recent progress in this area leverages *backstepping*, in which a change of variables allows for a closed-form solution of the boundary controller [8], e.g., for specific dynamics [117] such as the reaction-diffusion equation [9, 109, 122]. Beyond boundary, or system-specific control, a third, older set of results surrounds more general control in Hilbert spaces. While classical work exists on adaptive control [130], robust control [123], controllability [112], and stabilizing operators [40], practical tools are still needed to bridge the gap between theory and applications.

To address some of the challenges in these three sets of approaches to PDE control, we present a convex controller synthesis method for a large class of linear systems, including PDEs and integral

equations, using ideas from system level synthesis (SLS), a popular robust and distributed control tool [128, 22, 7, 61, 28], and generalize results from [47, 12] by formulating solutions in the weak sense (broadening the classes of evolution equations and solutions) and numerically considering the setting of non-spatially invariant systems. SLS was recently developed to parameterize state- and output-feedback controllers as a convex optimization problem, allowing structural constraints while preserving convexity [7, 128, 25]. Our method mitigates discretization issues by optimizing directly on infinite-dimensional Hilbert spaces; in general, optimizing and discretizing do not necessarily commute [72], and instead of optimizing over a large discrete state space, we use an optimize-then-discretize approach (see Figure 9.1). This requires conceptualization of the SLS framework in infinite dimensions, pushing beyond the current finite-dimensional theory and serving both more accurate theoretical analysis and expected practical performance gains.

Contributions: We propose a convex (output-feedback) controller parametrization for deterministic or stochastic evolution equations in infinite-dimensional spaces, incorporating convex structural constraints such as locality, sensor delay, and communication delays. We prove that optimizing over convex closed-loop maps is equivalent to optimizing over linear feedback controllers, which may not be convex. On a numerical implementation using an integral equation, we show how orthonormal basis functions reduce the problem to a finite-dimensional optimization of coefficients, with subproblems solved in parallel. This work lays the foundation for several natural generalizations, including time-varying operators in the dynamics and SLS for continuous-time evolution equations.

9.2 Problem Setup

9.2.1 Notation

Let $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denote the inner product in a real Hilbert space $\mathcal{H}$ and its topological dual $\mathcal{H}^* := \{f : \mathcal{H} \rightarrow \mathbb{R} \text{ continuous, linear functions}\}$. For a linear bounded operator $A : \mathcal{D}(A) \subseteq \mathcal{H}_1 \rightarrow \mathcal{H}_2$ between two Hilbert spaces with domain $\mathcal{D}(A)$ and range $\mathcal{R}(A)$, its dual operator $A^* : \mathcal{D}(A^*) \subseteq \mathcal{H}_2^* \rightarrow \mathcal{H}_1^*$ maps $f(\cdot) \in \mathcal{H}_2^*$ to the function $g(\cdot) = f(A\cdot) \in \mathcal{H}_1^*$. The Riesz map allows us to identify A^* with the adjoint $A' : \mathcal{D}(A') \subseteq \mathcal{H}_2 \rightarrow \mathcal{H}_1$ of A , which is defined via the relation $\langle x_2, Ax_1 \rangle_{\mathcal{H}_2} = \langle A'x_2, x_1 \rangle_{\mathcal{H}_1}$ for all $x_1 \in \mathcal{D}(A) \subseteq \mathcal{H}_1, x_2 \in \mathcal{D}(A') \subseteq \mathcal{H}_2$. Thanks to the Riesz map, $\mathcal{H}$ is isometrically isomorphic to $\mathcal{H}^*$. From now on, we identify $\mathcal{H}^*$ with $\mathcal{H}$ and A' with A^* . We remark that the operator domain is part of the definition of the operator itself, and therefore A and $\mathcal{D}(A)$ are always stated together.

Throughout, we will use that the state $x^t \in \mathcal{X}$, input $u^t \in \mathcal{U}$, output $y^t \in \mathcal{Y}$, disturbance $w_x^t \in \mathcal{X}$, and observation noise $w_y^t \in \mathcal{Y}$ at time t are in Hilbert spaces. For a Hilbert space $\mathcal{X}$ and discrete times $t \in \{0, \dots, T\}$, we use the script $\mathcal{X} := \mathcal{X}^{\otimes(T+1)}$ to denote the $T + 1$ dimensional vector of elements $x^0, \dots, x^T \in \mathcal{X}$. This is again a Hilbert space when endowed with the inner product

$\langle x, \bar{x} \rangle_{\mathcal{X}} := \sum_{t=0}^T \langle x^t, \bar{x}^t \rangle_{\mathcal{X}}$. In the context of PDEs, a very useful Hilbert space is the Sobolev space $H^k(\Omega)$ with $\Omega \subseteq \mathbb{R}^d$ denoting the class of functions whose weak derivatives exist up to order k and are in $L^2(\Omega) := \{f : \Omega \rightarrow \mathbb{R} \text{ s.t. } \int_{\Omega} |f(z)|^2 dz < \infty\}$. Finally, we denote by $\delta_{z_0}(\cdot)$ the Dirac delta at z_0 .

9.2.2 Strong and Weak Forms of the Dynamics

Consider the discrete-time dynamics

$$x^{t+1} = A^* x^t + B^* u^t + w_x^t, \quad (9.1a)$$

$$y^t = C^* x^t + w_y^t. \quad (9.1b)$$

with operators

$$\text{Dynamics} \quad A^* : \mathcal{D}(A^*) \rightarrow \mathcal{X}, \quad \mathcal{D}(A^*) \subseteq \mathcal{X},$$

$$\text{Control} \quad B^* : \mathcal{D}(B^*) \rightarrow \mathcal{X}, \quad \mathcal{D}(B^*) \subseteq \mathcal{U},$$

$$\text{Observation} \quad C^* : \mathcal{D}(C^*) \rightarrow \mathcal{Y}, \quad \mathcal{D}(C^*) \subseteq \mathcal{X}.$$

The Hilbert spaces $\mathcal{X}, \mathcal{U}, \mathcal{Y}$ may be finite- or infinite-dimensional¹. In infinite dimensions, a large class of PDEs and integral equations can be expressed as (9.1) for suitable choices of operators A^*, B^* and C^* . If $\{x^t\}_{t=0}^T = \{x^t(\cdot)\}_{t=0}^T$ and $\{w_x^t\}_{t=0}^T = \{w_x^t(\cdot)\}_{t=0}^T$ are functions over some space Ω , Equation (9.1a) is said to hold *strongly*, i.e. $x^{t+1}(z) = (A^* x^t)(z) + (B^* u^t)(z) + w_x^t(z)$ pointwise for all $z \in \Omega$. However, this notion of solution to (9.1) may be too restrictive in many application settings. For instance, one may want to consider cases where (9.1a) is the time-distretization of an evolution equation

$$\partial_t x(t, z) = L^* x(t, z) + B^* u(t, z) + w_x(t, z), \quad (9.2)$$

for some differential operator $L^* : \mathcal{D}(L^*) \subseteq \mathcal{X} \rightarrow \mathcal{X}$ (for Forward Euler discretization with time step Δt , one would recover (9.1a) by setting $A^* = I + \Delta t L^*$). If $(L^* x)(z) := v \cdot \nabla_z x(z)$ is the advection operator for some fixed velocity field $v \in \mathbb{R}^d$ and $B^* = 0, w_x = 0$, then a step function at the origin, $x^0(z) := \text{sign}(z) + 1$, is propagated through the dynamics and remains a step function at all times. This solution, which is physically relevant, does not solve (9.2) strongly since a step function is not differentiable in the classical sense. Or, consider for instance dynamics L^* such that the solution $x(t, z)$ develops shocks or discontinuities over time, loosing differentiability. To give a sense to 'solving (9.1)' in these cases, we use a more general notion of solution: $(x^t, u^t)_{t=0}^T$ solves (9.1) *weakly* if for all $f \in \mathcal{D}(A) \cap \mathcal{D}(B) \subseteq \mathcal{X}$ and $g \in \mathcal{D}(C) \subseteq \mathcal{Y}$,

$$\langle x^{t+1}, f \rangle_{\mathcal{X}} = \langle x^t, A f \rangle_{\mathcal{X}} + \langle u^t, B f \rangle_{\mathcal{U}} + \langle w_x^t, f \rangle_{\mathcal{X}}, \quad (9.3a)$$

$$\langle y^t, g \rangle_{\mathcal{Y}} = \langle x^t, C g \rangle_{\mathcal{X}} + \langle w_y^t, g \rangle_{\mathcal{Y}}. \quad (9.3b)$$

¹An alternative is to use general dual spaces, going beyond Hilbert spaces.

Here, f, g are considered test functions, and applying differential operators to test functions instead of the solution itself allows for less regular solutions, see Example 9.2. We will use the weak notion of solution throughout this paper, and all pointwise statements are thus to be understood in the almost everywhere sense.

Example 9.1. *In finite dimensions, the weak and strong forms are equivalent, so our framework reduces to the standard SLS setting. When $\mathcal{X} = \mathbb{R}^{n_x}$, $\mathcal{U} = \mathbb{R}^{n_u}$, and $\mathcal{Y} = \mathbb{R}^{n_y}$, the weak form of the dynamics (9.3) reduces to the standard (strong form) finite-dimensional control setting, since $\mathcal{X}$, $\mathcal{U}$ and $\mathcal{Y}$ are Hilbert spaces when equipped with the 2-norm. Observe that the domains of $A^* \in \mathbb{R}^{n_x \times n_x}$, $B^* \in \mathbb{R}^{n_x \times n_u}$, and $C^* \in \mathbb{R}^{n_y \times n_x}$ are the entire spaces $\mathcal{X}$ and $\mathcal{U}$, which is not necessarily true in the infinite dimensional setting.*

Example 9.2. *Consider $\mathcal{X} = \mathcal{U} = L^2(\mathbb{R})$, and the dynamics*

$$x^{t+1}(z) = \partial_z x^t(z) + \int b(z - z') u^t(z') dz',$$

for $z \in \mathbb{R}$, $b \in L^2(\mathbb{R})$ with $b(-z) = b(z)$. This corresponds to the strong form of the dynamics (9.1a) with $A^* = \partial_z$ the differentiation operator and B^* the convolution operator. Here, $\mathcal{D}(A^*) = H^1(\mathbb{R}) := \{f \in L^2(\mathbb{R}) \text{ s.t. } \int_{\mathbb{R}} |\partial_z f(z)|^2 dz < \infty\}$ and $B^* u^t = b * u^t$ with $\mathcal{D}(B^*) = \mathcal{U}$. The weak form (9.3a) of the dynamics is

$$\langle x^{t+1}, f \rangle_{\mathcal{X}} = -\langle x^t, \partial_z f \rangle_{\mathcal{X}} + \langle u^t, b * f \rangle_{\mathcal{U}},$$

for all $f \in \mathcal{D}(A) = H^1(\mathbb{R})$, where $A = -\partial_z$ and $Bf = b * f$. Explicitly writing the L^2 inner products, the above means

$$\int x^{t+1}(z) f(z) dz = - \int x^t(z) \partial_z f(z) dz + \iint u^t(z) b(z - z') f(z') dz dz', \quad \forall f \in H^1(\mathbb{R}).$$

The strong form requires $x^t \in H^1(\mathbb{R})$ to be differentiable, whereas for the weak form $x^t \in L^2(\mathbb{R})$ is sufficient, shifting regularity requirements to the test functions $f \in H^1(\mathbb{R})$ instead. In this way the weak form gives sense to a non-differentiable function x^t solving a differential equation.

9.2.3 System Level Synthesis

Finite-dimensional SLS can be used to optimize a linear controller while convexly enforcing locality, communication delay and speed constraints [7], which is not possible when optimizing over the (linear) controller gains directly. In particular, locality constraints allow for large-scale problems to be solved in parallel, with controllers using only local information, leading to a scalable controller implementation. We leverage SLS to generalize the convex controller synthesis to infinite dimensional settings, including deterministic and stochastic PDE dynamics and integral equations.

In the weak sense the control input u^t is given via a family of operators $K_x^{t,\tau} : \mathcal{D}(K_x^{t,\tau}) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ and $K_y^{t,\tau} : \mathcal{D}(K_y^{t,\tau}) \subseteq \mathcal{U} \rightarrow \mathcal{Y}$ for $t, \tau \in \{0, \dots, T\}$ such that for all test functions $h \in \cap_{\tau=0}^t \mathcal{D}(K_{x,y}^{t,t-\tau}) \subseteq \mathcal{U}$,

$$\begin{aligned} \langle u^t, h \rangle_{\mathcal{U}} &= \sum_{\tau=0}^t \langle x^\tau, K_x^{t,t-\tau} h \rangle_{\mathcal{X}} \quad \text{for state feedback,} \\ \langle u^t, h \rangle_{\mathcal{U}} &= \sum_{\tau=0}^t \langle y^\tau, K_y^{t,t-\tau} h \rangle_{\mathcal{Y}} \quad \text{for output feedback.} \end{aligned}$$

We will present results for both state feedback (SF) and output feedback (OF) settings. The controller is $u^t = \sum_{\tau=0}^t (K_x^{t,t-\tau})^* x^\tau$ in the SF setting, and $u^t = \sum_{\tau=0}^t (K_y^{t,t-\tau})^* y^\tau$ in the OF setting.

