Essays on Law and Economics

Thesis by

Matthew Estes

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

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended May 11, 2026

© 2026

Matthew Estes
ORCID: 0009-0004-3804-5383

All rights reserved

ACKNOWLEDGEMENTS

This dissertation benefited from multiple people whose guidance and support I greatly appreciate.

First, I would like to thank my thesis committee: Mike Alvarez, Jim Greiner, Jonathan Katz, and Bob Sherman. My academic work and career owe much to the assistance, feedback, guidance, and support of the committee members. Classes, discussions, and meetings with you all have been formative and rewarding. Committee feedback has also improved my work, highlighting substantive connections and methodological possibilities in each meeting. I am very appreciative of each member of the committee, and I look forward to continued collaboration.

Second, I extend thanks and gratitude to all who assisted with the job market. In addition to the substantial support of my thesis committee, I am fortunate to have received guidance and encouragement from numerous professors at the University of Chicago Law School. I would foremost like to thank Lisa Bernstein, William Hubbard, and Brian Leiter for their helpful advice and feedback throughout the job market season. I am also thankful to Douglas Baird, Lee Fennell, Tom Ginsburg, Jacob Goldin, and other participants at the 2025 University of Chicago Law Teaching Conference for job market advice. I am very grateful to these individuals and other past law professors, whose classes and scholarship inspire and shape my own work.

Third, I would like to thank my fellow graduate students, colleagues, and collaborators. I thank Ransi Clark and Mitchell Linegar, for being a considerable source of friendship and encouragement the last five years. I also appreciate other collaborators and friends, including Jack Adeney, Joe Clark, Alex Halsey, Kyle Nelson, Ronak Patel, Dylan Perlson, Conner Wells. Many thanks as well to the members of Mike's research group: I appreciate the suggestions and engagement in the Friday meetings from Beatrice, Danny, Jacob, Melina, Rafal, Raquel, and others.

Finally, a special thanks to my family for supporting me and these works, particularly my mother, aunt, brother, sister, niece, and grandmother.

ABSTRACT

This dissertation empirically studies problems in law and economics. Using data and econometric methods, the dissertation estimates causal effects in different legal settings. Chapter 1 studies eviction reduction policies, finding that reforms to the number of LA County eviction courts had limited effects. Chapter 2 further studies eviction court assignments using a regression discontinuity design, finding local gaps in eviction probability and testing the impact of two causal mechanisms. Chapter 3 develops a partial identification approach to regression discontinuity designs, offering local and extrapolated bounds, new estimators, simulation results, and an application to a prior property price study. Chapter 4 studies the relationship between courts and legal scholarship, finding that judicial mentions boosted citations to certain works. Overarching themes include specification robustness, partial identification, measurement, counterfactual simulations, and visualizations to demonstrate causal identification and model estimation.

PUBLISHED CONTENT AND CONTRIBUTIONS

Estes, M. (2026a). Eviction Reduction Policies: Do More Courts Reduce Defaults in LA County? [Submitted].
This paper is sole-authored by M.E. The paper is currently under review.

Estes, M., & Nelson, K. (2026). Justice Divided, Justice Denied? The Effects of Court Rules on Eviction Outcomes in Los Angeles County [Under preparation].
This paper is coauthored with Kyle Nelson. M.E. participated in every stage of this paper, including conceptualization, data collection, methodology, empirical analysis, and the writing of the manuscript.

Estes, M. (2026b). Monotone Regression Discontinuity: Partial Identification, Estimation, and Extrapolation in RDDs [Under preparation].
This paper is sole-authored by M.E.

Estes, M., & Clark, R. (2026). Courting the Academy: The Judicial Role in Popularizing Legal Scholarship [Accepted]. *Kansas Journal of Law and Public Policy*, *35*.
This paper is coauthored with Ransi Clark. The paper is forthcoming in Volume 35 of the *Kansas Journal of Law and Public Policy*. M.E. participated in every stage of this paper, including conceptualization, data collection, methodology, empirical analysis, and the writing of the manuscript.

TABLE OF CONTENTS

Acknowledgements . . . iii
Abstract . . . iv
Published Content and Contributions . . . v
Table of Contents . . . v
List of Illustrations . . . viii
List of Tables . . . xiii
Introduction . . . 1
Chapter I: Eviction Reduction Policies: Do More Courts Reduce Defaults in LA County? . . . 3
1.1 Introduction . . . 3
1.2 Eviction Reduction Strategies . . . 6
1.3 Eviction Process & Institutional Background . . . 12
1.4 Data Collection & Descriptive Evidence . . . 15
1.5 Court Expansion Study . . . 17
1.6 Discussion & Policy Implications . . . 33
Chapter II: Justice Divided, Justice Denied? The Effects of Court Rules on Eviction Outcomes in Los Angeles County . . . 44
2.1 Introduction . . . 44
2.2 Institutional Background . . . 48
2.3 Data and Outcomes . . . 54
2.4 Regression Discontinuity Design: Methods & Results . . . 57
2.5 Testing Mechanisms: Distance-to-Court and Procedures . . . 66
2.6 Conclusion . . . 73
Chapter III: Monotone Regression Discontinuity: Partial Identification, Estimation, and Extrapolation in RDDs . . . 78
3.1 Introduction . . . 78
3.2 RDD Identification: Local Treatment Effect . . . 80
3.3 Simulation Results . . . 94
3.4 Extensions: Extrapolation to ATE & Fuzzy Monotone RDD . . . 99
3.5 Application: Energy Efficiency and Property Price Discontinuities . 103
3.6 Conclusion . . . 106
Chapter IV: Courting the Academy: The Judicial Role in Popularizing Legal Scholarship . . . 109
4.1 Introduction . . . 109
4.2 Legal Theories in the Academy and in the Courts . . . 114
4.3 Data & Empirical Descriptions of the Movements . . . 118
4.4 Methodology: Bounding Causal Effects . . . 133
4.5 Results: How Courts Impact Legal Movements . . . 139
4.6 Conclusion . . . 147

Appendix A: Supplementary Material for Chapter 1 . . . 153
A.1 Types of Reforms . . . 153
A.2 Additional Eviction Courthouse Details . . . 155
A.3 Additional Result Plots . . . 156
A.4 Commute Times . . . 166
A.5 Relaxing Parallel Trends . . . 169
A.6 All Observed Evictions . . . 172
A.7 Covariate Specification Curves . . . 173
A.8 Comparative California Eviction Rules . . . 174
Appendix B: Supplementary Material for Chapter 2 . . . 180
B.1 Creating the Courthouse Assignment Map . . . 180
B.2 Docket Records (Scraping the LASC Portal, Sampling Rates, Geocoding Addresses) . . . 183
B.3 Time-to-Default Plots . . . 187
B.4 Total Case Time Plots . . . 192
B.5 Money Judgment Plots . . . 194
B.6 Data Overview . . . 198
B.7 Procedures in Docket Records . . . 199
B.8 Misfiled Cases . . . 207
B.9 Covariate Continuity Tests with `rdrobust` . . . 211
B.10 Additional Robustness Checks . . . 216
B.11 Mechanism Model . . . 221
Appendix C: Supplementary Material for Chapter 3 . . . 249
C.1 Sharp RDD Constant Treatment Effect . . . 249
C.2 Extrapolation to ATE . . . 251
C.3 Data Generating Processes (DGPs) . . . 255
C.4 Reverse Kernel Examples . . . 259
C.5 Local Simulation Study Results . . . 261
C.6 Local Simulation Study Coverage Results . . . 269
C.7 Additional Simulation Results . . . 273
C.8 Additional Application Results . . . 279
C.9 Fuzzy RDD . . . 283
Appendix D: Supplementary Material for Chapter 4 . . . 289
D.1 Additional Descriptive Plots . . . 289
D.2 Counterfactual Complexity & The Curse of Dimensionality . . . 294
D.3 Full Treatment Histories . . . 298
D.4 Counterfactual Framework and Identification . . . 301
D.5 Bounded Deviations Tuning Parameter . . . 307
D.6 Additional Results Plots . . . 308