Remark 9.3. Taking operator adjoints means switching the spaces for domains and ranges, so generalizing SLS to operators in Hilbert spaces switches the conventional source and target spaces from existing SLS theory. For example, if $B : \mathcal{X} \rightarrow \mathcal{U}$ then $B^* : \mathcal{U} \rightarrow \mathcal{X}$. For details, see [108].

9.3 Results

In this section, we show how to convexly synthesize a linear controller for the dynamics (9.3). Throughout, we consider a time horizon $T \in \mathbb{N}$ and denote

$$x := \begin{bmatrix} x^0 \\ \vdots \\ x^T \end{bmatrix} \in \mathcal{X}, \quad w_x := \begin{bmatrix} x^0 \\ w^0 \\ \vdots \\ w^{T-1} \end{bmatrix} \in \mathcal{X}, \quad u = \begin{bmatrix} u^0 \\ \vdots \\ u^T \end{bmatrix} \in \mathcal{U},$$

and operator matrices $\mathcal{A} : \mathcal{D}(\mathcal{A}) \subseteq \mathcal{X} \rightarrow \mathcal{X}$, $\mathcal{K}_x : \mathcal{D}(\mathcal{K}_x) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ given by

$$\mathcal{A} := \begin{bmatrix} 0 & A & & \\ & \ddots & \ddots & \\ & & 0 & A \\ & & & 0 \end{bmatrix}, \quad \mathcal{K}_x := \begin{bmatrix} K_x^{0,0} & \dots & K_x^{T,T} \\ & \ddots & \vdots \\ & & K_x^{T,0} \end{bmatrix}, \quad (9.4)$$

with $\mathcal{B} : \mathcal{D}(\mathcal{B}) \subseteq \mathcal{X} \rightarrow \mathcal{U}$ defined similarly to $\mathcal{A}$, $\mathcal{K}_y : \mathcal{D}(\mathcal{K}_y) \subseteq \mathcal{U} \rightarrow \mathcal{Y}$ defined similarly to $\mathcal{K}_x$, and $\mathcal{C} : \mathcal{D}(\mathcal{C}) \subseteq \mathcal{Y} \rightarrow \mathcal{X}$ given by $\mathcal{C} := \text{blkdiag}(C)$, that is, the matrix C repeated on the block diagonals; all operator matrices are of size $(T+1) \times (T+1)$. Denote $\bar{\mathcal{D}} := \mathcal{D}(\mathcal{A}) \cap \mathcal{D}(\mathcal{B}) = \mathcal{X} \times (\mathcal{D}(A) \cap \mathcal{D}(B))^{\otimes T} \subseteq \mathcal{X}$. Using the full time horizon, the dynamics (9.3) can be written, for all $f \in \bar{\mathcal{D}}$ and $g \in \mathcal{D}(C)$,

$$\langle x, f \rangle_{\mathcal{X}} = \langle x, \mathcal{A}f \rangle_{\mathcal{X}} + \langle u, \mathcal{B}f \rangle_{\mathcal{U}} + \langle w_x, f \rangle_{\mathcal{X}}, \quad (9.5a)$$

$$\langle y, g \rangle_{\mathcal{Y}} = \langle x, \mathcal{C}g \rangle_{\mathcal{X}} + \langle w_y, g \rangle_{\mathcal{Y}}. \quad (9.5b)$$

The relationship between u and x is given by

$$\langle u, h \rangle_{\mathcal{U}} = \langle x, \mathcal{K}_x h \rangle_{\mathcal{X}} \quad \forall h \in \mathcal{D}(\mathcal{K}_x) \text{ for SF,} \quad (9.6a)$$

$$\langle u, h \rangle_{\mathcal{U}} = \langle y, \mathcal{K}_y h \rangle_{\mathcal{Y}} \quad \forall h \in \mathcal{D}(\mathcal{K}_y) \text{ for OF.} \quad (9.6b)$$

The matrix structure in (9.4) arises by writing the strong form of the dynamics (9.1) in vector notation, and then taking adjoints of the resulting operators with respect to $\langle \cdot, \cdot \rangle_{\mathcal{X}, \mathcal{U}, \mathcal{Y}}$. Note that the controller being causal means $\mathcal{K}_x^*$ and $\mathcal{K}_y^*$ are lower block-triangular, and hence $\mathcal{K}_x$ and $\mathcal{K}_y$ are upper block-triangular.

9.3.1 State Feedback

We define closed-loop operators $\theta_x : \mathcal{D}(\theta_x) \subseteq \mathcal{X} \rightarrow \mathcal{X}$ and $\theta_u : \mathcal{D}(\theta_u) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ to map between the disturbance and the state and input respectively, in block-operator form

$$\theta_x := \begin{bmatrix} \theta_x^{0,0} & \theta_x^{1,1} & \dots & \theta_x^{T,T} \\ & \theta_x^{1,0} & & \theta_x^{T,T-1} \\ & & \ddots & \vdots \\ & & & \theta_x^{T,0} \end{bmatrix}, \quad (9.7)$$

where $\theta_x^{t,\tau} : \mathcal{D}(\theta_x^{t,\tau}) \subseteq \mathcal{X} \rightarrow \mathcal{X}$ and $\theta_u^{t,\tau} : \mathcal{D}(\theta_u^{t,\tau}) \subseteq \mathcal{U} \rightarrow \mathcal{X}$ are defined similarly. The adjoints θ_x^*, θ_u^* of the CLMs θ_x and θ_u are lower block-triangular to ensure causality, hence θ_x, θ_u are upper block-triangular, and they parameterize the trajectories $x \in \mathcal{X}$ and $u \in \mathcal{U}$ via

$$\langle x, f \rangle_{\mathcal{X}} = \langle w_x, \theta_x f \rangle_{\mathcal{X}} \quad \forall f \in \mathcal{D}(\theta_x) \subseteq \mathcal{X}, \quad (9.8a)$$

$$\langle u, h \rangle_{\mathcal{U}} = \langle w_x, \theta_u h \rangle_{\mathcal{X}} \quad \forall h \in \mathcal{D}(\theta_u) \subseteq \mathcal{U}. \quad (9.8b)$$

In the following theorem, we state one of the main result of this paper: the relationship between the CLMs θ_x, θ_u and the SF controller $\mathcal{K}_x$ is given by

$$\langle f, \theta_x^{-1} \theta_u h \rangle_{\mathcal{X}} = \langle f, \mathcal{K}_x h \rangle_{\mathcal{X}}, \quad (9.9)$$

for appropriate test functions f, h . The SLP in Theorem 9.4 parameterizes all possible causal θ_x, θ_u .

Theorem 9.4 (SLP-SF). *Fix disturbance function realization $w_x \in \mathcal{X}$ and operators $\mathcal{A}, \mathcal{B}$.*

(I) *If $\mathcal{K}_x$ of form (9.4) is given such that $\mathcal{R}(\mathcal{A}), \mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\mathcal{K}_x)$ and $\mathcal{R}(\mathcal{K}_x) \subseteq \bar{\mathcal{D}}$, then any trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ satisfying the closed-loop dynamics (9.5a)-(9.6a) also satisfies (9.8) with some causal CLMs θ_x, θ_u satisfying $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\theta_x)$, $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_u)$ and*

$$\langle f, \theta_x \hat{f} \rangle_{\mathcal{X}} = \langle f, \theta_x \mathcal{A} \hat{f} \rangle_{\mathcal{X}} + \langle f, \theta_u \mathcal{B} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}}, \quad \forall f \in \mathcal{X}, \quad \forall \hat{f} \in \bar{\mathcal{D}}. \quad (9.10)$$

(II) Let $\theta_x : \mathcal{D}(\theta_x) \rightarrow \mathcal{X}$, $\theta_u : \mathcal{D}(\theta_u) \rightarrow \mathcal{X}$ be arbitrary CLMs satisfying (9.10) such that θ_x, θ_u are upper block-triangular (implying that θ_x is invertible), $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\theta_x)$, $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_u)$ and $\mathcal{R}(\theta_u) \subseteq \mathcal{D}(\theta_x^{-1})$. Then the corresponding trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ computed with (9.8) also satisfies (9.5a), (9.6a) with the controller $\mathcal{K}_x$ defined by $\mathcal{K}_x := \theta_x^{-1}\theta_u$ and $\mathcal{D}(\mathcal{K}_x) := \mathcal{D}(\theta_u)$.

9.3.2 Output Feedback

In the output feedback setting, the controller does not have access to full state information, only the observation data y . The CLM operator matrices

$$\begin{aligned} \theta_{xx} : \mathcal{D}(\theta_{xx}) \subseteq \mathcal{X} &\rightarrow \mathcal{X}, & \theta_{ux} : \mathcal{D}(\theta_{ux}) \subseteq \mathcal{U} &\rightarrow \mathcal{X}, \\ \theta_{xy} : \mathcal{D}(\theta_{xy}) \subseteq \mathcal{X} &\rightarrow \mathcal{Y}, & \theta_{uy} : \mathcal{D}(\theta_{uy}) \subseteq \mathcal{U} &\rightarrow \mathcal{Y}, \end{aligned}$$

are defined as block-triangular operator matrices with the same structure (9.7) as θ_x and θ_u in the state feedback setting. As shown in the following theorem, the trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ is parameterized by

$$\langle x, f \rangle_{\mathcal{X}} = \langle w_x, \theta_{xx}f \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy}f \rangle_{\mathcal{Y}}, \quad (9.11a)$$

$$\langle u, h \rangle_{\mathcal{U}} = \langle w_x, \theta_{ux}h \rangle_{\mathcal{X}} + \langle w_y, \theta_{uy}h \rangle_{\mathcal{Y}}, \quad (9.11b)$$

$$\forall f \in \mathcal{D}(\theta_{xx}) \cap \mathcal{D}(\theta_{xy}) \subseteq \mathcal{X}, \quad \forall h \in \mathcal{D}(\theta_{ux}) \cap \mathcal{D}(\theta_{uy}) \subseteq \mathcal{U}.$$

We will show a relationship between the closed-loop maps and the controller $\mathcal{K}_y : \mathcal{D}(\mathcal{K}_y) \rightarrow \mathcal{Y}$ given by

$$\langle g, (\theta_{uy} - \theta_{xy}\theta_{xx}^{-1}\theta_{ux})h \rangle_{\mathcal{Y}} = \langle g, \mathcal{K}_y h \rangle_{\mathcal{Y}}, \quad (9.12)$$

for appropriate test functions g, h .

Theorem 9.5 (SLP-OF). Fix disturbance function realizations $w_x \in \mathcal{X}$, $w_y \in \mathcal{Y}$ and operators $\mathcal{A}, \mathcal{B}, \mathcal{C}$ with $\mathcal{R}(\mathcal{A}) \subseteq \bar{\mathcal{D}}$ and $\mathcal{R}(\mathcal{C}) = \bar{\mathcal{D}}$.

(I) If $\mathcal{K}_y$ of form (9.4) is given such that $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\mathcal{K}_y)$ and $\mathcal{R}(\mathcal{K}_y) \subseteq \mathcal{D}(\mathcal{C})$, then any trajectory $(x, u) \in \mathcal{X} \times \mathcal{U}$ satisfying the closed-loop dynamics (9.5a), (9.5b), (9.6b) also satisfies (9.11) with some CLMs $(\theta_{xx}, \theta_{xy}, \theta_{ux}, \theta_{uy})$ satisfying

$$\begin{aligned} \mathcal{D}(\theta_{xx}) = \mathcal{D}(\theta_{xy}) = \bar{\mathcal{D}}, & \quad \mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy}) = \mathcal{D}(\mathcal{K}_y), \\ \mathcal{R}(\theta_{xx}) = \mathcal{R}(\theta_{ux}) = \bar{\mathcal{D}}, & \quad \mathcal{R}(\theta_{xy}) = \mathcal{R}(\theta_{uy}) = \mathcal{R}(\mathcal{K}_y), \end{aligned}$$

and given by

$$\langle f, \theta_{xx}\hat{f} \rangle_{\mathcal{X}} = \langle f, \theta_{xx}\mathcal{A}\hat{f} \rangle_{\mathcal{X}} + \langle f, \theta_{ux}\mathcal{B}\hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}}, \quad (9.13a)$$

$$\langle g, \theta_{xy} \hat{f} \rangle_{\mathcal{Y}} = \langle g, \theta_{xy} \mathcal{A} \hat{f} \rangle_{\mathcal{Y}} + \langle g, \theta_{uy} \mathcal{B} \hat{f} \rangle_{\mathcal{Y}}, \quad (9.13b)$$

$$\langle f, \theta_{xx} \hat{f} \rangle_{\mathcal{X}} = \langle f, \mathcal{A} \theta_{xx} \hat{f} \rangle_{\mathcal{X}} + \langle f, C \theta_{xy} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}}, \quad (9.13c)$$

$$\langle f, \theta_{ux} h \rangle_{\mathcal{X}} = \langle f, \mathcal{A} \theta_{ux} h \rangle_{\mathcal{X}} + \langle f, C \theta_{uy} h \rangle_{\mathcal{X}}, \quad (9.13d)$$

$$\forall f \in \mathcal{X}, g \in \mathcal{Y}, \hat{f} \in \bar{\mathcal{D}}, h \in \mathcal{D}(\theta_{ux}).$$

(II) Let $(\theta_{xx}, \theta_{xy}, \theta_{ux}, \theta_{uy})$ be arbitrary CLMs satisfying (9.13) with causal upper block-triangular structure (implying that θ_{xx} is invertible), with

$$\mathcal{D}(\theta_{xx}) = \mathcal{D}(\theta_{xy}) = \bar{\mathcal{D}}, \quad \mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy}),$$

$$\mathcal{R}(\theta_{xx}) = \mathcal{R}(\theta_{ux}) = \bar{\mathcal{D}}, \quad \mathcal{R}(\theta_{xy}) = \mathcal{R}(\theta_{uy}) = \mathcal{D}(C).$$

Then the corresponding trajectory $(x, u) \in X \times \mathcal{U}$ computed with (9.11) and output $y \in \mathcal{Y}$ computed with (9.6b) also satisfy (9.5), with the controller $\mathcal{K}_y$ defined by

$$\mathcal{K}_y := \theta_{uy} - \theta_{xy}(\theta_{xx}^{-1})\theta_{ux},$$

and $\mathcal{D}(\mathcal{K}_y) := \mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy})$; it follows that $\mathcal{R}(\mathcal{K}_y) = \mathcal{D}(C)$.

9.3.3 Optimization and Structural Constraints

Classical control problems aim to construct controllers that minimize some cost, usually in terms of the state x and input u . In the SLS framework, it was shown in [7, Section 2.2.2] that optimizing over a weighted Frobenius norm of (finite-dimensional) θ_x and θ_u solves the linear quadratic regulator (LQR) problem. Other common objectives such as $\mathcal{H}_\infty$ and L_1 can also be posed in terms of closed-loop maps (see, e.g., [22, 7, 28] for SLS-based robustness guarantees). In the same vein, we propose optimizing over (infinite-dimensional) θ_x and θ_u instead of the state and input. In the state feedback setting, the (infinite-dimensional) convex optimization problem given a cost J is

$$\min_{\theta_x, \theta_u} J(\theta_x, \theta_u) \quad \text{such that (9.10) holds,}$$

and in the output feedback setting,

$$\min_{\theta_{xx}, \theta_{ux}, \theta_{uy}, \theta_{xy}} J(\theta_{xx}, \theta_{ux}, \theta_{uy}, \theta_{xy}) \quad \text{such that (9.13) holds.}$$

Let θ be the set of closed-loop maps, with $\theta = (\theta_x, \theta_u)$ in the state feedback setting and $\theta = (\theta_{xx}, \theta_{ux}, \theta_{uy}, \theta_{xy})$ in the output feedback setting. A key feature of the original finite-dimensional SLS formulation [7] is its ability to impose structural constraints such as communication delay, sensor delay and locality constraints by constraining the support of θ , while maintaining the convexity of the optimization problem. This proposed infinite-dimensional formulation has the same properties. When restricting to a constraint set S , the optimization problem then becomes

$$\min_{\theta \in S} J(\theta) \quad \text{s.t. (9.10) or (9.13) holds.}$$

The following example shows how to construct such constraint sets S .

Example 9.6 (Integral Dynamics with Constraints). *Let $\Omega \subset \mathbb{R}^2$ be a compact set. Consider the system with full state measurements $x^t \in L^2(\Omega)$, $u^t \in \mathbb{R}^{n_u}$, parameterized by $a \in L^2(\Omega)$ and $b : \Omega \rightarrow \mathbb{R}^{n_u}$,*

$$\begin{aligned} A^* x^t(\cdot) &= (a * x^t)(\cdot) \in L^2(\Omega), \quad \forall x^t \in L^2(\Omega), \\ B^* u^t &= \sum_{l=1}^{n_u} b(l, z) u_l^t \in L^2(\Omega), \quad \forall u^t \in \mathbb{R}^{n_u}. \end{aligned}$$

This means (9.3) is a simplified scattering equation, also known as linear Boltzmann, and is used to model phenomena like bacterial movement [95, Sec 5.6]. The functions a and b have compact support; for a , this specifies that $\mathcal{A}$ is local, i.e., the state at z affects only nearby states at the next time step. For b , the control affects states around the location of each actuator, where the location of the i^{th} of n_u actuators is located at $\hat{z}^{(i)} \in \mathbb{R}^2$. Assume $\text{supp } a(\cdot) = \{z \in \Omega \mid \|z\|_2 \leq r\}$ and $\text{supp } b(l, \cdot) = \{z \in \Omega \mid \|z - \hat{z}^{(l)}\|_2 \leq r\}$, where r is sufficiently small. For this problem set-up, we optimize over causal operators $\theta^* = (\theta_x^*, \theta_u^*)$ that can be expressed via kernels: they are upper block-triangular as in (9.7), and for $f \in L^2(\Omega)$ and $z \in \Omega$,

$$\begin{aligned} ((\theta_x^{t,\tau})^* f)(z) &= \int_{\Omega} \vartheta_x^{t,\tau}(\tilde{z}, z) f(\tilde{z}) d\tilde{z}, \\ ((\theta_u^{t,\tau})^* f)_l &= \int \vartheta_u^{t,\tau}(\tilde{z}, \hat{z}^{(l)}) f(\tilde{z}) d\tilde{z}, \end{aligned}$$

with kernels $\vartheta_x^{t,\tau}(\tilde{z}, z) \in L^2(\Omega \times \Omega)$, $\vartheta_u^{t,\tau}(\tilde{z}, z) \in L^2(\Omega, C(\Omega))$. To impose locality constraints on θ_x^* and θ_u^* , we fix the constraint set S to operators with restricted support of their corresponding kernels: $\text{supp } \vartheta_x^{t,\tau} = \text{supp } \vartheta_u^{t,\tau} = \{\tilde{z}, z \in \Omega \mid \|\tilde{z} - z\|_2 \leq r\}$.

Note that in the general linear operator setting, the system is often not controllable due to under-actuation. This makes it more difficult to ensure that the dynamics and the support constraints are feasible, i.e., there exists some θ^* that satisfies both. An interesting question for future work is rigorously defining notions of controllability on subspaces.

9.3.4 Parallel Solving for Integral Transforms

Consider the state space $\mathcal{X} = L^2(\Omega)$ for some $\Omega \subset \mathbb{R}^d$, and consider CLMs with the following integral transform structure: for $f \in \mathcal{X}$,

$$\begin{aligned} (\theta_x^{t,\tau})^* f &= \int_{\Omega} \vartheta_x^{t,\tau}(\tilde{z}, \cdot) f(\tilde{z}) d\tilde{z}, \\ (\theta_u^{t,\tau})^* f &= T_u \circ \int_{\Omega} \vartheta_u^{t,\tau}(\tilde{z}, \cdot) f(\tilde{z}) d\tilde{z}, \end{aligned}$$

where $T_u : \mathcal{V} \rightarrow \mathcal{U}$ for a suitable functional space $\mathcal{V}$ over Ω (not necessarily a Hilbert space), and kernels $\vartheta_x^{t,\tau} \in L^2(\Omega \times \Omega)$, $\vartheta_u^{t,\tau} \in L^2(\Omega, \mathcal{V})$. We call $\mathcal{M}$ the set of all such admissible causal kernels $\vartheta = (\vartheta_x^{t,\tau}, \vartheta_u^{t,\tau})$, and define for fixed $\tilde{z} \in \Omega$ the set $\mathcal{M}_1 := \{\vartheta \mid \vartheta_x^{t,\tau}(\tilde{z}, \cdot) \in \mathcal{X}, \vartheta_u^{t,\tau}(\tilde{z}, \cdot) \in \mathcal{V}\}$ and for fixed $z \in \Omega$ the set $\mathcal{M}_2 := \{\vartheta \mid \vartheta_x^{t,\tau}(\cdot, z), \vartheta_u^{t,\tau}(\cdot, z) \in \mathcal{X}\}$. Optimizing over CLMs is equivalent to optimizing over kernels $\vartheta \in \mathcal{M}$. When there exists a feasible set S of locality constraints, it is possible to solve this optimization problem in parallel.

Consider cost functions $J_i : \mathcal{M}_i \rightarrow \mathbb{R} \cup \{+\infty\}$ and $J : \mathcal{M} \rightarrow \mathbb{R} \cup \{+\infty\}$ for $i = \{1, 2\}$. Denote the constraints by P , where P collects the controller parametrization (9.10) or (9.13), and the structural constraints S .

Definition 9.7 (Disturbance-Parallelizable). *J and P are disturbance-parallelizable if there exists an increasing function φ_1 and norm $\|\cdot\|_1$ such that $J(\vartheta) = \varphi_1(\|\psi_1\|_1)$ for $\psi_1(\tilde{z}) := J_1(\vartheta(\tilde{z}, \cdot))$, and*

$$\bar{\vartheta} = \arg \min_{\vartheta(\tilde{z}, \cdot) \in \mathcal{M}_1} J_1(\vartheta(\tilde{z}, \cdot)) \quad \text{s.t. } P(\vartheta(\tilde{z}, \cdot)) = 0 \quad \forall \tilde{z},$$

results in the same minimizer as the full problem, that is,

$$\bar{\vartheta} = \arg \min_{\vartheta \in \mathcal{M}} J(\vartheta) \quad \text{s.t. } P(\vartheta) = 0.$$

Definition 9.8 (xuy -Parallelizable). *J and P are xuy -parallelizable if there exists an increasing function φ_2 and norm $\|\cdot\|_2$ such that $J(\vartheta) = \varphi_2(\|\psi_2\|_2)$ for $\psi_2(z) := J_2(\vartheta(z, \cdot))$, and*

$$\bar{\vartheta} = \arg \min_{\vartheta(\cdot, z) \in \mathcal{M}_2} J_2(\vartheta(\cdot, z)) \quad \text{s.t. } P(\vartheta(\cdot, z)) = 0 \quad \forall z,$$

results in the same minimizer as the full problem, that is,

$$\bar{\vartheta} = \arg \min_{\vartheta \in \mathcal{M}} J(\vartheta) \quad \text{s.t. } P(\vartheta) = 0.$$

Disturbance-parallelizable is analogous to column-separable in finite dimensional SLS; xuy -parallelizable is analogous to row-separable in finite dimensional SLS.

Remark 9.9. *If the system is parallelizable or separable into disturbance- and xuy -parallelizable components, results from [15], and finite-dimensional SLS [7, 6] suggest Alternating Direction Method of Multipliers (ADMM) as a potential distributed solver implementation.*

We will use a disturbance-parallelizable objective and constraint in the following section and illustrate results with parallel computation.

9.4 Control Implementation

In what follows, consider the dynamics from Example 9.6,

$$x^{t+1}(z) = \int a(z - z')x^t(z')dz' + \sum_{l=1}^{n_u} b(l, z)u_l^t. \quad (9.14)$$

Here, we can take $\mathcal{V} = C(\Omega)$ the continuous functions on Ω , and $T : C(\Omega) \rightarrow \mathbb{R}^{n_u}$ to be the evaluation operator at the actuators, $(Tv)_l := v(\hat{z}^{(l)})$.

From (9.10), the dynamics for the CLM kernels ϑ are

$$\begin{aligned} \vartheta_x^{t+1, \tau+1}(\tilde{z}, z) &= (A^* \vartheta_x^{t, \tau}(\tilde{z}, \cdot))(z) + (B^* T_u \circ \vartheta_u^{t, \tau}(\tilde{z}, \cdot))(z) \\ &= \int a(z - z') \vartheta_x^{t, \tau}(\tilde{z}, z') dz' + \sum_{l=1}^{n_u} b(l, z) \vartheta_u^{t, \tau}(\tilde{z}, \hat{z}^{(l)}), \\ \vartheta_x^{t, 0}(\tilde{z}, z) &= \delta_0(\tilde{z} - z), \end{aligned}$$

where $\tilde{z}$ is the input, z is the output. Define the cost function

$$J(\vartheta) := \sum_{t, \tau} \left(Q \iint |\vartheta_x^{t, \tau}(\tilde{z}, z)|^2 d\tilde{z} dz + R \int \|T_u \circ \vartheta_u^{t, \tau}(\tilde{z}, \cdot)\|_2^2 d\tilde{z} \right),$$

for scalars $Q > 0$ and $R \geq 0$, analogous to LQR for finite-dimensional SLS. Note that this means we seek a minimum operator norm solution for $(\vartheta_x^{t, \tau})^*$ and $(\vartheta_u^{t, \tau})^*$, and with this choice the optimization problem and constraints are disturbance-parallelizable according to Definition 9.7, that is, we can optimize for each $\tilde{z} \in \Omega$ separately.

We use real Fourier basis functions for ϑ_x , a , and b ². The basis allows us to solve the optimization problem over Fourier coefficients, rather than over function spaces, with truncation giving an approximate solution. Additionally, we constrain the support of ϑ_x and ϑ_u without affecting the convexity of the problem, as shown in Example 9.6. The constraint enforces locality on the disturbance propagation, and the controller. This locality constraint is enforced via basis functions with matching support.

9.4.1 Simulation Results

The simulation results show selecting an expressive basis and a small number of approximation functions ($k = 12$) are sufficient for the results to be accurate within 0.3%. The time horizon is $T = 5$ with $n_u = 16$ controllers over a domain $[-2, 2]$ in both directions. The controller topology and state response to a disturbance at position $(-0.26, 0.56)$ for the first two time steps are shown in Figure 9.2. The state is hit with a disturbance at time $t = 0$, and the nearest controller responds at time $t = 1$. By the second time step, the state is nearly back to zero.

²For details about the choice of basis, and the CVXPY [31] quadratic program implementation of constraint (9.13) for the dynamics (9.14), see [Appendix 9.B](#) and <https://github.com/LEConger/SLS-evolution-equations>.