LIST OF ILLUSTRATIONS

Number *Page*

1.1 LA County Courthouse Expansion . . . 19
1.2 Changes in Distance-to-Court from 2017 Court Expansion Policy . . 21
1.3 Checking for Discontinuous Shocks near Policy Reform . . . 22
1.4 Cohort Averages: Piecewise Fits . . . 23
1.5 Dynamic DID Estimates: Comparisons with Controls . . . 27
1.6 Dynamic DID Estimates: Above-Median Comparisons with Controls 29
1.7 Model Summary: Aggregate DID Estimates . . . 33
1.8 California County Distance-to-Court Densities . . . 36
2.1 Eviction Timeline (Simplified) . . . 52
2.2 Courthouse Assignment . . . 54
2.3 Pasadena & West Covina Map . . . 59
2.4 LATE Estimates by Courthouse Pair (`rdrobust`) . . . 61
2.5 Distribution of Mechanism Beta Estimates (Distance, Non-Standardized) 71
3.1 Sharp RDD: Continuous CEFs with Constant Treatment Effect . . . 83
3.2 Bounding the Treatment Effect at the Cutoff $C = 3$ . . . 86
3.3 Local Extremum Estimation (DGP1 Example) . . . 90
3.4 Kernel Estimation: Parabolic & Reverse Parabolic Kernels . . . 92
3.5 Data Generating Processes for Simulation (DGPs) . . . 95
3.6 Local RDD Treatment Effect Estimates by Cutoff . . . 105
4.1 Cumulative Realist Counts . . . 122
4.2 Realist & Originalist Court Mentions . . . 123
4.3 Cumulative Originalist Counts . . . 125
4.4 ChatGPT Legal Subject Classifications of Citing Articles . . . 130
4.5 Overview of Court Mentions and Academic Citations . . . 131
4.6 Bounded Deviations Example . . . 138
4.7 Individual (Unit-Level) Treatment Effects Illustration . . . 140
4.8 Cohort Average Citations and Divergence Periods . . . 142
4.9 Realists: Mid vs Never-Treated Cohorts . . . 143
4.10 Realists: Late vs Never-Treated Cohorts . . . 144
4.11 Realists: Mid vs Never-Treated Cohorts . . . 146
4.12 Originalists: Late vs Never-Treated Cohorts . . . 147

A.1 Eviction Timeline . . . 155
A.2 LA County Eviction Court Assignment . . . 155
A.3 Observed Defaults & Descriptive Findings . . . 157
A.4 Reduced Form (LA County Assessor): 5-mile Binned Averages . . . 159
A.5 Reduced Form Relationships: Default Probability & Default Counts . 159
A.6 Difference-in-Differences (DID) Method . . . 160
A.7 Reduced Form (LA City Controller Data) . . . 161
A.8 Dynamic DID Estimates (LA City Controller Addresses Only) . . . 161
A.9 Dynamic DID Estimates Pre-2020 (LA City Controller Addresses Only) . . . 162
A.10 Dynamic DID Estimates (Assessor Unit-Weighted) . . . 163
A.11 Dynamic DID Estimates Pre-2020 (Assessor Unit-Weighted) . . . 163
A.12 Dynamic DID Estimates (pre-Aug 2017 Buildings Only) . . . 164
A.13 Dynamic DID Estimates: Above-Median Comparisons (pre-Aug 2017 Buildings Only) . . . 165
A.14 Commute Time Histogram, Averages, and Correlations . . . 166
A.15 Dynamic DID Estimates: Comparisons with Controls (Commute Times) . . . 167
A.16 Dynamic DID Estimates: Above-Median Comparisons with Controls (Commute Times) . . . 168
A.17 Partial Identification in the Bounded Deviations Approach . . . 170
A.18 Relaxing Parallel Trends: Bounded Deviations . . . 171
A.19 Dynamic DID Estimates: All Observed Outcomes . . . 172
A.20 Covariate Specification Curve (Defaults Only) . . . 173
A.21 Covariate Specification Curve (All Outcomes) . . . 173
A.22 Additional Eviction Court Locations . . . 176
A.23 Commute-to-Court Times by CA County . . . 179
B.1 Sampling Rates using FOIA Data: 2016-2023 . . . 184
B.2 Defaults Geocoding & Matching Rates . . . 186
B.3 Time-to-Default: Early-Estimated Default Entered Date . . . 188
B.4 Time-to-Default: Late-Estimated Default Entered Date . . . 189
B.5 Time-to-Default: By Courthouses . . . 190
B.6 Time-to-Default: Using Earliest Case Dates . . . 191
B.7 Total Case Time (< 1 year cases only) . . . 192
B.8 Total Case Time (< 3 years cases only) . . . 193
B.9 Yearly Money Judgments . . . 194

B.10 Money Judgments by Courthouse (< $5,000 cases only) 195
B.11 Money Judgments by Courthouse (< $10,000 cases only) 196
B.12 Money Judgments by Courthouse (< $20,000 cases only) 197
B.13 Eviction Filing Volume (CA & LA County) and Disposition Times (CA) 198
B.14 PCA Procedure Index Histograms by Court (procedure all) 204
B.15 PCA Procedure Index Histograms by Court (procedure extended) . . 205
B.16 PCA Procedure Index Histograms by Court (procedure core) 206
B.17 Misfiling: Zip 90043 Example & Geolocating Check 208
B.18 Misfiled Cases: Distance to Observed vs Assigned Court 209
B.19 LATE Estimates by Courthouse Pair 217
B.20 JW Similarity Scores 218
B.21 Mechanism Simulation 223
B.22 Reduced Form Relationships Across Boundary Segments 224
B.23 Distance-to-Court and Generalized Cost Relationship 225
B.24 Generalized Costs Over Time 226
B.25 Distribution of Mechanism Beta Estimates (Distance, Standardized) . 227
B.26 Distribution of Mechanism Beta Estimates (Generalized Cost, Non-Standardized) 228
B.27 Distribution of Mechanism Beta Estimates (Generalized Cost, Standardized) 229
B.28 Distribution of Unweighted Mechanism Beta Estimates (Distance-to-Court, Non-Standardized) 231
B.29 Distribution of Unweighted Mechanism Beta Estimates (Distance-to-Court, Standardized) 232
B.30 Distribution of Unweighted Mechanism Beta Estimates (Generalized Cost, Non-Standardized) 233
B.31 Distribution of Unweighted Mechanism Beta Estimates (Generalized Cost, Standardized) 234
C.1 Bounding the ATE Under Global Monotonicity 254
C.2 Data Generating Process 1 (DGP1) 255
C.3 Data Generating Process 2 (DGP2) 256
C.4 Data Generating Process 3 (DGP3) 257
C.5 Data Generating Process 4 (DGP4) 258
C.6 Reverse Kernel Estimation 260
C.7 DGP1: Local Effect Upper Bound Simulation Results 261

C.8 DGP2: Local Effect Upper Bound Simulation Results 262
C.9 DGP3: Local Effect Upper Bound Simulation Results 263
C.10 DGP4: Local Effect Upper Bound Simulation Results 264
C.11 DGP1: Local Effect Lower Bound Simulation Results 265
C.12 DGP2: Local Effect Lower Bound Simulation Results 266
C.13 DGP3: Local Effect Lower Bound Simulation Results 267
C.14 DGP4: Local Effect Lower Bound Simulation Results 268
C.15 Upper Bound Coverage 269
C.16 Lower Bound Coverage 270
C.17 Upper Bound Coverage of True Effect 271
C.18 Lower Bound Coverage of True Effect 272
C.19 More Bandwidths: Coverage Rates 273
C.20 More Bandwidths: Partial ID Region Widths 274
C.21 Missing Data: Coverage Rates 275
C.22 Missing Data: Partial ID Region Widths 276
C.23 Measurement Error: Coverage Rates 277
C.24 Measurement Error: Partial ID Region Widths 278
C.25 Sejas-Portillo, Moro, and Stowasser 2025, Figure 3 Reproduction . . 279
C.26 Placebo Jump Distributions near E to D Cutoff 280
C.27 Placebo Jump Distributions near F to E Cutoff 281
C.28 Placebo Jump Distributions near G to F Cutoff 282
C.29 Bounding ATE in Fuzzy RDD Under Global Monotonicity 287
D.1 Cumulative Citations 289
D.2 Yearly Citations 290
D.3 Yearly CourtListener Court Mentions 291
D.4 Legal Realist Court Mentions 292
D.5 Word Clouds by Legal Movement 293
D.6 Potential Outcomes: Combinatorial Explosion 295
D.7 Legal Realists: Court Mentions 298
D.8 Originalists: Court Mentions (Part 1) 299
D.9 Originalists: Court Mentions (Part 2) 299
D.10 Originalists: Court Mentions (Part 3) 300
D.11 Originalists: Court Mentions (Part 4) 300
D.12 Bounded Deviations (Fixed $C = 1$) 303
D.13 Bounded Deviations & DID Comparison 305
D.14 ATT Bounds as a Function of C 307