Figure 9.2: Optimal state response (continuous green shade) and control magnitude (red dot) after $t = 1$, and $t = 2$.

Accuracy of basis functions: To determine the accuracy of the basis functions, the state sequence given the controllers resulting from the basis function coefficient optimization was computed using numerical integration and compared with θ_x . The accuracy at each time step is summarized in Table 9.1, with at most 0.23% error from the exact solution.

Table 9.1: Relative error and performance gain at each time step.

time step	error (%)	perf. gain (%)
1	0.16	42.23
2	0.11	61.23
3	0.17	70.93
4	0.21	77.39
5	0.23	82.16

Comparison with uncontrolled system: To determine how much the controllers influence the state of the system, we compare the optimal state response with its uncontrolled counterpart, i.e., $\theta_u = 0$. The average 2-norm of the state at each time step is given in Table 9.1. Compared with the state response with no control, the controller causes the state norm to decrease, e.g., by 42.23% at $t = 1$.

Parallel control with constraints: To illustrate how to compute controllers locally for large-scale systems, we increase the state domain and controllers, and we constrain the controllers to respond only to disturbances within a 1-norm radius of length 2. The SLS optimization problem is solved in parallel for each of the 15 disturbances. For controllers responding to multiple disturbances, the linearity of the CLMs allows us to add the maps together after synthesis. See Figure 9.3 to observe the locality constraint held for each disturbance. As is the case for discrete-space SLS [136], the disturbance-separability allows us to compute responses to multiple disturbances in parallel.

Benefit of continuous space control: To highlight the benefits of our approach, we compare its




Figure 9.3: The state and control responses to each of the 15 disturbances are computed in parallel, and then summed to obtain the final control values. The closed-loop maps are constrained to have support locally around each disturbance.

performance against a classic discretize-then-optimize SLS formulation. After spatially discretizing the dynamics (9.14) for different discretization steps dx , the optimal discrete controller computed with finite-dimensional SLS [7] was integrated exactly in the dynamics. In Table 9.2, finite-dimensional SLS and continuous space SLS are compared at each time step against the exact response of a zeros input controller.

In particular, our approach leads to higher performance for this numerical example, compared to its finite-dimensional counterpart [7] for all considered discretizations, highlighting the benefits of our approach for continuous-space systems. Importantly, tailored (finite-dimensional) SLS solvers achieve a state dimension scalability of $\mathcal{O}(n_x^3)$ [61], whereas our approach remains independent of the discretization used.

Table 9.2: Performance comparison with finite-dimensional SLS [7].

discretization step	avg perf. gain (%)	state dimension n_x
continuous (our approach)	42.79	/
$dx = 0.1$	32.26	1600
$dx = 0.2$	30.54	400
$dx = 0.25$	31.91	256
$dx = 0.5$	37.36	64

9.5 Conclusion

We proposed a new convex output-feedback controller synthesis framework for infinite-dimensional evolution equations, based on SLS. The framework allows for convex structural constraints, such

as sensor delay, communication delays, and locality constraints. Due to its similarities with finite-dimensional SLS, we expect many SLS properties to translate to this infinite-dimensional setting, e.g., efficient distributed solving, robustness guarantees, or robust constraint satisfaction.

9.5.1 Acknowledgements

The authors thank James Anderson, Samir Bennani, John C. Doyle and Dennice Gayme for valuable discussions.

Appendix

9.A Theorem Proofs

Proof of Theorem 9.4. Proof of (I): Using the weak form of the dynamics (9.5a) and relation (9.6a), x satisfies

$$\begin{aligned} \langle x, f \rangle_{\mathcal{X}} &= \langle x, (\mathcal{A} + \mathcal{K}_x \mathcal{B})f \rangle_{\mathcal{X}} + \langle w_x, f \rangle_{\mathcal{X}} \quad \forall f \in \bar{\mathcal{D}} \\ \Rightarrow \quad \langle x, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})f \rangle_{\mathcal{X}} &= \langle w_x, f \rangle_{\mathcal{X}} \quad \forall f \in \bar{\mathcal{D}}. \end{aligned}$$

The operator matrix $\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B}$ has domain and range $\mathcal{D}(\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B}) = \mathcal{R}(\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B}) = \bar{\mathcal{D}}$, is upper block-triangular with identity operators on the diagonal, and therefore invertible with inverse given explicitly in terms of A, B and $K_x^{l,\tau}$. Note that also $\mathcal{D}((\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1}) = \mathcal{R}((\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1}) = \bar{\mathcal{D}}$. Denote $\hat{f} := (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})f$. Hence

$$\langle x, \hat{f} \rangle_{\mathcal{X}} = \langle w_x, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} \hat{f} \rangle_{\mathcal{X}},$$

Define $\theta_x := (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1}$ with $\mathcal{D}(\theta_x) := \bar{\mathcal{D}}$ and $\theta_u := \theta_x \mathcal{K}_x$ with $\mathcal{D}(\theta_u) = \mathcal{D}(\mathcal{K}_x)$. By assumption, $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\mathcal{K}_x) \subseteq \bar{\mathcal{D}}$, and so $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{D}(\theta_x)$. Trivially, $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_u)$. Finally, θ_x and θ_u satisfy (9.10): for all $f \in \mathcal{X}$, $\hat{f} \in \bar{\mathcal{D}}$,

$$\begin{aligned} \langle f, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} \hat{f} \rangle_{\mathcal{X}} &= \langle f, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} \mathcal{A} \hat{f} \rangle_{\mathcal{X}} \\ &\quad + \langle f, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} \mathcal{K}_x \mathcal{B} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}}, \\ \Leftrightarrow \quad \langle f, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B}) \hat{f} \rangle_{\mathcal{X}} &= \langle f, \hat{f} \rangle_{\mathcal{X}}. \end{aligned}$$

Proof of (II): Since θ_x, θ_u are upper-block-diagonal, and satisfy (9.10), it follows that $\theta_x^{t,0} = \mathbf{I}$ for all $t \in \{0, \dots, T\}$; hence θ_x is invertible. Define $\mathcal{K}_x := \theta_x^{-1} \theta_u$ and $\mathcal{D}(\mathcal{K}_x) := \mathcal{D}(\theta_u)$. We begin by showing (9.6a). Indeed, using (9.8), for any $h \in \mathcal{D}(\theta_u) = \mathcal{D}(\mathcal{K}_x)$, we have

$$\langle u, h \rangle_{\mathcal{U}} = \langle w_x, \theta_u h \rangle_{\mathcal{X}} = \langle w_x, \theta_x \mathcal{K}_x h \rangle_{\mathcal{X}} = \langle x, \mathcal{K}_x h \rangle_{\mathcal{X}}.$$

To show (9.5a), recall the weak form of the SLP,

$$\langle f, \theta_x \hat{f} \rangle_{\mathcal{X}} = \langle f, \theta_x \mathcal{A} \hat{f} \rangle_{\mathcal{X}} + \langle f, \theta_u \mathcal{B} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}},$$

for all $f \in \mathcal{X}$, $\hat{f} \in \bar{\mathcal{D}}$, which can be rewritten as

$$\begin{aligned} \langle f, \theta_x(\mathbf{I} - \mathcal{A})\hat{f} \rangle_{\mathcal{X}} - \langle f, \theta_u \mathcal{B}\hat{f} \rangle_{\mathcal{X}} &= \langle f, \hat{f} \rangle_{\mathcal{X}} \\ \Rightarrow \langle f, \theta_x \tilde{f} \rangle_{\mathcal{X}} &= \langle f, (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})^{-1} \tilde{f} \rangle_{\mathcal{X}}, \end{aligned}$$

where we defined $\tilde{f} := (\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})\hat{f}$. In fact, thanks to the block-triangular structure of the operators, $\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B}$ is invertible and therefore for any $\tilde{f} \in \mathcal{X}$ there exists $\hat{f} \in \bar{\mathcal{D}}$ such that this relationship holds. Then selecting $f = w_x \in \mathcal{X}$, we have for any x solving (9.8a),

$$\langle x, \tilde{f} \rangle_{\mathcal{X}} = \langle w_x, \theta_x(\mathbf{I} - \mathcal{A} - \mathcal{K}_x \mathcal{B})\hat{f} \rangle_{\mathcal{X}} = \langle w_x, \hat{f} \rangle_{\mathcal{X}}.$$

Substituting for $\tilde{f}$ on the left-hand side and rearranging,

$$\begin{aligned} \langle x, \hat{f} \rangle_{\mathcal{X}} &= \langle x, \mathcal{A}\hat{f} \rangle_{\mathcal{X}} + \langle x, \mathcal{K}_x \mathcal{B}\hat{f} \rangle_{\mathcal{X}} + \langle w_x, \hat{f} \rangle_{\mathcal{X}} \\ &= \langle x, \mathcal{A}\hat{f} \rangle_{\mathcal{X}} + \langle u, \mathcal{B}\hat{f} \rangle_{\mathcal{U}} + \langle w_x, \hat{f} \rangle_{\mathcal{X}} + \left(\langle x, \mathcal{K}_x \mathcal{B}\hat{f} \rangle_{\mathcal{X}} - \langle u, \mathcal{B}\hat{f} \rangle_{\mathcal{U}} \right). \end{aligned}$$

The last term in brackets vanishes by (9.6a), and so the desired system response is achieved. $\square$

Proof of Theorem 9.5. Proof of (I): Given a causal $\mathcal{K}_y$ and using (9.6b), the dynamics (9.5) can be rearranged: for any $f \in \bar{\mathcal{D}}$,

$$\begin{aligned} \langle x, (\mathbf{I} - \mathcal{A})f \rangle_{\mathcal{X}} &= \langle y, \mathcal{K}_y \mathcal{B}f \rangle_{\mathcal{Y}} + \langle w_x, f \rangle_{\mathcal{X}} = \langle x, C\mathcal{K}_y \mathcal{B}f \rangle_{\mathcal{X}} + \langle w_y, \mathcal{K}_y \mathcal{B}f \rangle_{\mathcal{Y}} + \langle w_x, f \rangle_{\mathcal{X}} \\ \Rightarrow \langle x, (\mathbf{I} - \mathcal{A} - C\mathcal{K}_y \mathcal{B})f \rangle_{\mathcal{X}} &= \langle w_y, \mathcal{K}_y \mathcal{B}f \rangle_{\mathcal{Y}} + \langle w_x, f \rangle_{\mathcal{X}}. \end{aligned}$$

The operator matrix $\tilde{\mathcal{A}} := \mathbf{I} - \mathcal{A} - C\mathcal{K}_y \mathcal{B}$ has domain $\mathcal{D}(\tilde{\mathcal{A}}) = \bar{\mathcal{D}}$, range $\mathcal{R}(\tilde{\mathcal{A}}) = \bar{\mathcal{D}}$, and is invertible because it is upper block-triangular with identity operators on the block diagonal. Define $\hat{f} := (\mathbf{I} - \mathcal{A} - C\mathcal{K}_y \mathcal{B})f$. Then

$$\langle x, \hat{f} \rangle_{\mathcal{X}} = \langle w_y, \mathcal{K}_y \mathcal{B}\tilde{\mathcal{A}}^{-1}\hat{f} \rangle_{\mathcal{Y}} + \langle w_x, \tilde{\mathcal{A}}^{-1}\hat{f} \rangle_{\mathcal{X}}.$$

By (9.6b) and since $\mathcal{R}(\mathcal{K}_y) \subseteq \mathcal{D}(C)$, for all $h \in \mathcal{D}(\mathcal{K}_y)$,

$$\langle u, h \rangle_{\mathcal{U}} = \langle y, \mathcal{K}_y h \rangle_{\mathcal{Y}} = \langle x, C\mathcal{K}_y h \rangle_{\mathcal{X}} + \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}}.$$

Define the closed-loop maps

$$\theta_{xx} := \tilde{\mathcal{A}}^{-1}, \quad \theta_{xy} := \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1}, \quad \theta_{ux} := \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y, \quad \theta_{uy} := \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y + \mathcal{K}_y.$$

Note that all conditions on domains and ranges for these operators are satisfied thanks to the assumptions. Next, we will show that the closed-loop maps satisfy the SLP (9.13). We will check each of the four equalities. (9.13a) holds if for all $f \in \mathcal{X}$, $\hat{f} \in \bar{\mathcal{D}}$,

$$\langle f, \tilde{\mathcal{A}}^{-1}\hat{f} \rangle_{\mathcal{X}} = \langle f, \tilde{\mathcal{A}}^{-1}\mathcal{A}\hat{f} \rangle_{\mathcal{X}} + \langle f, \tilde{\mathcal{A}}^{-1}C\mathcal{K}_y \mathcal{B}\hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}},$$

which is equivalent to $\langle f, \tilde{\mathcal{A}}^{-1} \tilde{\mathcal{A}} \hat{f} \rangle_{\mathcal{X}} = \langle f, \hat{f} \rangle_{\mathcal{X}}$. Eqn (9.13b) holds if for all $g \in \mathcal{Y}$, $\hat{f} \in \bar{\mathcal{D}}$,

$$\langle g, \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} \hat{f} \rangle_{\mathcal{Y}} = \langle g, \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} \mathcal{A} \hat{f} \rangle_{\mathcal{Y}} + \langle g, (\mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} C + I) \mathcal{K}_y \mathcal{B} \hat{f} \rangle_{\mathcal{Y}} ,$$

which reduces to $\langle g, \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} \tilde{\mathcal{A}} \hat{f} \rangle_{\mathcal{Y}} = \langle g, \mathcal{K}_y \mathcal{B} \hat{f} \rangle_{\mathcal{Y}}$. (9.13c) holds if for all $f \in \mathcal{X}$, $\hat{f} \in \bar{\mathcal{D}}$,

$$\langle f, \tilde{\mathcal{A}}^{-1} \hat{f} \rangle_{\mathcal{X}} = \langle f, \mathcal{A} \tilde{\mathcal{A}}^{-1} \hat{f} \rangle_{\mathcal{X}} + \langle f, C \mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} \hat{f} \rangle_{\mathcal{X}} + \langle f, \hat{f} \rangle_{\mathcal{X}} ,$$

which reduces to $\langle f, \tilde{\mathcal{A}} \tilde{\mathcal{A}}^{-1} \hat{f} \rangle_{\mathcal{X}} = \langle f, \hat{f} \rangle_{\mathcal{X}}$. (9.13d) holds if for all $f \in \mathcal{X}$ and $h \in \mathcal{D}(\mathcal{K}_y)$,

$$\langle f, \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y h \rangle_{\mathcal{X}} = \langle f, \mathcal{A} \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y h \rangle_{\mathcal{X}} + \langle f, C (\mathcal{K}_y \mathcal{B} \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y + \mathcal{K}_y) h \rangle_{\mathcal{X}} ,$$

which reduces to $\langle f, \tilde{\mathcal{A}} \tilde{\mathcal{A}}^{-1} C \mathcal{K}_y h \rangle_{\mathcal{X}} = \langle f, C \mathcal{K}_y h \rangle_{\mathcal{X}}$.