D.15 Realists: Early vs Mid Cohorts 308
D.16 Realists: Mid vs Late Cohorts 309
D.17 Realists: Early vs Late Cohorts 309
D.18 Originalists: Mid vs Late Cohorts 310
D.19 Realists: Mid vs Never-Treated Cohorts 310
D.20 Realists: Mid vs Late Cohorts 311

LIST OF TABLES

Number *Page*

1.1 Zip Code Table for Eviction Cases . . . 14
1.2 Cohorts Definitions and Expectations . . . 20
2.1 Zip Code Table for Unlawful Detainer Cases . . . 53
2.2 RDD LATE Robustness Checks . . . 64
3.1 Results of Simulation Study (Upper Bound) . . . 97
3.2 Results of Simulation Study (Lower Bound) . . . 98
3.3 Local Linear RDD Price Discontinuity Estimates (Sejas-Portillo, Moro, and Stowasser (2025) Table 3) . . . 104
4.1 Legal Realists and Their Works . . . 121
4.2 Sampled Articles by Legal Category . . . 128
4.3 Articles with Category in Title or Keywords (Regex Check) . . . 129
4.4 Cohort Distribution for Realists and Originalists . . . 135
A.1 Comparison of Procedural and Structural Eviction Reforms . . . 154
A.2 Courthouses and Locations for Estimating Distance-to-Court Densities 175
A.3 Contra Costa County Filing Rule . . . 177
A.4 San Bernardino County Filing Rule . . . 178
B.1 Simplifying Assignment: Unique Courthouses . . . 181
B.2 Simplified Zip Code Table . . . 181
B.3 Procedural Events in Docket Records (Collected by Regex) . . . 199
B.4 Procedure Averages by Courthouse . . . 200
B.5 Procedure Averages by Courthouse (cont.) . . . 201
B.6 Chatsworth and Van Nuys East: Covariate Continuity Checks . . . 211
B.7 Compton and Norwalk: Covariate Continuity Checks . . . 212
B.8 Compton and Stanley Mosk: Covariate Continuity Checks . . . 212
B.9 Inglewood and Long Beach: Covariate Continuity Checks . . . 213
B.10 Pasadena and Stanley Mosk: Covariate Continuity Checks . . . 213
B.11 Pasadena and West Covina: Covariate Continuity Checks . . . 214
B.12 Santa Monica and Stanley Mosk: Covariate Continuity Checks . . . 214
B.13 McCrary / rddensity Tests by Courthouse Pair . . . 215
B.14 Additional Outcome LATE Estimates . . . 219
B.15 Additional Outcome LATE Estimates (w/Covariates) . . . 220
B.16 Chatsworth and Van Nuys East (Distance Specifications) . . . 235
B.17 Chatsworth and Van Nuys East (Cost Specifications) . . . 236

B.18 Compton and Norwalk (Distance Specifications) 237
B.19 Compton and Norwalk (Cost Specifications) 238
B.20 Compton and Stanley Mosk (Distance Specifications) 239
B.21 Compton and Stanley Mosk (Cost Specifications) 240
B.22 Inglewood and Long Beach (Distance Specifications) 241
B.23 Inglewood and Long Beach (Cost Specifications) 242
B.24 Pasadena and Stanley Mosk (Distance Specifications) 243
B.25 Pasadena and Stanley Mosk (Cost Specifications) 244
B.26 Pasadena and West Covina (Distance Specifications) 245
B.27 Pasadena and West Covina (Cost Specifications) 246
B.28 Santa Monica and Stanley Mosk (Distance Specifications) 247
B.29 Santa Monica and Stanley Mosk (Cost Specifications) 248
C.1 Linear Regression Results 249

INTRODUCTION

This dissertation empirically studies problems in law and economics. Using econometric methods, the chapters focus on estimating causal effects using simulations and real data on housing courts, property prices, and legal citations. Overarching themes include specification robustness, partial identification, measurement, and using simulations and visualizations to demonstrate counterfactual identification and model estimation. Using various methodologies (difference-in-differences, regression discontinuity, bounding), the chapters here identify the impact of legal phenomena on certain outcomes and explore the robustness of counterfactual conclusions.

Chapter 1 studies eviction reduction policies using a difference-in-differences method. Eviction is a leading cause of housing instability and eviction reduction is a policy priority. Yet many strategies to reduce evictions lack empirical footing and the effectiveness of any particular policy is likely to hinge on the mechanisms determining eviction outcomes. Empirical work quantifying how and how much different policies causally effect eviction outcomes is therefore needed. Using data on LA County eviction filings, this paper empirically examines how court-based procedural policies impact eviction. By exploiting an August 2017 change in court policy, this paper quantifies the extent to which changing the number of eviction courts causally effects default evictions via changes in tenant costs. The paper shows that the reform to the number of courts had limited effects on eviction outcomes. Aggregate treatment effect estimates are often insignificant, do not align with expectations, or imply small effect sizes across model specifications. The findings highlight the need for both short-term and long-term strategies to address eviction and affordability.

Chapter 2 further studies how court assignment influences housing outcomes using a regression discontinuity design. More than 40,000 households in Los Angeles County face eviction each year. Pursuing policies that reduce the number of evictions is of growing importance to state and local policymakers, but the causes and consequences of eviction are poorly understood. By collecting eviction docket records and linking them to administrative data, we are able to study an important institutional determinant of eviction in LA County: how courts assign cases. Because eviction cases are assigned to courthouses based on a unique spatial assignment rule, we test the effect of court assignment on default eviction probability using a regression discontinuity design. We show that courthouse assignment can increase the probability

of default eviction by 0.05–13.34 percentage points. Estimates are robust to some model specifications and sensitive to others. Finally, we implement a spatial boundary segment mechanism model to disentangle the effects of two causal mechanisms: tenant costs and court procedure. Using outcome and mechanism variation across boundary segments, we find small local effects of tenant costs and court procedure on default outcomes but also a high degree of specification uncertainty. Additional studies may more definitively isolate eviction mechanisms.

Chapter 3 develops a partial identification method to bound treatment effects in regression discontinuity designs. Unlike the standard regression discontinuity approach, the partial identification results here rely on monotonicity and boundedness assumptions rather than a continuity assumption. This paper develops two novel estimation strategies to bound the local average treatment effect in such designs where the data generating processes are monotone near the running variable. In simulation studies, these estimators perform well with good coverage properties and are robust to bandwidth selection. But there is an efficiency cost, particularly in missing data or measurement error settings. Next, the paper extends the monotone partial identification approach to fuzzy designs and to extrapolate bounds on the average treatment effect away from the cutoff. Finally, I apply Monotone RDD to a prior regression discontinuity study, investigating the effect of energy efficiency ratings on property prices (Sejas-Portillo, Moro, & Stowasser, 2025).

Chapter 4 studies the relationship between the courts and legal academia. Legal scholars propose theories that judges use to support their opinions. Soon other academics join the fray, giving their own assessments of the judicial use of the theory, raising the academic profile of the legal theory and the legal scholar. Selective in their acknowledgments, judges become not just arbiters of justice but also academic success. The judge-academic dynamic, though widely noted, is under-studied empirically. This paper closes the gap by introducing a causal framework to explore how legal theories become legal canon. Combining data on academic citations and judicial mentions, we quantify the extent to which judges impacted the popularity of two legal movements: legal realism and originalism. Exploiting timing differences in judicial mentions between legal works, we apply a conservative bounding approach to estimate the degree to which academic popularity is spurred on by judicial acknowledgment. We find that judicial mentions increased academic citations to legal realist works but not for originalist works. Because the movements are not contemporary, the findings do not show the relative judicial role net of other factors.