Proof of (II): Consider $\theta = (\theta_{xx}, \theta_{xy}, \theta_{ux}, \theta_{uy})$ solving the SLP (9.13). Since all θ operators have the structure (9.7), it follows that $\theta_{xx}^{t,0} = I$ for all $t \in \{0, \dots, T\}$; hence θ_{xx} is invertible. Defining the controller $\mathcal{K}_y = \theta_{uy} - \theta_{xy}(\theta_{xx}^{-1})\theta_{ux}$ with $\mathcal{D}(\mathcal{K}_y) = \mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy})$, we will show that (x, u) defined via (9.11) and output feedback y defined via (9.6b) solve (9.5). First, note that $\mathcal{R}(C) = \mathcal{R}(\theta_{ux})$, θ_{xx}^{-1} is bijective on $\bar{\mathcal{D}}$ and $\mathcal{R}(\theta_{xy}) = \mathcal{R}(\theta_{uy}) = \mathcal{D}(C)$; hence $\mathcal{D}(C) = \mathcal{R}(\mathcal{K}_y)$. From the SLP (9.13),

$$\theta_{ux} \mathcal{B} = \theta_{xx} (I - \mathcal{A}) - I, \quad (9.15)$$

$$\theta_{uy} \mathcal{B} = \theta_{xy} (I - \mathcal{A}), \quad (9.16)$$

$$C \theta_{xy} = (I - \mathcal{A}) \theta_{xx} - I, \quad (9.17)$$

weakly hold on $\bar{\mathcal{D}}$, and

$$C \theta_{uy} = (I - \mathcal{A}) \theta_{ux}, \quad (9.18)$$

weakly on $\mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy})$. To show (9.5a), we use (9.11a) to compute for all $f \in \bar{\mathcal{D}}$,

$$\begin{aligned} \langle x, (I - \mathcal{A})f \rangle_{\mathcal{X}} - \langle u, \mathcal{B}f \rangle_{\mathcal{U}} - \langle w_x, f \rangle_{\mathcal{X}} &= \langle w_x, \theta_{xx} (I - \mathcal{A})f \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy} (I - \mathcal{A})f \rangle_{\mathcal{Y}} \\ &\quad - \langle u, \mathcal{B}f \rangle_{\mathcal{U}} - \langle w_x, f \rangle_{\mathcal{X}} . \end{aligned}$$

The right-hand side vanishes, since by (9.11b) together with (9.15) and (9.16),

$$\begin{aligned} \langle u, \mathcal{B}f \rangle_{\mathcal{U}} &= \langle w_x, \theta_{ux} \mathcal{B}f \rangle_{\mathcal{X}} + \langle w_y, \theta_{uy} \mathcal{B}f \rangle_{\mathcal{Y}} \\ &= \langle w_x, \theta_{xx} (I - \mathcal{A})f \rangle_{\mathcal{X}} - \langle w_x, f \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy} (I - \mathcal{A})f \rangle_{\mathcal{Y}} , \end{aligned}$$

using $\mathcal{R}(\mathcal{A}) \subseteq \bar{\mathcal{D}}$ and $\mathcal{R}(\mathcal{B}) \subseteq \mathcal{D}(\theta_{ux}) = \mathcal{D}(\theta_{uy})$.

To show (9.5b): Since $\mathcal{D}(C) = \mathcal{R}(\mathcal{K}_y)$, for any $g \in \mathcal{D}(C)$ there exists $h \in \mathcal{D}(\mathcal{K}_y)$ such that $g = \mathcal{K}_y h$. Hence from (9.6b) we have

$$\langle y, g \rangle_{\mathcal{Y}} - \langle x, Cg \rangle_{\mathcal{X}} - \langle w_y, g \rangle_{\mathcal{Y}} = \langle u, h \rangle_{\mathcal{U}} - \langle x, C\mathcal{K}_y h \rangle_{\mathcal{X}} - \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}}.$$

We will show that the right-hand side vanishes. Note that $C\mathcal{K}_y h = Cg \in \bar{\mathcal{D}}$ since $\mathcal{R}(C) \subseteq \bar{\mathcal{D}}$. Then from (9.11a) together with (9.17)-(9.18), it follows that

$$\begin{aligned} \langle x, C\mathcal{K}_y h \rangle_{\mathcal{X}} &= \langle w_x, \theta_{xx} C\mathcal{K}_y h \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy} C\mathcal{K}_y h \rangle_{\mathcal{Y}} \\ &= \langle w_x, \theta_{xx} C\theta_{uy} h \rangle_{\mathcal{X}} - \langle w_x, \theta_{xx} C\theta_{xy} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy} C\theta_{uy} h \rangle_{\mathcal{Y}} - \langle w_y, \theta_{xy} C\theta_{xy} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{Y}} \\ &= \langle w_x, \theta_{xx} (I - \mathcal{A}) \theta_{ux} h \rangle_{\mathcal{X}} - \langle w_x, \theta_{xx} (I - \mathcal{A}) \theta_{xx} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{X}} + \langle w_x, \theta_{ux} h \rangle_{\mathcal{X}} \\ &\quad - \langle w_y, \theta_{xy} (I - \mathcal{A}) \theta_{xx} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{Y}} + \langle w_y, \theta_{xy} (I - \mathcal{A}) \theta_{ux} h \rangle_{\mathcal{Y}} + \langle w_y, \theta_{xy} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{Y}} \\ &= \langle w_x, \theta_{ux} h \rangle_{\mathcal{X}} + \langle w_y, \theta_{xy} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{Y}}. \end{aligned}$$

Using (9.11b), we have

$$\begin{aligned} \langle u, h \rangle_{\mathcal{U}} - \langle x, C\mathcal{K}_y h \rangle_{\mathcal{X}} - \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}} &= \langle w_x, \theta_{ux} h \rangle_{\mathcal{X}} + \langle w_y, \theta_{uy} h \rangle_{\mathcal{Y}} - \langle w_x, \theta_{ux} h \rangle_{\mathcal{X}} - \langle w_y, \theta_{xy} \theta_{xx}^{-1} \theta_{ux} h \rangle_{\mathcal{Y}} - \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}} \\ &= \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}} - \langle w_y, \mathcal{K}_y h \rangle_{\mathcal{Y}} = 0. \end{aligned}$$

□

9.B Implementation Details

In this section, we provide implementation details for our numeric example, particularly regarding the basis function selection and optimization setup. Recall the dynamics from Example 9.6,

$$x^{t+1}(z) = \int a(z - z') x^t(z') dz' + \sum_{l=1}^{n_u} b(l, z) u_l^t.$$

From (9.10), the dynamics for the CLM kernels ϑ are

$$\begin{aligned} \vartheta_x^{t+1, \tau+1}(\tilde{z}, z) &= (A^* \vartheta_x^{t, \tau}(\tilde{z}, \cdot))(z) + (B^* T_u \circ \vartheta_u^{t, \tau}(\tilde{z}, \cdot))(z) \\ &= \int a(z - z') \vartheta_x^{t, \tau}(\tilde{z}, z') dz' + \sum_{l=1}^{n_u} b(l, z) \vartheta_u^{t, \tau}(\tilde{z}, \hat{z}^{(l)}), \\ \vartheta_x^{t, 0}(\tilde{z}, z) &= \delta_0(\tilde{z} - z), \end{aligned}$$

where $\tilde{z}$ is the input, z is the output. Recall also the cost function

$$J(\vartheta) := \sum_{t, \tau} \left(Q \iint |\vartheta_x^{t, \tau}(\tilde{z}, z)|^2 d\tilde{z} dz + R \int \|T_u \circ \vartheta_u^{t, \tau}(\tilde{z}, \cdot)\|_2^2 d\tilde{z} \right),$$

for scalars $Q > 0$ and $R \geq 0$, analogous to LQR for finite-dimensional SLS. We use real Fourier basis functions for ϑ_x and ϑ_u . For $z = [z_1, z_2]$,

$$\begin{aligned}\varphi_{mn1}(z) &= \sin \frac{2\pi m z_1}{\lambda_1} \sin \frac{2\pi n z_2}{\lambda_2}, & \varphi_{mn2}(z) &= \cos \frac{2\pi m z_1}{\lambda_1} \sin \frac{2\pi n z_2}{\lambda_2}, \\ \varphi_{mn3}(z) &= \sin \frac{2\pi m z_1}{\lambda_1} \cos \frac{2\pi n z_2}{\lambda_2}, & \varphi_{mn4}(z) &= \cos \frac{2\pi m z_1}{\lambda_1} \cos \frac{2\pi n z_2}{\lambda_2}.\end{aligned}$$

The approximations for a , b , and ϑ_x are

$$\begin{aligned}a(z) &\approx \sum_{m,n=0}^k \sum_{i=1}^4 a_{mni} \varphi_{mni}(z), & b(z) &\approx \sum_{m,n=0}^k \sum_{i=1}^4 \begin{bmatrix} b_{mn1} \varphi_{mni}(z) \\ \vdots \\ b_{mn4} \varphi_{mni}(z) \end{bmatrix} \in \mathbb{R}^{n_u}, \\ \vartheta_x^t(\tilde{z}, z) &\approx \sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}(\tilde{z}) \varphi_{mni}(z),\end{aligned}$$

for coefficients a_{mni} , b_{mni} , α_{mni} defined as

$$\begin{aligned}a_{mni} &= \frac{\kappa}{\lambda_1 \lambda_2} \int_{-\lambda_2/2}^{\lambda_2/2} \int_{-\lambda_1/2}^{\lambda_1/2} a(z) \varphi_{mni}(z) dz_1 dz_2, & b_{mni} &= \frac{\kappa}{\lambda_1 \lambda_2} \int_{-\lambda_2/2}^{\lambda_2/2} \int_{-\lambda_1/2}^{\lambda_1/2} b(l, z) \varphi_{mni}(z) dz_1 dz_2, \\ \kappa &= \begin{cases} 1 & \text{if } m = 0 \text{ and } n = 0 \\ 2 & \text{if } m = 0 \text{ or } n = 0 \\ 4 & \text{otherwise.} \end{cases}\end{aligned}$$

Note that this problem is input-separable (analogous to column-separable in finite dimensions) for a Frobenius norm cost function, so we can fix $\tilde{z}$ and solve for the operator in terms of z only. For implementation purposes, λ_1 and λ_2 must be at least twice the size of the domain in the x and y directions, respectively, to prevent aliasing. The function $a(z)$ is

$$a(z) = \begin{cases} \frac{1}{2} + \frac{1}{2} \cos \frac{\pi \|z\|_2}{r} & \text{if } \|z\|_2 \leq r \\ 0 & \text{otherwise.} \end{cases}$$

We numerically integrate to compute the coefficients a_{mni} . Because this takes some time, we save the coefficients so that we can load them instead of computing them in the future. Next we compute the coefficients for $b(z)$, which we select to be equal to

$$b(z) = \begin{bmatrix} -a(z - u_{loc}^{(1)}) \\ \vdots \\ -a(z - u_{loc}^{(n_u)}) \end{bmatrix},$$

where $u_{loc}^{(l)}$ is the location of the l^{th} controller. Since this function has n_u components, it is slower to compute the coefficients than for $a(z)$, so we suggest running the file `compute_b_coefficients_parallel.py` which parallelizes the computation of the coefficients. We use the system-level parameterization (SLP) to relate all of the coefficients and controller values. At time $t = 0$, the first state response is

$$\vartheta_x^0(\tilde{z}, z) = \delta(z - \tilde{z}).$$

The state response at time $t = 1$, as specified by the SLP, is

$$\vartheta_x^1(\tilde{z}, z) = \iint a(z - z') \vartheta_x^0(\tilde{z}, z') dz + b(z)^\top \vartheta_u(\tilde{z}) = a(z - \tilde{z}) + b(z)^\top \vartheta_u(\tilde{z}).$$