Chapter 1

EVICTION REDUCTION POLICIES: DO MORE COURTS REDUCE DEFAULTS IN LA COUNTY?

1.1 Introduction

The lack of affordable housing is a growing crisis across the United States. By some estimates, "tens of millions of families, across red and blue states, struggle with rent and home prices" (Dougherty, 2024). By U.S. Department of Housing and Urban Development (2024) metrics, more than half of renting households are cost-burdened: they spend more than 30% of their income on housing costs. And according to Pew Research Center (2024) polling, 69% of U.S. adults are now "very concerned" about rising housing costs.[1] Consequently, housing policy is an increasingly pressing priority for elected officials.[2]

One key consequence of the housing crisis is a growing number of evictions. Because a growing portion of household income is spent on rent, evictions are rising.[3] The repercussions for evicted tenants are significant: reduced credit access; reduced earnings; increased hospital visits; lower educational performance; and increased housing instability, including homelessness (Collinson, Humphries, Mader, et al., 2024; Desmond, 2012, 2017; Desmond & Kimbro, 2015; Lens, Nelson, & Gromis, 2020; Meyer, Wyse, & Logani, 2025). Because eviction is linked to these varied and pernicious social ills, understanding the causes and consequences of eviction is an important goal for researchers. Yet many factors impact eviction outcomes, so understanding which policies deliver temporary versus lasting relief to renters is empirically challenging. Consequently, carefully tailoring eviction and housing policy is an important but practically difficult endeavor.

The focus in this paper is assessing available interventions for effectiveness at decreasing evictions, particularly *default* evictions. Default evictions occur when

[1] The impact is felt unevenly across demographic groups: for instance, younger renters (18-29) are particularly likely (55%) to say "the availability of affordable housing is a major problem in their local community." Pew Research Center (2022).

[2] In California, for example, Governor Newsom has signed into law numerous housing reforms aimed at reducing rent and increasing housing supply (Office of Governor Gavin Newsom, 2024). See also Office of Mayor Karen Bass (2022).

[3] The most common reason for eviction is non-payment of rent (Desmond, 2017; Los Angeles City Controller's Office, 2025).

tenants do not show up to housing court and landlords win *by default*. Default is not easily reversed and accelerates the process of lock out. Although some level of default is inevitable in any court system, eliminating unnecessary or unfair defaults is a key part of the access-to-justice literature (Greiner & Matthews, 2015). Optimally tailoring the default rate matters for efficiency and welfare reasons. Zero default may strain court resources with little upside if benefits accrue only to tenants with a persistent inability-to-pay (Abramson, 2024). On the other hand, landlords assured default wins may hastily (or improperly) evict tenants with legal or practical grounds for staying in the rental unit.[4] In equilibrium, eviction protections tend to increase tenant rents, further complicating welfare analysis.[5]

These issues therefore necessitate empirical analysis to understand how tenants, landlords, and courts respond to changes in the housing policy apparatus. This paper begins by mapping the empirical terrain, reviewing legal and economic studies on the effectiveness of various interventions at reducing eviction and defaults in theory and in practice. There are three intervention approaches with prior empirical support: (1) representation-based approaches focused on providing tenants with legal resources at or before eviction court, (2) tenant cost approaches focused on reducing the cost of using or getting to court, and (3) market-based approaches aimed at increasing housing supply and reducing rents.

This paper next empirically studies how reforms to part of the eviction machine impact default evictions via the tenant cost channel. LA County is the largest local trial court system in the U.S. and is acutely affected by housing and eviction issues.[6] But reforms to LA's housing and eviction rules are understudied. Using a staggered difference-in-differences (DID) design, I estimate how expanding the number of LA County eviction courts impacted eviction outcomes. The eviction court expansion enables testing how tenant cost changes causally effect LA County evictions over time.

[4] If landlords face a legal rule which assures them of a tenant default, they may evict paying tenants for discriminatory or otherwise illegal reasons. See *infra* (discussing composition effects from procedural reforms). Similarly, landlords may be too quick to evict if tenants experience only temporary income shocks. In this case, society may have practical reasons to give tenants who are likely-to-repay a chance to do so before starting the eviction process.

[5] The impact across renters is largely due to an *ex ante* screening problem (Abramson, 2024): landlords do not know with certainty what "type" of tenant you are, so they increase rents to account for the possibility (and added cost) of evicting types that are unable-to-pay rent. In partial equilibrium, policy changes which impact the default probability may affect the total number of evictions in unexpected ways. See *infra* Section 1.2.

[6] See, e.g., Lanzas (2024): "Los Angeles, the nation's second-largest city, consistently ranks amongst the least affordable housing markets and has one of the highest eviction rates."

Specifically, the empirical study focuses on a court policy shock in LA County beginning in August 2017 which expanded the number of eviction courts from 8 to 11. The court expansion shocked tenants costs (measured by distance-to-court or commute times) for some but not all tenants. The policy shock resulted in some tenants facing increased costs (longer distance-to-court), some facing decreased costs (shorter distance-to-court), and some unaffected (no change in distance-to-court). Comparing these cohorts, I find mixed results from the reform. Although before-after comparisons show the observed number of defaults increased for those with increased costs and decreased for those with decreased costs, the DID analysis cannot rule out null effects under many model specifications. Additional robustness checks suggest that, in the aggregate, estimates are either insignificant, do not align with expectations, or are small relative to outcome variation.

The takeaway is that default judgments are persistent and resistant to policy reforms. Additional legal resources (more courts, more judges, more lawyers) or procedural reforms are not guaranteed to meaningfully halt broader legal and socioeconomic trends, at least in the case of eviction default and relatively modest legal reforms. This is part of an emerging set of findings that procedural legal interventions that increase legal resources, reduce litigant costs, or otherwise alter the form of the proceeding may have only limited observed effects in the aggregate.

This paper also compares eviction institutions and filing rules across jurisdictions. Using Eviction Lab data on California eviction filings, I estimate the distance-to-court distributions for tenants in several California counties. Using these estimates, I find that there are many counties where average tenant costs of getting to court (in the form of distance-to-court) are higher than LA County. The comparative institutional exercise suggests that, although court expansion had small effects on LA County default evictions, reforms in other California counties may be more impactful. This may especially be the case in counties served by only one eviction courthouse.

This paper is also part of a series of works to determine the relative effectiveness of procedural versus structural housing reforms in combating eviction.[7] The overarching research goal is understanding how different levers (e.g., court-based procedural

[7]Because tenants often face a large resource disadvantage in eviction cases, procedural reforms aim to "level the playing field" for tenants in cases against better-resourced landlords. By contrast, because non-payment of rent is the primary reason for eviction, structural reforms tend to focus on reducing rent and housing prices by increasing the housing supply. See Appendix A.1 for further discussion. See also Teresa et al. (2025) describing upstream and downstream eviction prevention policies.

policies and market-based structural reforms) can reduce evictions and housing instability for vulnerable populations. This paper contributes to the procedural half of the broader scholarly agenda to understand how to reduce evictions. But eviction scholars may need to reorient focus away from procedures and toward the structure of rental markets. In other words, legal procedure may simply patch fundamental and long-term problems in housing supply.

The remainder of the paper is structured as follows. Section 1.2 reviews the empirical and theoretical findings from the law, economics, policy, sociology, and urban planning literatures, focusing on empirical studies of methods to reduce evictions generally and defaults specifically. Section 1.3 reviews the eviction process and court assignment in LA County, describing the court assignment mechanism which is key to the empirical exercise. Next, in Section 1.4, I briefly describe the collected data qualitatively and quantitatively. Section 1.5 is the heart of the paper: I empirically test whether changes in court assignment policy reduced default evictions. Finally, I discuss implications for how policies may or may not reduce evictions in Section 1.6.

1.2 Eviction Reduction Strategies

This section briefly reviews the main findings from a diverse, interdisciplinary eviction literature. I begin with the access-to-justice problem posed by default eviction judgments and possible solutions suggested by the literature. I then turn to detailing the emerging economic picture of the costs and benefits of procedural reforms: although strengthening eviction protections for tenants may give them time and/or improve case outcomes, the economic costs of eviction protections in equilibrium are less discussed. Finally, I note the extensive descriptive literature on eviction correlates, which sheds light on how eviction disproportionately affects vulnerable communities.

The Problem of Default & Possible Solutions

From an access-to-justice perspective, default judgments are a continuing problem across the law. Default judgments occur when plaintiffs obtain judgments against defendants because defendants do not contest the case. Because unrepresented defendants often default at higher rates, "such behavior presents an access to justice problem in that low- and middle-income people are benefiting from—and are in fact disproportionately suffering adverse consequences from—the formal administration of justice." (Greiner & Matthews, 2015). Defaults are often considered, then, to be

a legal harm in themselves (Greiner & Matthews, 2015):

> "From a systemic point of view, court actors and judges consider high default rates a public harm. The administration of justice suffers when parties do not meet in an adversarial proceeding in order to resolve claims."

Understanding the legal and non-legal factors that determine default is important to craft policy solutions. Although other procedural interventions, including mediation and/or technology policy, merit further empirical study,[8] I focus here on three legal factors inspired by the literature that partially explain the eviction and default phenomena: (1) lack of resources (esp. legal representation), (2) tenant costs of getting to court, and (3) market structure (esp. rental prices).

Representation-Based Interventions (Civil *Gideon*)

Likely the most common policy suggestion to assist tenants in eviction court is to provide some form of legal representation. Because tenants are often unrepresented and eviction is a summary proceeding, some scholars believe evictions are "patently unfair" or even "violent" (Scherer, 2022, 2023). Empirical scholars, accordingly, have studied whether some forms of legal representation can remedy the disparities between landlord and tenant in eviction proceedings.

Empirical studies have used a variety of methods from different legal settings across the U.S. For example, Greiner, Pattanayak, and Hennessy (2012) study the effects of legal representation interventions on eviction outcomes in a randomized controlled trial (RCT). These studies find mixed results. In Greiner, Pattanayak, and Hennessy (2012), the authors find:

> "no statistically significant evidence that [the offer] of a traditional attorney-client relationship, as compared to a referral to the [lawyer for the day] program, had a large (or any) effect on the likelihood that the occupant would retain possession; on the financial consequences of

[8] For example, it is possible that mediation would "help parties reach better outcomes in eviction cases," at least for smaller, non-corporate landlords (Armstrong & Ryan Jr., 2024; Bieretz, Burrowes, & Bramhall, 2020). Another possibility is that technology (artificial intelligence, in particular) will level the informational playing field in eviction and reduce defaults. But see Simshaw (2022) for AI skepticism. The effectiveness of such possibilities may turn on what mechanisms most impact default and eviction outcomes. See also Elliott, 2025 for a discussion of eviction mediation.

> the dispute; on the judicial involvement in or attention to the litigation cases; or on any other outcome."

In the other study (Greiner, Pattanayak, & Hennessy, 2013), the authors find that full legal representation following an initial intake and screening process mattered: 34% of treated occupants (who received full legal representation from a legal services staff attorney) versus 62% of control occupants (who received an instructional clinic only) lost possession of their units. Observational studies on the effectiveness of NYC's access-to-counsel program are similar, finding that tenants are less likely to face adverse eviction judgments, especially those at the highest risk of possessory judgment (Cassidy & Currie, 2023). Other studies find high tenant demand for lawyers in eviction cases, perhaps especially for out-of-court reasons (Caspi & Rafkin, 2023).

Still, there are limits to the effectiveness of legal representation. For example, some eviction defenses may be unlikely to benefit tenants in practice.[9] And the direct costs can be large: in 2024, NYC spent around $179.5 million on tenant legal services, including the Universal Access to Counsel (UAC) program.[10]

Tenant Costs & Spatial Considerations

Another factor from the literature that may cause defaults are spatial costs (see, e.g., Clair, Orozco, and Zhang (2025)). As Prescott (2017) puts it, the idea is that:

> "[T]he inability to access justice is rooted in something more physical, more mundane: the many and varied costs of getting to and physically using a brick-and-mortar courthouse."

Put differently, the idea is that factors (longer commutes, e.g.) which increase litigant costs or effort of getting to court tend to increase defaults.

Other studies confirm that the cost of getting to court is important in eviction cases. In Philadelphia, Hoffman and Strezhnev (2023) find that a one hour increase in travel time increases the probability of eviction default by 3.9 to 8.6 percentage points, arguing that distance-to-court causes default evictions. The authors find that this

[9] See, e.g., Summers (2020), finding that most represented tenants with likely meritorious warrants of habitability defenses do not win rent abatement.

[10] NYC Gov (2024). See also Pollock (2024) and Woodward (2025) discussing the tenant right-to-counsel, including its history.

"transit effect" disappeared when courts made virtual accommodations in eviction proceedings during the COVID-19 pandemic.

In LA County, the specter of getting to court looms large. This is especially true for eviction proceedings and for vulnerable renter populations, as argued by tenants-plaintiffs in *Miles v. Wesley*, 801 F.3d 1060 (9th Cir. 2015):

> "[R]educing the number of courthouses handling unlawful detainer cases disproportionately impacts poor, disabled, and minority residents. ... [B]ecause individuals with disabilities and minorities are disproportionately renters who rely on public transportation, the closure of these courtrooms would have a disparate impact on these communities. [T]he importance of neighborhood court access is heightened in light of the expedited timeline of unlawful detainer actions, the fact that most low-income tenants are not represented by counsel, and the prospect that a default judgment could render a tenant homeless.

Shortly after *Miles*, the LA Superior Court system increased the number of regional eviction courthouses. Whether the distance concerns cited in *Miles* persisted is studied here: using data before and after a court expansion policy shock, I test whether the number of courts significantly affects default outcomes.[11]

Market Structure, Housing Supply, and Reducing Rents

A final factor that is important but relatively neglected by eviction scholars is the structure of housing markets themselves. Although less discussed by eviction scholars, supply-and-demand arguments suggest that supply-side housing reforms should reduce the number of evictions *ceteris paribus* because they reduce rents (see generally Baum-Snow and Duranton (2025)). Indeed, as shown in Mast (2023), new market-rate construction of even more expensive rental housing improves affordability in the market for low-income housing.[12] The intuition for the mechanism—called a "migration chain" by Mast (2023)—is that constructing new housing, even at the high-end of the market, may have spillover effects (Mast, 2023): "geographically

[11] Because the policy shock changed the tenant distance-to-court distribution for some but not all tenants, I study whether increased (resp. decreased) tenant costs increase (resp. decrease) default outcomes.

[12] See also Bratu, Harjunen, and Saarimaa (2023), which provides "empirical evidence on how the moving chain mechanism unfolds in a European city."

localized shocks are likely to have ripple effects on other parts of the region." Consequently, failure to construct new housing can impact renters at the lower end of the income distribution disproportionately, with rents increasing the most in the lowest-income zip codes.[13]

Although there are laws and regulations at all governing levels that affect the supply of housing, one set of rules of growing interest for reform are zoning and land use regulations.[14] Such reforms could reduce evictions: because the proximate cause of an eviction filing is most often non-payment of rent, interventions which reduce rental prices should (on average) reduce the number of evictions holding constant the renter population.[15] Eviction scholars to-date have not empirically tested these possibilities because:

> "[A]lthough researchers have established the centrality of zoning policies to a wide array of social problems and conditions, they lack access to national longitudinal zoning and land use data, which would help researchers identify the causal effects of exclusionary zoning." (Mleczko & Desmond, 2023)

As new data becomes available, researchers will be able to test how housing and eviction outcomes respond to changes in zoning and land use regulations.[16] The impact of housing density preferences on such regulations is especially important (Gyourko & McCulloch, 2024), particularly when certain municipal land use plans or regulations may fail to aggravate rather than alleviate segregation (Monkkonen et al., 2024).