Since $\vartheta_x^0(\tilde{z}, z) = \delta(z - \tilde{z})$ is not a free variable, we start by parameterizing $\vartheta_z^1(\tilde{z}, z)$ in terms of $\vartheta_u^0(\tilde{z})$. We will use the basis function approximation. We use the trigonometry identities

$$\begin{aligned} \sin \frac{2\pi m}{\lambda_i} (z_i - \tilde{z}_i) &= \sin \frac{2\pi m}{\lambda_i} z_i \cos \frac{2\pi m}{\lambda_i} \tilde{z}_i - \cos \frac{2\pi m}{\lambda_i} z_i \sin \frac{2\pi m}{\lambda_i} \tilde{z}_i \\ \cos \frac{2\pi m}{\lambda_i} (z_i - \tilde{z}_i) &= \cos \frac{2\pi m}{\lambda_i} z_i \cos \frac{2\pi m}{\lambda_i} \tilde{z}_i + \sin \frac{2\pi m}{\lambda_i} z_i \sin \frac{2\pi m}{\lambda_i} \tilde{z}_i \end{aligned}$$

to express $a(z - \tilde{z})$ in terms of basis functions

$$a(z - \tilde{z}) \approx \sum_{m,n=0}^k \sum_{i=1}^4 a_{mni} \varphi_{mni}(z - \tilde{z}) = \sum_{m,n=0}^k \sum_{i=1}^4 \hat{a}_{mni}(\tilde{z}) \varphi_{mni}(z).$$

The coefficients $\hat{a}_{mni}$ are

$$\hat{a}_{mni} = \begin{cases} a_{mn1} \varphi_4(\tilde{z}) + a_{mn2} \varphi_3(\tilde{z}) + a_{mn3} \varphi_2(\tilde{z}) + a_{mn4} \varphi_1(\tilde{z}) & \text{if } i = 1 \\ -a_{mn2} \varphi_3(\tilde{z}) + a_{mn2} \varphi_4(\tilde{z}) - a_{mn3} \varphi_1(\tilde{z}) + a_{mn4} \varphi_2(\tilde{z}) & \text{if } i = 2 \\ -a_{mn3} \varphi_2(\tilde{z}) - a_{mn2} \varphi_1(\tilde{z}) + a_{mn3} \varphi_4(\tilde{z}) + a_{mn4} \varphi_3(\tilde{z}) & \text{if } i = 3 \\ a_{mn4} \varphi_1(\tilde{z}) - a_{mn2} \varphi_2(\tilde{z}) - a_{mn3} \varphi_3(\tilde{z}) + a_{mn4} \varphi_4(\tilde{z}) & \text{if } i = 4 \end{cases}.$$

Therefore $\vartheta_x^{(1)}(\tilde{z})$ is

$$\vartheta_x^1(\tilde{z}) = \sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}^{(1)}(\tilde{z}) \varphi_{mni}(z) = \sum_{m,n=0}^k \sum_{i=1}^4 \left(\hat{a}_{mni}(\tilde{z}) + \sum_{l=1}^{n_u} b_{mnil} \vartheta_u^{(0)}(\tilde{z})_l \right) \varphi_{mni}(z).$$

Since this must hold for z almost everywhere in the domain, the coefficients must satisfy

$$\alpha_{mni}^{(1)}(\tilde{z}) - \hat{a}_{mni}(\tilde{z}) - \sum_{l=1}^{n_u} b_{mnil} \vartheta_u^{(0)}(\tilde{z})_l = 0 \quad \forall m, n \in [0, \dots, k], i \in [4]. \quad (9.19)$$

For all other time steps, the following parameterization holds:

$$\begin{aligned}\vartheta_x^{(t+1)}(\tilde{z}, z) &\approx \sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}^{(t+1)}(\tilde{z}) \varphi_{mni}(z) \\ &= \iint \left(\sum_{m,n=0}^k \sum_{i=1}^4 a_{mni} \varphi_i(z - z') \right) \left(\sum_{\hat{m}, \hat{n}=0}^k \sum_{\hat{i}=1}^4 \alpha_{\hat{m}\hat{n}\hat{i}}^{(t)}(\tilde{z}) \varphi_{\hat{m}\hat{n}\hat{i}}(z') \right) dz' + b(z)^\top \vartheta_u^{(t)}(\tilde{z}).\end{aligned}$$

Since the basis functions are orthogonal for $(m, n, i) \neq (\hat{m}, \hat{n}, \hat{i})$ for the domain over which the integral is defined, only the $(m, n, i) = (\hat{m}, \hat{n}, \hat{i})$ terms are nonzero:

$$\begin{aligned}\sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}^{(t+1)}(\tilde{z}) \varphi_{mni}(z) &= \sum_{m,n=0}^k \sum_{i=1}^4 \iint a_{mni} \varphi_{mni}(z - z') \alpha_{mni}^{(t)}(\tilde{z}) \varphi_{mni}(z') dz' + b(z)^\top \vartheta_u^{(t)}(\tilde{z}) \\ &= \sum_{m,n=0}^k \sum_{i=1}^4 a_{mni} \alpha_{mni}^{(t)}(\tilde{z}) \iint \varphi_{mni}(z - z') \varphi_{mni}(z') dz' + b(z)^\top \vartheta_u^{(t)}(\tilde{z}).\end{aligned}$$

A character-building computation, which can be done with software such as Mathematica, reduces the integral

$$\sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}^{(t+1)}(\tilde{z}) \varphi_{mni}(z) = \sum_{m,n=0}^k \sum_{i=1}^4 4\lambda_1 \lambda_2 \bar{A}_{mn}[i, :] \alpha_{mn}^{(t)}(\tilde{z}) \varphi_{mni}(z) + b(z)^\top \vartheta_u^{(t)}(\tilde{z}),$$

where $\bar{A}_{mn}$, $\alpha_{mn}^{(t)}$ are given by

$$\bar{A}_{mn} = \begin{bmatrix} a_{mn4} & a_{mn3} & a_{mn2} & a_{mn1} \\ -a_{mn3} & a_{mn4} & -a_{mn1} & a_{mn2} \\ -a_{mn2} & -a_{mn1} & a_{mn4} & a_{mn3} \\ a_{mn1} & -a_{mn2} & -a_{mn3} & a_{mn4} \end{bmatrix}, \quad \alpha_{mn}^{(t)}(\tilde{z}) = \begin{bmatrix} \alpha_{mn1}^{(t)}(\tilde{z}) \\ \alpha_{mn2}^{(t)}(\tilde{z}) \\ \alpha_{mn3}^{(t)}(\tilde{z}) \\ \alpha_{mn4}^{(t)}(\tilde{z}) \end{bmatrix}.$$

Since this also must hold for z almost everywhere in the domain, the constraint can be written in terms of coefficients

$$\alpha_{mni}^{(t+1)}(\tilde{z}) - 4\lambda_1 \lambda_2 \bar{A}_{mn}[i, :] \alpha_{mn}^{(t)}(\tilde{z}) - \sum_{l=1}^{n_u} b_{mnil} \vartheta_u^{(t)}(\tilde{z})_l = 0 \quad \forall m, n \in [0, \dots, k], i \in [4], \quad (2)$$

The optimization problem we would like to solve, written in terms of the basis functions, is

$$\begin{aligned}J(\vartheta(\tilde{z})) &= \sum_{t=0}^T \left(Q \iint |\vartheta_x^{(t)}(\tilde{z}, z)|^2 d\tilde{z} + R |\vartheta_u^{(t)}(\tilde{z})|^2 \right) \\ &\approx \sum_{t=0}^T \left(Q \iint \left| \sum_{m,n=0}^k \sum_{i=1}^4 \alpha_{mni}^{(t)} \varphi_{mni}(z) \right|^2 d\tilde{z} + R |\vartheta_u^{(t)}(\tilde{z})|^2 \right),\end{aligned}$$

and since the basis functions are orthonormal, the cost can be written in terms of the coefficients $\alpha_{mni}^{(t)}$:

$$J(\vartheta_u(\tilde{z}), \alpha) \approx \sum_{t=0}^T \sum_{m,n=0}^k \sum_{i=1}^4 \left(Q(\alpha_{mni}^{(t)})^2 + R|\vartheta_u^{t,\tau}(\tilde{z})|^2 \right) .$$

The optimization problem that we solve is

$$\arg \min_{\vartheta_u(\tilde{z}), \alpha} J(\vartheta_u(\tilde{z}), \alpha)$$

such that (1), (2) hold.

PART II BIBLIOGRAPHY

- [1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Online least squares estimation with self-normalized processes: an application to bandit problems. *arXiv preprint arXiv:1102.2670*, 2011.
- [2] Y. Abbasi-Yadkori and C. Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26. JMLR Workshop and Conference Proceedings, 2011.
- [3] M. Abeille and A. Lazaric. Thompson sampling for linear-quadratic control problems. In *Artificial intelligence and statistics*, pages 1246–1254. PMLR, 2017.
- [4] V. Adetola and M. Guay. Robust adaptive mpc for constrained uncertain nonlinear systems. *International Journal of Adaptive Control and Signal Processing*, 25(2):155–167, 2011.
- [5] H. Akçay. The size of the membership-set in a probabilistic framework. *Automatica*, 40(2):253–260, 2004.
- [6] C. Amo Alonso, J. S. Li, J. Anderson, and N. Matni. Distributed and localized model-predictive control—part I: synthesis and implementation. *IEEE Trans. on Control of Network Systems*, 10(2):1058–1068, 2023. doi: 10.1109/TCNS.2022.3219770.
- [7] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni. System level synthesis. *Annual Reviews in Control*, 47:364–393, 2019.
- [8] P. Ascencio, A. Astolfi, and T. Parisini. Backstepping PDE Design: A Convex Optimization Approach. *IEEE TAC*, 63(7):1943–1958, July 2018. issn: 1558-2523. doi: 10.1109/TAC.2017.2757088.
- [9] M. Ayamou, N. Espitia, A. Polyakov, and E. Fridman. Finite-dimensional homogeneous boundary control for a 1D reaction-diffusion equation. In *Proc. 63rd IEEE CDC*, 2024.
- [10] E.-W. Bai, H. Cho, and R. Tempo. Convergence properties of the membership set. *Automatica*, 34(10):1245–1249, 1998.
- [11] E.-W. Bai, R. Tempo, and H. Cho. Membership set estimators: size, optimal inputs, complexity and relations with least squares. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 42(5):266–277, 1995.
- [12] B. Bamieh, F. Paganini, and M. Dahleh. Distributed control of spatially invariant systems. *IEEE TAC*, 47(7):1091–1107, July 2002. issn: 1558-2523. doi: 10.1109/TAC.2002.800646.
- [13] J. R. Benevides, M. A. Paiva, P. V. Simplício, R. S. Inoue, and M. H. Terra. Disturbance observer-based robust control of a quadrotor subject to parametric uncertainties and wind disturbance. *IEEE Access*, 10:7554–7565, 2022.
- [14] D. P. Bertsekas. *Control of uncertain systems with a set-membership description of the uncertainty*. PhD thesis, Massachusetts Institute of Technology, 1971.

- [15] R. I. Boş and E. R. Csetnek. ADMM for monotone operators: convergence analysis and rates. en. *Advances in Comp. Mathematics*, 45(1):327–359, 2019. ISSN: 1572-9044. DOI: 10.1007/s10444-018-9619-3.
- [16] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [17] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig. Safe learning in robotics: from learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5:411–444, 2022.
- [18] M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli. Adaptive stochastic mpc under time-varying uncertainty. *IEEE Transactions on Automatic Control*, 66(6):2840–2845, 2020.
- [19] M. C. Campi and E. Weyer. Finite sample properties of system identification methods. *IEEE Transactions on Automatic Control*, 47(8):1329–1334, 2002.
- [20] L. N. Cattafesta and M. Sheplak. Actuators for Active Flow Control. en. *Annual Review of Fluid Mechanics*, 43(1):247–272, 2011. ISSN: 0066-4189, 1545-4479. DOI: 10.1146/annurev-fluid-122109-160634.
- [21] T.-J. Chang and S. Shahrampour. Regret analysis of distributed online lqr control for unknown LTI systems. *IEEE Transactions on Automatic Control*, 69(1):667–673, 2024.
- [22] S. Chen, V. M. Preciado, M. Morari, and N. Matni. Robust model predictive control with polytopic model uncertainty through system level synthesis. *Automatica*, 162:111431, 2024.
- [23] X. Chen and E. Hazan. Black-box control for linear dynamical systems. In *Conference on Learning Theory*, pages 1114–1143. PMLR, 2021.
- [24] L. Conger, A. P. Leeman, and F. Hoffmann. Convex Constrained Controller Synthesis for Evolution Equations. In *2025 American Control Conference (ACC)*, pages 1169–1176, July 2025. DOI: 10.23919/ACC63710.2025.11108072. ISSN: 2378-5861.
- [25] L. Conger, Y. Lin, A. Wierman, and E. Mazumdar. Characterizing Controllability and Observability for Systems with Locality, Communication, and Actuation Constraints. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pages 8491–8497, Dec. 2024. DOI: 10.1109/CDC56724.2024.10886635. ISSN: 2576-2370.
- [26] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*:1–47, 2019.
- [27] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. *Advances in Neural Information Processing Systems*, 31:4188–4197, 2018.
- [28] S. Dean, N. Matni, B. Recht, and V. Ye. Robust guarantees for perception-based control. In *Proc. L4DC*, pages 350–360. PMLR, 2020.
- [29] S. Dean, S. Tu, N. Matni, and B. Recht. Safely learning to control the constrained linear quadratic regulator. In *2019 American Control Conference (ACC)*, pages 5582–5588. IEEE, 2019.