---

[13] Pew Research Ctr., *New Housing Slows Rent Growth Most for Older, More Affordable Units* (Jul. 31, 2025); see also id. at Figure 2 ("Rents in Older, Less Expensive Apartments Decreased Most in High-Supply Metropolitan Areas"). The Pew researchers' explanation for the finding sounds in migration chain reasoning: "When not enough homes are built in high-income neighborhoods, people who would have lived in those neighborhoods can usually afford to move into middle-income neighborhoods, and middle-income residents can usually afford to move into low-income neighborhoods, but residents of low-income neighborhoods have nowhere to turn."

[14] See, e.g., Klein and Thompson (2025). But see Elmendorf, Nall, and Oklobdzija (2024), which notes that support for zoning liberalization is "much less stable than support for housing price controls and demand subsidies."

[15] In general equilibrium, the effect could go either way because lower rents tend to induce in-migration. See, e.g., Kim and Lee (2025).

[16] See, e.g., Rollet (2025) for a study of urban redevelopment in response to zoning reforms.

The Costs and Benefits of Reform: Equilibrium Rents & Composition Effects

Although well-intended, the welfare effects of access-to-justice interventions can depend on the policy method chosen. In general, Collinson, Humphries, Kestelman, et al. (2024) offer a model where legal assistance will "delay or prevent eviction," which is the primary gain accruing to tenants facing eviction.[17] Yet, legal assistance programs generate direct and indirect costs. On the direct side, some interventions may have high direct costs (e.g., right-to-counsel reforms) because they involve costly lawyering. And, on the indirect side, legal assistance to renters tends to raise rents in equilibrium: because landlords face a screening problem when deciding to whom they rent, reducing their ability to quickly re-rent in non-payment situations will tend to increase market rents.[18] Collinson, Humphries, Kestelman, et al. (2024) finds that the right-to-counsel "cost to tenants due to higher rent prices more than offsets the benefits from better eviction court outcomes."

Other studies find that procedural changes may backfire rather than improve access-to-justice in different ways. Consider, for example, Niblett and Yoon (2017), which studies outcomes from a Canadian policy shock which increased the recoverable amount in small claims cases. The change was meaningful because the informal small claims court was less expensive than Ontario Superior Court. The authors found that plaintiff suits did not increase significantly, but there were *compositional* effects. Specifically, they found that lower dollar claims (those less than the previous cap $10,000) decreased following the jurisdictional change (see Figure 2). The authors conclude that access-to-justice efforts can have unintended regressive effects, find that evidence suggesting court congestion is driving the regressive effect is mixed, and note that the optimal structure of access-to-justice interventions may depend on the court design perspective.[19]

To fully understand the welfare effects of eviction reduction policies, a necessary ingredient is credible evidence for how tenants, landlords, and markets react. As noted by Collinson, Humphries, Kestelman, et al. (2024), quasi-experimental evidence of

[17] Bell and Parchomovsky (2004) state that: "[I]t is clear that there is considerable value in stable ownership for both tenant and landlord, in particular, because tenants will often develop sentimental value for their leased premises, while landlords will frequently be better suited to extract value (due to specialized knowledge) from their premises than anyone else in the market."

[18] Landlords could also respond by increasing screening efforts or reducing maintenance costs, but Collinson, Humphries, Kestelman, et al. (2024) find little evidence for this.

[19] According to Niblett and Yoon (2017): "While economists and legal scholars have looked at the optimal structure of courts from the perspective of appeals, error correction, and *ex ante* rule setting, few have looked at the question of the optimal structure of courts from the perspective of *ex post* dispute resolution for the purposes of access to justice."

legal assistance interventions is difficult because "it is often not clear how to choose a reasonable comparison group." One possible strategy, Collinson, Humphries, Kestelman, et al. (2024) note, is leveraging a difference-in-differences design for nearby zip codes treated at different times. This paper similarly uses a difference-in-differences strategy by exploiting quasi-experimental variation stemming from changes in distance-to-court due to a court policy shock.

The Correlates of Eviction

Finally, this paper also contributes to the correlational research on factors affecting eviction outcomes (including default), such as individual tenant and landlords characteristics. Studies in this vein look at many characteristics correlated with eviction, including: demographic and income correlates of eviction (Desmond, 2012, 2017; Desmond & Kimbro, 2015), including race (Summers & Steil, 2025) and gender (Buhler & MacLean, 2024); neighborhood and spatial correlates with eviction (Lens, Nelson, & Gromis, 2020); and other case characteristics related to eviction and default (Larson, 2006), including urban-rural considerations (Armstrong & Ryan Jr., 2024). Below, I detail descriptive findings, including whether observable demographic factors are balanced across courthouses and whether they change sharply or not near the 2017 reform.

1.3 Eviction Process & Institutional Background

This section gives a brief overview of the typical eviction case in LA County. First, I summarize the standard timeline for eviction cases, adapted from the explanation in Estes and Nelson (2026). Next, I explain how eviction cases are assigned to courthouses and how the number of eviction courts evolved over time. It is the change in eviction court assignment policy over time which is the key quasi-experimental variation I utilize for causal inference in this paper.

Eviction Process

The process begins with an eviction notice, most of which are for non-payment of rent.[20] Eviction notices give tenants a few days to "cure" their lease breach by paying unpaid rent to the landlord,[21] although some landlords may permit informal arrangements for repaying rental arrears (such as negotiated payment plans outside

[20] In LA City, over 96% of the notices in 2023 were for non-payment of rent (Estes & Nelson, 2026; Los Angeles City Controller's Office, 2025).

[21] The amount of unpaid rent and number of eviction notices vary at the zip-code and building levels. Lens, Nelson, and Gromis (2020) discuss spatial autocorrelation in eviction variables.

of formal court proceedings). If the tenant does not cure non-payment, the landlord may then file an eviction (unlawful detainer) lawsuit.[22]

Tenants then file an Answer after they are served the Summons and Complaint forms.[23] The tenant is supposed to file the Answer in-person rather than by mail.[24] If they do not file the Answer, landlords may then petition for default judgment. Otherwise, a trial date is scheduled. Failure to appear at trial will also result in default, but if tenants appear there will typically be a judgment on the merits. Judgments in favor of the landlord generally awards past due rent. Landlords may then enforce the judgment via a writ of execution.[25] The typical process is represented graphically in Appendix Figure A.1 (Estes & Nelson, 2026).

The timeline here may not reflect the precise process for every eviction case, but the steps illustrate time pressures tenants face once an eviction notice has been served.[26] With the timeline of typical events specified, I now turn to explaining the court assignment procedure.

LA County Eviction Courts: Assignment Rule

According to LA County court rules,[27] eviction cases are assigned to courthouses based on a unique spatial mechanism. Under the rules, eviction cases:

> "must be filed in the courthouse serving the location and proper United States Postal Service zip code of the property in dispute using the Zip Code Table for [Eviction] cases."

Since late 2017, the eleven courthouses where eviction cases are filed include: Chatsworth, Compton, Governor George Deukmejian (Long Beach), Inglewood,

[22] In the LA Superior Court system the filing costs $240-$385. *See* Los Angeles Superior Court, *Civil Fee Schedule* (2024) (nos. 11 and 14).

[23] In LA County, tenants pre-2025 had 5 days to file the Answer. The five days do not include weekends or holidays. Additionally, tenants may have longer to respond if they are improperly served. But see Assembly Bill 2347, which as of Jan 2025 gives tenants 10 days to respond.

[24] As explained in Estes and Nelson (2026), the tenant will default if the Answer does not arrive. The official self-help page for California Courts therefore strongly recommends against mailing the Answer, and instead says you should show up to file the Answer at the relevant court. *See* California Courts (2024). Answer filing fees vary over time, but in LA County the 2024 fee in eviction actions where the contested amount of rent is less than $12,500 is $225. *See* Los Angeles Superior Court, *Civil Fee Schedule* (2024) (no. 15). But note that there are fee waiver applications available.

[25] See Estes and Nelson (2026) for further discussion of enforcement.

[26] See Estes and Nelson (2026), Lens, Nelson, and Gromis (2020), Nelson (2023), and Nelson et al. (2021).

[27] LASC Local Rule 2.3(a)(2).

Michael Antonovich Antelope Valley, Norwalk, Pasadena, Santa Monica, Stanley Mosk, Van Nuys East, and West Covina. Table 2.1 below illustrates how the courthouse assignment procedure works for the first few zip codes in Los Angeles County.

Table 1.1: Zip Code Table for Eviction Cases

Zip Code	City/Neighborhood	Modifier	Courthouse
90001	FLORENCE		Stanley Mosk
90001	HUNTINGTON PARK		Norwalk
90001	LOS ANGELES		Stanley Mosk
90002	FLORENCE		Compton
90002	LOS ANGELES		Compton
90002	LYNWOOD		Norwalk
90002	WATTS		Compton
90003	LOS ANGELES	North of Manchester	Stanley Mosk
90003	LOS ANGELES	South of Manchester	Compton
90004	LOS ANGELES		Stanley Mosk

Generally, the city-zip code pairs completely determine the assigned courthouse. However, in some zip codes the assignment is further determined relative to a particular street. For example, in Table 2.1 above, eviction cases arising in the 90003 zip code are assigned to the Stanley Mosk or Compton Courthouse if the tenant's address is north or south of Manchester Avenue, respectively.

The full assignment map for LA County is shown in the left panel of Appendix Figure A.2. The court districts differ in the land area they cover, but all courts other than the Stanley Mosk courthouse in downtown Los Angeles hear roughly the same number of cases each year. This map has been in effect since late 2017, but previously the number of courthouse districts fluctuated due to local budget constraints.[28] This is shown in the right panel of Appendix Figure A.2.

In the causal analysis of Section 1.5, I leverage the expansion of the number of courthouses which began in August 2017. Because the addition of new courthouses changed the distance-to-court for many tenants, I use quasi-experimental variation from the 2017 change in court procedure to isolate the causal effect of increasing or

[28] In particular, the 2008 financial crisis resulted in a dramatic reduction in the number of eviction courts. The number of courthouses that hear eviction cases has ranged between a low of five courthouses (2014) and a high of twenty-six courthouses (2009-2011) (Estes & Nelson, 2026; Nelson, 2023).

decreasing the distance-to-court. Companion work (Estes & Nelson, 2026) utilizes the other source of variation from the court assignment policy: the spatial nature of the assignment mechanism, as shown in the left panel of Appendix Figure A.2.

1.4 Data Collection & Descriptive Evidence

This section briefly describes the data collection process, including information collected from primary sources (docket records), linked datasets (e.g., Assessor data, Census variables), and imputation methods (e.g., racial characteristics). Next, I describe the general characteristics of the data, including looking at demographic balance across courthouses. Finally, this section estimates shocks to observed variables and discusses regression evidence of the relationship between distance-to-court and eviction outcomes.

LA Default Eviction Records Data

The data I use here includes tens of thousands of individual docket records in eviction cases scraped from the LA Superior Courts from 2016-2025.[29] The docket records contained both structured and unstructured information, including the following fields at the top of the record:

```
Case Number:  [10 character case number]

Filing Courthouse:  [court]

Filing Date:  [date]

Case Type:  [Unlawful Detainer Residential]

Status:  [Default]

Plaintiff:  [landlord name]

Defendant:  [tenant name]

Attorney for Plaintiff:  [attorney name]

Attorney for Defendant:  [attorney name]
```

[29] The years 2020-2022 are excluded from all analysis and discussion because of overlapping moratoria. See Estes and Nelson (2026).

From the individual-level text files, I used regular expressions (regex) to extract this information for each individual record, along with the following variables from the unstructured "register" of case actions: address information, monetary awards, judge, and other case timing information. Particularly important is the address information, which includes the full address (i.e., street number, street, city, state, zip code) for the majority of cases across years.[30]

Using the individual-level case address information, I geocode the address to get location information for each docket record[31] which allows encoding[32] how many default evictions are observed at each address i at time t:

$$\text{default}_{it} = \text{number of default evictions at address } i \text{ in month } t.$$

In what follows, I generally focus only on buildings that have at least one observed default eviction (or an observed eviction notice) unless otherwise noted.

Mapping the Data & Descriptive Findings

Using the records data, I then link each case to further information. Estes and Nelson (2026) linked data to building assessment information from the LA County Assessor. Here, I supplement the data further in two ways. First, I included several Census variables at the census block group level, including: median gross rent (in dollars) for renter-occupied housing units, median household income (in dollars), and demographic variables (e.g., gender, race, total census block group population).[33] Second, I used Bayesian imputation software from the R packages `wru` (Khanna et al., 2024) and `gender` (Mullen, 2021) along with tenant names to impute race and gender for each observed default eviction. The `wru` software, for instance, works to "predict individual race/ethnicity using surname, first name, middle name, geolocation, and other attributes" by "utiliz[ing] Bayes' Rule (with optional measurement error correction) to compute the posterior probability of each racial category for any given individual" (Khanna et al., 2024).

Appendix Figure A.3 details the final dataset in spatial and descriptive terms. In the top left panel of Appendix Figure A.3, I plot the geocoded defaults across LA

[30] See Estes and Nelson (2026) for further details on geocoding addresses.

[31] This can also be linked to LA County Assessor data. The Assessor data includes building-level covariates, including, e.g., the age of the building.

[32] *See* Estes and Nelson (2026) for more details on data collection and linking, including information on the case availability, FOIA data on the total number of evictions in a courthouse each month, geocoding information, and linking to LA County assessment data.

[33] *See also* Cassidy and Currie (2023) for Census variables included in their study.

County (2016–2025). The map underneath the defaulting address points is the 2017–2025 court assignment map, with eviction court district shaded in the same colors as the observed defaults but transparently. The top right table of Appendix Figure A.3 shows the average distance-to-court (in miles) for cases by observed filing courthouse under two court assignment policies. For example, for cases actually filed in Compton courthouse pre-reform, the average distance-to-court was 8.93 miles; by contrast, the average distance-to-court post-reform would be 3.47 miles. Whether the change in the distance-to-court distribution induced by the 2017 court expansion reform meaningfully impacted default evictions is studied below in Section 1.5. Note that in some cases, the decrease in average distance-to-court was large whereas in other cases it is modest (or, in the Stanley Mosk case, a slight increase in the average distance-to-court).

The bottom table of Appendix Figure A.3 shows averages in several covariates by courthouse. For example, the average imputed white probability (following the imputation procedure described above) is 31% for observed eviction cases filed in Antelope Valley courthouse. Average incomes across the filing courthouses range from $50,193–$83,420 approximately. Average rents are positively correlated with the median income and range from $1250–$1939. The percentage of observed cases with tenant legal representation is around 1-5% across filing courts. The observed data is consistent with the standard empirical portrait of eviction: tenants are most often minority renters with below median incomes who lack legal representation at eviction proceedings.

Finally, I also use the data to estimate the general relationship between distance-to-court and default outcomes using binned averages and non-parametric regressions. In Appendix A.3 (and, in particular, Appendix Figure A.7), I estimate the default probability as a function of distance-to-court pooling across years. The estimated relationship shows that the default probability is low but roughly increasing as distance-to-court increases (at least up to a certain distance).

1.5 Court Expansion Study

This section uses the LA County eviction data and changes in court assignment policy over time. Specifically, I estimate the impact of expanding the number of eviction courts in the second half of 2017. Because this expansion changed the distance-to-court for some but not all locations, I model defaults as a function of the change in distance-to-court. I show that units with an increased distance-to-

court following the court expansion policy had higher average defaults post-reform relative to control units (no change in distance-to-court) and units with a decreased distance-to-court had lower defaults. But I caution against causally interpreting naive before-after comparisons.

Instead, I describe the the staggered difference-in-differences (DID) counterfactual framework to assess causality. The DID framework is popular in empirical law and economics studies (≈ 40% of 2016–2024 Law and Economics NBER working papers per Goldsmith-Pinkham (2026)) and the intuition underlying it is briefly explained. Using this framework, I find that DID estimates imply directional treatment effects consistent with the distance-as-tenant-cost model, but only for units with large changes in distance-to-court. The effect is subject to model specification uncertainty and aggregate estimates are not always aligned with expectations under the tenant cost theory.