- [30] K. H. Degue and J. Le Ny. Differentially Private Interval Observer Design with Bounded Input Perturbation. In *2020 American Control Conference (ACC)*, pages 1465–1470, July 2020. DOI: 10.23919/ACC45564.2020.9147726. ISSN: 2378-5861.
- [31] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 2016.
- [32] J. Eising and J. Cortes. When sampling works in data-driven control: informativity for stabilization in continuous time. *arXiv preprint arXiv:2301.10873*, 2023.
- [33] J. Eising, S. Liu, S. Martínez, and J. Cortés. Using data informativity for online stabilization of unknown switched linear systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 8–13. IEEE, 2022.
- [34] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis. Finite-time adaptive stabilization of linear systems. *IEEE Transactions on Automatic Control*, 64(8):3498–3505, 2018.
- [35] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [36] E. Fogel and Y.-F. Huang. On the value of information in system identification—bounded noise case. *Automatica*, 18(2):229–238, 1982.
- [37] D. Foster, T. Sarkar, and A. Rakhlin. Learning nonlinear dynamical systems from a single trajectory. In *Learning for Dynamics and Control*, pages 851–861. PMLR, 2020.
- [38] Y. Gao, Y. Tang, H. Qi, and H. Yang. Closure: fast quantification of pose uncertainty sets. *arXiv preprint arXiv:2403.09990*, 2024.
- [39] D. Georges. The use of observability and controllability gramians or functions for optimal sensor and actuator location in finite-dimensional systems. In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, volume 4, 3319–3324 vol.4, Dec. 1995. DOI: 10.1109/CDC.1995.478999. ISSN: 0191-2216.
- [40] J. S. Gibson. The Riccati Integral Equations for Optimal Control Problems on Hilbert Spaces. en. *SIAM Journal on Control and Optimization*, 17(4):537–565, 1979. ISSN: 0363-0129, 1095-7138. DOI: 10.1137/0317039.
- [41] E. G. Gilbert. Controllability and Observability in Multivariable Control Systems. en. *Journal of the Society for Industrial and Applied Mathematics Series A Control*, 1(2):128–151, Jan. 1963. ISSN: 0887-4603. DOI: 10.1137/0301009.
- [42] J. P. Hespanha. *Linear systems theory*. Princeton university press, 2018.
- [43] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger. Learning-based model predictive control: toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:269–296, 2020.
- [44] M. Hinze and A. Rösch. *Discretization of Optimal Control Problems*. Springer, 2012. ISBN: 978-3-0348-0133-1. DOI: 10.1007/978-3-0348-0133-1_21.

- [45] D. Ho, H. Le, J. Doyle, and Y. Yue. Online robust control of nonlinear systems with large uncertainty. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 130, pages 3475–3483. PMLR, Apr. 2021.
- [46] D. Hsu, S. M. Kakade, and T. Zhang. Random design analysis of ridge regression. In *Conference on learning theory*, pages 9–1. JMLR Workshop and Conference Proceedings, 2012.
- [47] E. Jensen. *Topics in Optimal Distributed Control*. Thesis, University of California Santa Barbara, 2020.
- [48] L. Jinna, Z. Qingling, and L. Yuan. Controllability and observability of networked control systems with time-varying delays. *Journal of Systems Engineering and Electronics*, 20(4):800–806, Aug. 2009. ISSN: 1004-4132.
- [49] R. D. Joslin. Aircraft Laminar Flow Control. en. *Annual Review of Fluid Mechanics*, 30(1):1–29, 1998. ISSN: 0066-4189, 1545-4479. DOI: 10.1146/annurev.fluid.30.1.1.
- [50] T. Kailath. *Linear systems*. eng. Englewood Cliffs, N.J. : Prentice-Hall, 1980. ISBN: 9780135369616.
- [51] R. E. Kalman, Y. C. Ho, and K. C. Narendra. Controllability of Linear Dynamical Systems. *Contributions to Differential Equations*, 1(1), 1961.
- [52] T. Kargin, S. Lale, K. Azizzadenesheli, A. Anandkumar, and B. Hassibi. Thompson sampling achieves $\tilde{O}\sqrt{T}$ regret in linear quadratic control. In *Conference on Learning Theory*, pages 3235–3284. PMLR, 2022.
- [53] Y. Kawano and M. Cao. Design of Privacy-Preserving Dynamic Controllers. *IEEE Transactions on Automatic Control*, 65(9):3863–3878, Sept. 2020. ISSN: 1558-2523. DOI: 10.1109/TAC.2020.2994030.
- [54] Y. Kawano and M. Cao. Revisit Input Observability: A New Approach to Attack Detection and Privacy Preservation. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 7095–7100, Dec. 2018. DOI: 10.1109/CDC.2018.8618666. ISSN: 2576-2370.
- [55] W. Kitamura and Y. Fujisaki. Convergence properties of the membership set in the presence of disturbance and parameter uncertainty. *Transactions of the Society of Instrument and Control Engineers*, 39(4):382–387, 2003.
- [56] W. Kitamura, Y. Fujisaki, and E.-W. Bai. The size of the membership set in the presence of disturbance and parameter uncertainty. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 5698–5703. IEEE, 2005.
- [57] J. Köhler, E. Andina, R. Soloperto, M. A. Müller, and F. Allgöwer. Linear robust adaptive model predictive control: computational complexity and conservatism. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1383–1388. IEEE, 2019.
- [58] S. Lale, K. Azizzadenesheli, B. Hassibi, and A. Anandkumar. Reinforcement learning with fast stabilization in linear dynamical systems. In *International Conference on Artificial Intelligence and Statistics*, pages 5354–5390. PMLR, 2022.

- [59] M. Lauricella and L. Fagiano. Set membership identification of linear systems with guaranteed simulation accuracy. *IEEE Transactions on Automatic Control*, 65(12):5189–5204, 2020.
- [60] J. Le Ny. Privacy-preserving nonlinear observer design using contraction analysis. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 4499–4504, Dec. 2015. DOI: 10.1109/CDC.2015.7402922.
- [61] A. P. Leeman, J. Köhler, F. Messerer, A. Lahr, M. Diehl, and M. N. Zeilinger. Fast system level synthesis: robust model predictive control using Riccati recursions. *IFAC-PapersOnLine*, 58(18):173–180, 2024. ISSN: 2405-8963. DOI: <https://doi.org/10.1016/j.ifacol.2024.09.027>.
- [62] J. S. Li and C. Amo Alonso. Global Performance Guarantees for Localized Model Predictive Control. *IEEE Open Journal of Control Systems*, 2:325–336, 2023. ISSN: 2694-085X. DOI: 10.1109/OJCSYS.2023.3308009.
- [63] Y. Li, X. Chen, and N. Li. Online optimal control with linear dynamics and predictions: algorithms and regret analysis. In *Advances in Neural Information Processing Systems*, pages 14887–14899, 2019.
- [64] Y. Li, S. Das, and N. Li. Online optimal control with affine constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35 of number 10, pages 8527–8537, 2021.
- [65] Y. Li, S. Das, J. Shamma, and N. Li. Safe adaptive learning-based control for constrained linear quadratic regulators with regret guarantees. *arXiv preprint arXiv:2111.00411*, 2021.
- [66] Y. Li, J. A. Preiss, N. Li, Y. Lin, A. Wierman, and J. S. Shamma. Online switching control with stability and regret guarantees. In *Learning for Dynamics and Control Conference*, pages 1138–1151. PMLR, 2023.
- [67] Y. Li, Y. Tang, R. Zhang, and N. Li. Distributed reinforcement learning for decentralized linear quadratic control: a derivative-free policy optimization approach. *IEEE Transactions on Automatic Control*, 67(12):6429–6444, 2021.
- [68] Y. Li, J. Yu, L. Conger, T. Kargin, and A. Wierman. Learning the Uncertainty Sets of Linear Control Systems via Set Membership: A Non-asymptotic Analysis. en. In *Proceedings of the 41st International Conference on Machine Learning*, pages 29234–29265. PMLR, July 2024.
- [69] Y. Li, T. Zhang, S. Das, J. Shamma, and N. Li. Non-asymptotic system identification for linear systems with nonlinear policies. *arXiv preprint arXiv:2306.10369*, 2023.
- [70] Y. Lin, J. Gan, G. Qu, Y. Kanoria, and A. Wierman. Decentralized online convex optimization in networked systems. In *International Conference on Machine Learning*, pages 13356–13393. PMLR, 2022.
- [71] Y. Lin, G. Qu, L. Huang, and A. Wierman. Multi-agent reinforcement learning in stochastic networked systems. *Advances in neural information processing systems*, 34:7825–7837, 2021.

- [72] J. Liu and Z. Wang. Non-commutative discretize-then-optimize algorithms for elliptic PDE-constrained optimal control problems. *Journal of Comp. and Applied Mathematics*, 362:596–613, 2019. ISSN: 0377-0427. DOI: 10.1016/j.cam.2018.07.028.
- [73] X. Liu, Z. Yang, and L. Ying. Online nonstochastic control with adversarial and static constraints. *arXiv preprint arXiv:2302.02426*, 2023.
- [74] Y. Liu, Y. Zhao, and F. Wu. Ellipsoidal state-bounding-based set-membership estimation for linear system with unknown-but-bounded disturbances. *IET Control Theory & Applications*, 10(4):431–442, 2016.
- [75] M. M. Livstone and M. A. Dahleh. Asymptotic properties of set membership identification algorithms. *Systems & control letters*, 27(3):145–155, 1996.
- [76] L. Ljung. *System identification*. Springer, 1998.
- [77] B. T. Lopez, J.-J. E. Slotine, and J. P. How. Robust adaptive control barrier functions: an adaptive and data-driven approach to safety. *IEEE Control Systems Letters*, 5(3):1031–1036, 2020.
- [78] M. Lorenzen, M. Cannon, and F. Allgöwer. Robust mpc with recursive model update. *Automatica*, 103:461–471, 2019.
- [79] X. Lu and M. Cannon. Robust adaptive model predictive control with persistent excitation conditions. *Automatica*, 152:110959, 2023.
- [80] X. Lu, M. Cannon, and D. Koksál-Rivet. Robust adaptive model predictive control: performance and parameter estimation. *International Journal of Robust and Nonlinear Control*, 2019.
- [81] H. Mania, S. Tu, and B. Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*. Volume 32, pages 10154–10164. Curran Associates, Inc., 2019.
- [82] A. Manzoni, A. Quarteroni, and S. Salsa. Quadratic Control Problems Governed by Linear Elliptic PDEs. en. In *Optimal Control of Partial Differential Equations: Analysis, Approximation, and Applications*, pages 103–165. Springer International Publishing, 2021. ISBN: 9783030772260. DOI: 10.1007/978-3-030-77226-0_5.
- [83] N. Matni and V. Chandrasekaran. Regularization for design. In *53rd IEEE Conference on Decision and Control*, pages 1111–1118, Dec. 2014. DOI: 10.1109/CDC.2014.7039530. ISSN: 0191-2216.
- [84] D. Q. Mayne, M. M. Seron, and S. Rakovic. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.
- [85] A. Mete, R. Singh, and P. Kumar. Augmented rbmle-ucb approach for adaptive control of linear quadratic systems. *Advances in Neural Information Processing Systems*, 35:9302–9314, 2022.
- [86] M. Milanese, J. Norton, H. Piet-Lahanier, and É. Walter. *Bounding approaches to system identification*. Springer Science & Business Media, 2013.

- [87] M. Milanese and A. Vicino. Optimal estimation theory for dynamic systems with set membership uncertainty: an overview. *Automatica*, 27(6):997–1009, 1991.
- [88] Y. Nakahira, Q. Liu, N. Bernat, T. Sejnowski, and J. Doyle. Theoretical foundations for layered architectures and speed-accuracy tradeoffs in sensorimotor control. In *2019 American Control Conference (ACC)*, pages 809–814. IEEE, 2019.
- [89] K. Narendra and A. Annaswamy. Robust adaptive control in the presence of bounded disturbances. *IEEE Transactions on Automatic Control*, 31(4):306–315, 1986.
- [90] K. S. Narendra and A. M. Annaswamy. Persistent excitation in adaptive systems. *International Journal of Control*, 45(1):127–160, 1987.
- [91] C. O. Nwankpa, J. C. Jimenez, and S. Jayasuriya. Observability of network-delayed multi-converter power systems. In *2015 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 930–933, May 2015. DOI: 10.1109/ISCAS.2015.7168787. ISSN: 2158-1525.
- [92] S. Oymak and N. Ozay. Non-asymptotic identification of lti systems from a single trajectory. In *2019 American control conference (ACC)*, pages 5655–5661. IEEE, 2019.
- [93] A. Parsi, A. Iannelli, and R. S. Smith. Active exploration in adaptive model predictive control. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 6186–6191. IEEE, 2020.
- [94] A. Parsi, A. Iannelli, M. Yin, M. Khosravi, and R. S. Smith. Robust adaptive model predictive control with worst-case cost. *IFAC-PapersOnLine*, 53(2):4222–4227, 2020.
- [95] B. Perthame. *Transport Equations in Biology*. en. Frontiers in Mathematics. Birkhäuser, 2007. ISBN: 9783764378417. DOI: 10.1007/978-3-7643-7842-4.
- [96] M. Petrik and R. H. Russel. Beyond confidence regions: tight bayesian ambiguity sets for robust mdps. *Advances in neural information processing systems*, 32, 2019.
- [97] J. Qi, G. He, S. Mei, and F. Liu. Power system set membership state estimation. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–7. IEEE, 2012.
- [98] A. Rantzer. Concentration bounds for single parameter adaptive control. In *2018 Annual American Control Conference (ACC)*, pages 1862–1866. IEEE, 2018.
- [99] J. B. Rawlings and D. Q. Mayne. *Model predictive control: Theory and design*. Nob Hill Pub., 2009.
- [100] J. B. Rawlings, D. Q. Mayne, and M. Diehl. *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI, 2017.
- [101] C. Rogers. Covering a sphere with spheres. *Mathematika*, 10(2):157–164, 1963.
- [102] T. Sarkar and A. Rakhlin. Near optimal finite time identification of arbitrary linear dynamical systems. In *International Conference on Machine Learning*, pages 5610–5618. PMLR, 2019.
- [103] A. Sarker, P. Fisher, J. E. Gaudio, and A. M. Annaswamy. Accurate parameter estimation for safety-critical systems with unmodeled dynamics. *Artificial Intelligence*:103857, 2023.

- [104] A. Sasfi, M. N. Zeilinger, and J. Köhler. Robust adaptive mpc using control contraction metrics. *arXiv preprint arXiv:2209.11713*, 2022.
- [105] Y. Sattar, S. Oymak, and N. Ozay. Finite sample identification of bilinear dynamical systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 6705–6711. IEEE, 2022.
- [106] C. R. Shapiro, P. Bauweraerts, J. Meyers, C. Meneveau, and D. F. Gayme. Model-based receding horizon control of wind farms for secondary frequency regulation. *Wind Energy*, 20(7):1261–1275, 2017. doi: <https://doi.org/10.1002/we.2093>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.2093>.
- [107] S. Shin, Y. Lin, G. Qu, A. Wierman, and M. Anitescu. Near-optimal distributed linear-quadratic regulator for networked systems. *SIAM Journal on Control and Optimization*, 61(3):1113–1135, 2023.
- [108] R. E. Showalter. Hilbert Space Methods for Partial Differential Equations. en. *Electronic Journal of Diff. Equations*:01–214, 2009. issn: 1072-6691. doi: 10.58997/ejde.mon.01.
- [109] Y.-C. Si, C.-K. Xie, and N. Zhao. Boundary control for a class of reaction-diffusion systems. en. *International Journal of Automation and Computing*, 15(1):94–102, 2018. issn: 1751-8520. doi: 10.1007/s11633-016-1012-4.
- [110] M. Simchowitz and D. Foster. Naive exploration is optimal for online lqr. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020.
- [111] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht. Learning without mixing: towards a sharp analysis of linear system identification. In *Conference On Learning Theory*, pages 439–473. PMLR, 2018.
- [112] M. Slemrod. A Note on Complete Controllability and Stabilizability for Linear Control Systems in Hilbert Space. en. *SIAM Journal on Control*, 12(3):500–508, 1974. issn: 0036-1402. doi: 10.1137/0312038.
- [113] A. Smyshlyaev and M. Krstic. *Adaptive Control of Parabolic PDEs*. en. Princeton University Press, 2010. isbn: 9781400835362.
- [114] J. Taft and P. De Martini. Sensing and Measurement Architecture for Grid Modernization. Technical report, Pacific Northwest National Laboratory, 2016.
- [115] M. Tanaskovic, L. Fagiano, R. Smith, P. Goulart, and M. Morari. Adaptive model predictive control for constrained linear systems. In *2013 European Control Conference (ECC)*, pages 382–387. IEEE, 2013.
- [116] Y. Tang, J.-B. Lasserre, and H. Yang. Uncertainty quantification of set-membership estimation in control and perception: revisiting the minimum enclosing ellipsoid. *arXiv preprint arXiv:2311.15962*, 2023.
- [117] F. Tröltzsch. *Optimal control of partial differential equations: theory, methods, and applications*, volume 112. AMS, 2010.
- [118] A. Tsiamis, I. Ziemann, N. Matni, and G. J. Pappas. Statistical learning theory for control: a finite sample perspective. *arXiv preprint arXiv:2209.05423*, 2022.

- [119] S. L. Tu. *Sample complexity bounds for the linear quadratic regulator*. University of California, Berkeley, 2019.
- [120] P. Van Overschee and B. De Moor. *Subspace identification for linear systems: Theory—Implementation—Applications*. Springer Science & Business Media, 2012.
- [121] H. J. van Waarde, J. Eising, M. K. Camlibel, and H. L. Trentelman. The informativity approach to data-driven analysis and control. *arXiv preprint arXiv:2302.10488*, 2023.
- [122] R. Vazquez and M. Krstic. Boundary control and estimation of reaction-diffusion equations on the sphere under revolution symmetry conditions. en. *Int. Journal of Control*, 92(1):2–11, 2019. ISSN: 0020-7179, 1366-5820. DOI: 10.1080/00207179.2017.1286691.
- [123] S. R. Venkatesh, A. Megretski, and M. A. Dahleh. On robust control synthesis and analysis in a Hilbert space. *Systems & Control Letters*, 39(1):1–12, 2000. ISSN: 0167-6911. DOI: 10.1016/S0167-6911(99)00078-X.
- [124] J.-L. Verger-Gaugry. Covering a ball with smaller equal balls in $\mathbb{R}^n$. *Discrete and Computational Geometry*, 33:143–155, 2005.
- [125] M. Vidyasagar and R. L. Karandikar. A learning theory approach to system identification and stochastic adaptive control. *Probabilistic and randomized methods for design under uncertainty*:265–302, 2006.
- [126] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger. Data-driven safety filters: hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems. *IEEE Control Systems Magazine*, 43(5):137–177, 2023.
- [127] A. Wagenmaker and K. Jamieson. Active learning for identification of linear dynamical systems. In *Conference on Learning Theory*, pages 3487–3582. PMLR, 2020.
- [128] Y.-S. Wang, N. Matni, and J. C. Doyle. Separable and Localized System-Level Synthesis for Large-Scale Systems. *IEEE TAC*, 63(12):4234–4249, 2018. ISSN: 1558-2523. DOI: 10.1109/TAC.2018.2819246.
- [129] Z. Wang and Q. Yuan. A method for analyzing the state controllability of linear discrete time- varying time-delay systems. In *2020 IEEE 9th Data Driven Control and Learning Systems Conference (DDCLS)*, pages 1176–1181, Nov. 2020. DOI: 10.1109/DDCLS49620.2020.9275035.
- [130] J. T.-Y. Wen and M. J. Balas. Robust adaptive control in Hilbert Space. *Journal of Mathematical Analysis and Applications*, 143(1):1–26, 1989. ISSN: 0022-247X. DOI: 10.1016/0022-247X(89)90025-5.
- [131] C. Wu, K. L. Teo, and S. Wu. Min–max optimal control of linear systems with uncertainty and terminal state constraints. *Automatica*, 49(6):1809–1815, 2013.
- [132] H. Xu and Y. Li. On the convergence rates of set membership estimation of linear systems with disturbances bounded by general convex sets. *arXiv preprint arXiv:2406.00574*, 2024.

- [133] C. Yeh, J. Yu, Y. Shi, and A. Wierman. Robust online voltage control with an unknown grid topology. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, pages 240–250, 2022.
- [134] J. Yu, V. Gupta, and A. Wierman. Online stabilization of unknown linear time-varying systems. *arXiv preprint arXiv:2304.02878*, 2023.
- [135] J. Yu, D. Ho, and A. Wierman. Online Adversarial Stabilization of Unknown Networked Systems. en. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(1):1–43, Feb. 2023. ISSN: 2476-1249. DOI: 10.1145/3579452.
- [136] J. Yu, Y.-S. Wang, and J. Anderson. Localized and distributed $\mathcal{H}_2$ state feedback control. In *2021 American Control Conference (ACC)*, pages 2732–2738. IEEE, 2021.
- [137] K. Zhang, Q. Sun, and Y. Shi. Trajectory tracking control of autonomous ground vehicles using adaptive learning mpc. *IEEE Transactions on Neural Networks and Learning Systems*, 32(12):5554–5564, 2021.
- [138] R. Zhang, Y. Li, and N. Li. On the regret analysis of online lqr control with predictions. In *2021 American Control Conference (ACC)*, pages 697–703. IEEE, 2021.
- [139] X. Zhang, W. Shi, X. Li, B. Yan, A. Malkawi, and N. Li. Decentralized temperature control via HVAC systems in energy efficient buildings: an approximate solution procedure. In *Proceedings of 2016 IEEE Global Conference on Signal and Information Processing*, pages 936–940, 2016.
- [140] Y. Zhang, G. Qu, P. Xu, Y. Lin, Z. Chen, and A. Wierman. Global convergence of localized policy iteration in networked multi-agent reinforcement learning. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(1):1–51, 2023.
- [141] Y. Zhang, R. Cheng, and Y. Xia. Observability Blocking for Functional Privacy of Linear Dynamic Networks. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 7469–7474, Dec. 2023. DOI: 10.1109/CDC49753.2023.10384211. ISSN: 2576-2370.
- [142] Z. Zhao and Q. Li. Adaptive sampling methods for learning dynamical systems. In *Mathematical and Scientific Machine Learning*, pages 335–350. PMLR, 2022.
- [143] Y. Zheng and N. Li. Non-asymptotic identification of linear dynamical systems using multiple trajectories. *IEEE Control Systems Letters*, 5(5):1693–1698, 2020.
- [144] B. Zhou, Z. Wang, Y. Zhai, and H. Yuan. Data-Driven Analysis Methods for Controllability and Observability of A Class of Discrete LTI Systems with Delays. In *2018 IEEE 7th Data Driven Control and Learning Systems Conference (DDCLS)*, pages 380–384, May 2018. DOI: 10.1109/DDCLS.2018.8515909.

INDEX

A

ADMM, *see also* Alternating Direction Method of Multipliers, 46–49, 348, 407

Alternating Direction Method of Multipliers, 40

B

Bakry-Émery approach, 13, 14, 16, 99

barycentric convexity, 24, 26, 192, 196, 197, 202, 213, 218, 232

best response, 29, 31, 35, 37, 261, 263, 280, 283–285, 289–292

C

CLM, *see also* closed-loop maps, 57, 58, 403–407, 409

closed-loop maps, 42, 337, 340, 399, 403–405, 412

controllability, 40, 45, 46, 50, 51, 58, 338, 339, 341–345, 347, 349, 406

D

Danskin's theorem, 30, 31, 35, 37, 84, 109, 117, 125, 260, 261, 307, 308, 310

displacement convexity, *see also* geodesic convexity, 12, 13, 16, 19, 21, 22, 30, 34, 74, 90, 137, 144, 156, 258, 270, 273, 303, 306, 307

distribution shift, 20, 28, 29, 32, 33, 35, 68, 70–73, 80, 84, 87, 88, 250–252, 255, 263

E

EVI, *see also* evolution variational inequality, 9–11, 13, 26–28

evolution variational inequality, 9, 26, 187–189

G

geodesic convexity, 6, 12, 18, 25, 137, 146, 156, 173, 174, 210, 218, 270, 315

geodesics, 6, 10–14, 16, 25, 26, 31, 110, 144, 188, 190, 198–200, 202, 204, 208–210, 218, 232, 237, 238, 241, 242

H

HWI inequality, 14–16, 19, 277, 278, 288

K

KL divergence, *see also* Kullback-Leibler divergence, 81, 85, 86, 91, 255, 257, 258, 266, 273, 300–304, 307, 311, 314–317

Kullback-Leibler divergence, 33

L

linear quadratic regulator, 40, 43, 50, 368, 369, 405
 log Sobolev inequality, 95, 278

M

minimizing movement scheme, 10, 13, 23, 187, 236
 monotonicity, 18, 20–22, 24, 26, 27, 137–139, 148, 160, 187, 192, 193, 195–200, 217–219, 221, 236, 239–242

N

Nash equilibrium, 18, 70, 72, 76, 79, 102, 137, 138, 145, 148, 154, 187, 191, 193, 199, 217–219, 233, 237, 238, 301

O

observability, 40, 45, 46, 50, 51, 58, 337–339
 optimal transport, 21, 28, 35, 36, 93, 137, 270, 304
 OT, *see also* optimal transport, 36–38, 301, 308, 313

R

relative energy, 75, 256

S

set membership estimation, 52
 SLP, *see also* system level parameterization, 42, 43, 57, 403, 404, 411–413, 416
 SLS, *see also* System Level Synthesis, 40, 41, 43, 45, 47–51, 54–57, 337–340, 355, 358, 366, 367, 397–399, 401, 402, 405, 407–411, 415
 SME, *see also* set membership estimation, 53, 54, 353–364, 366–369, 371–373, 388, 392, 394, 395
 Stackelberg equilibrium, 31, 32
 system level parameterization, 42, 45
 System Level Synthesis, 2, 39–41, 46, 48, 49, 52, 54

T

Talagrand inequality, 15, 96, 116, 259, 273, 279
 timescale separation, 29, 31, 32, 34, 37, 72, 81–83, 255, 263, 305, 308

V

variational movement scheme, 23, 187, 191–193, 197, 201, 233, 234