Expanding the Number of Eviction Courthouses & Key Expectations

Over time, LA County has expanded the number of eviction courthouses.[34] In 2015, seven eviction courthouses heard eviction cases in LA County.[35] But in September 2015, the LA County Superior Court system increased the number of eviction courts from seven to eight, splitting off part of the Pasadena eviction court district for coverage by the Pomona courthouse. This is shown in the left 2016–Early 2017 map of Figure 1.1. Then, beginning in August 2017, LA County expanded the number of courthouses more substantially, adding 3 new courthouses and sending Pomona court cases to the West Covina courthouse.[36]

[34] *See* Figure A.2 (right panel).

[35] Note that for 2015 and 2016, due to eviction data unavailability we cannot know with certainty the contours of the eviction map in the northern-most parts of LA County or in some of the south-eastern regions. Only zip codes with uniquely-assigned courts and no potentially misfiled cases under the then-operating eviction assignment rule are shown in the left and middle maps. *See also* Estes and Nelson (2026) discussion of misfiling, map creation, and the assignment rule.

[36] Formally, this is a staggered rollout: the change from Pomona to West Covina happened in August 2017. The addition of the Chatsworth, Compton, and Inglewood courthouses occurred in September 2017. Remaining changes to the assignment rule were effective Oct 10, 2017.

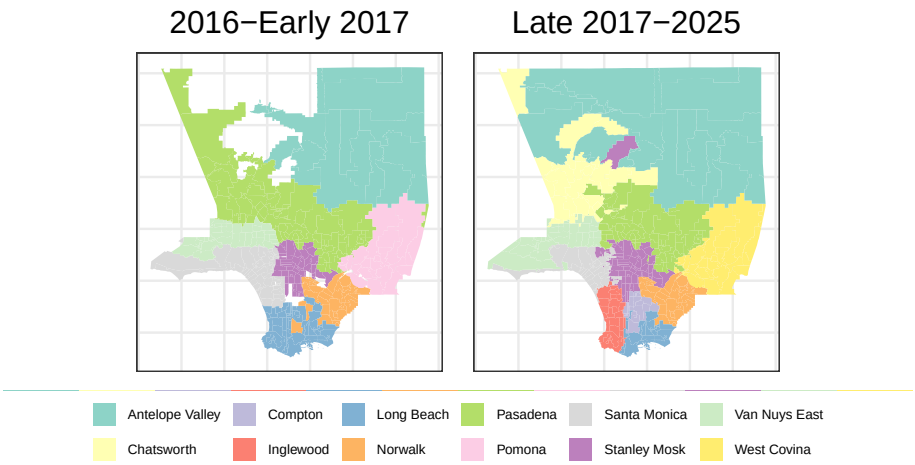

Figure 1.1: LA County Courthouse Expansion

Note: All maps omit the Catalina courthouse, where there are few evictions and no observed defaults.

I utilize all available data collected (from 2016 onwards) to study the effect of the 2017 policy change in court assignment on tenants. Because the court assignment policy change induced changes in distance-to-court, I test whether default evictions responded to changes consistent with the tenant cost theory of defaults. Recall that on this theory, defaults are caused by the myriad costs of getting to court, with distance-to-court being a large obstacle in LA County specifically.[37] We therefore expect that tenants with increased costs (higher difficulty in getting to court) should experience higher rates of default. Conversely, lowering costs by assigning tenants to closer courthouses should tend to decrease defaults.

Accordingly, I begin by splitting tenants into three cohorts.[38] The first group I call the Increase-Treated group: tenants in this group experienced an increase in distance-to-court because of the 2017 court expansion policy shock, which represents an increase in the tenant cost of getting to court. The second group is the Decrease-Treated group faced decreased "costs" post-August 2017, meaning they experienced a decrease in distance-to-court following the court expansion. Finally, the last group

[37] *See supra* (discussing *Miles v. Wesley*).

[38] To formalize the cohorts, define the following variable:

$$\Delta_i = D_{i,Post-Expansion} - D_{i,Pre-Expansion}$$

where $D_{i,Post-Expansion}$ (resp. $D_{i,Pre-Expansion}$) is the distance from unit i to the assigned courthouse after (resp. before) the 2017 courthouse expansion policy. The cohorts are defined as: Increase-Treated ($\Delta_i > 0$)), Decrease-Treated ($\Delta_i < 0$), and Control ($\Delta_i = 0$).

is the Control group, which did not undergo a change in court assignment and therefore experienced the same distance-to-court before and after the 2017 reform. The cohort definitions and expectations under the tenant cost theory are summarized below in Table 1.2.

Cohort	**Definition**	**Tenant Cost Expectation**
Control	Distance Before = Distance After	
Increased	Distance After > Distance Before	Defaults Up ↑
Decreased	Distance After < Distance Before	Defaults Down ↓

Table 1.2: Cohorts Definitions and Expectations

The changes from the August 2017 policy shock on distance-to-court are also summarized in Figure 1.2. In the left panel of Figure 1.2, I show the histogram of the changed distances across all of LA County. The large black bar (apprx. 23,000 addresses) are Controls: they did not experience any change in distance-to-court from the 2017 court expansion. The dark blue bars plotted to the left (less than zero) are Decrease-Treated units: tenants living at these addresses were closer to court because of the 2017 court expansion. Although fewer in number than the Control group, the Decrease-Treated cohort outnumbers the Increase-Treated cohort: units with an increased distance-to-court following the 2017 court expansion are shown in dark red bars on the histogram.

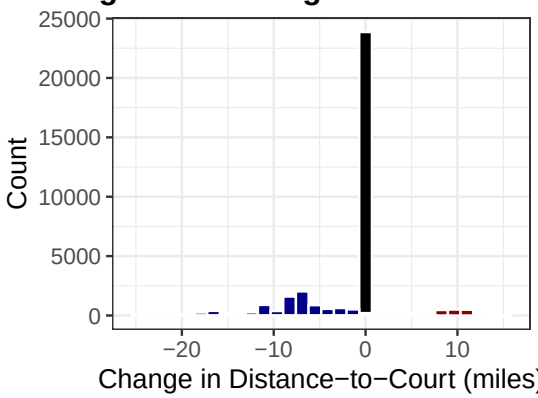

Changed Distance by Court

Court 2016	Control	Decreased	Increased
Antelope Valley	2053	0	0
Long Beach	4223	1643	262
Norwalk	3632	608	0
Pasadena	3083	2038	0
Pomona	0	2174	1088
Santa Monica	1936	2167	150
Stanley Mosk	7050	424	171
Van Nuys East	1695	118	979

Figure 1.2: Changes in Distance-to-Court from 2017 Court Expansion Policy

Note: The left panel shows a histogram of changes in distance-to-court following the 2017 Court Expansion: controls are shown in black, decrease-treated units in blue, and increase-treated units in red. The right panel table shows the cohort counts by 2016 assigned courts.

The number of addresses within each cohort are broken up by assigned 2016 court in the right panel of Figure 1.2. Four of the 2016 court districts had some units in all three cohorts (Long Beach, Santa Monica, Stanley Mosk, Van Nuys East), two were only in the Control or Decrease-Treated cohorts (Norwalk & Pasadena), one had no Control units (Pomona), and one had only Control units (Antelope Valley).

To confirm that the 2017 Court Expansion policy shocked tenant costs but not other important eviction determinants, I measure twelve variables across eviction records before and after the policy shock in Figure 1.3. Fitting separate LOESS regressions to before and after the 2017 reform, I investigate whether there are significant discontinuities in important covariates other than the tenant cost of getting to court. If a significant difference in, e.g., legal representation rates around the policy treatment timing were discovered, researchers may suspect that change is related to potential changes in outcomes-of-interest rather than the change in the policy treatment. Figure 1.3 shows that tenant costs (measured by either distance-to-court or commute-to-court) significantly and immediately changed at the policy reform timing, which is indicated by the vertical dashed line. By contrast, other socioeconomic variables (predicted race, predicted gender, imputed income, imputed rent), legal representation variables (plaintiff and defendant representation rates), and procedural motions are not significantly different at the 2017 policy shock date.

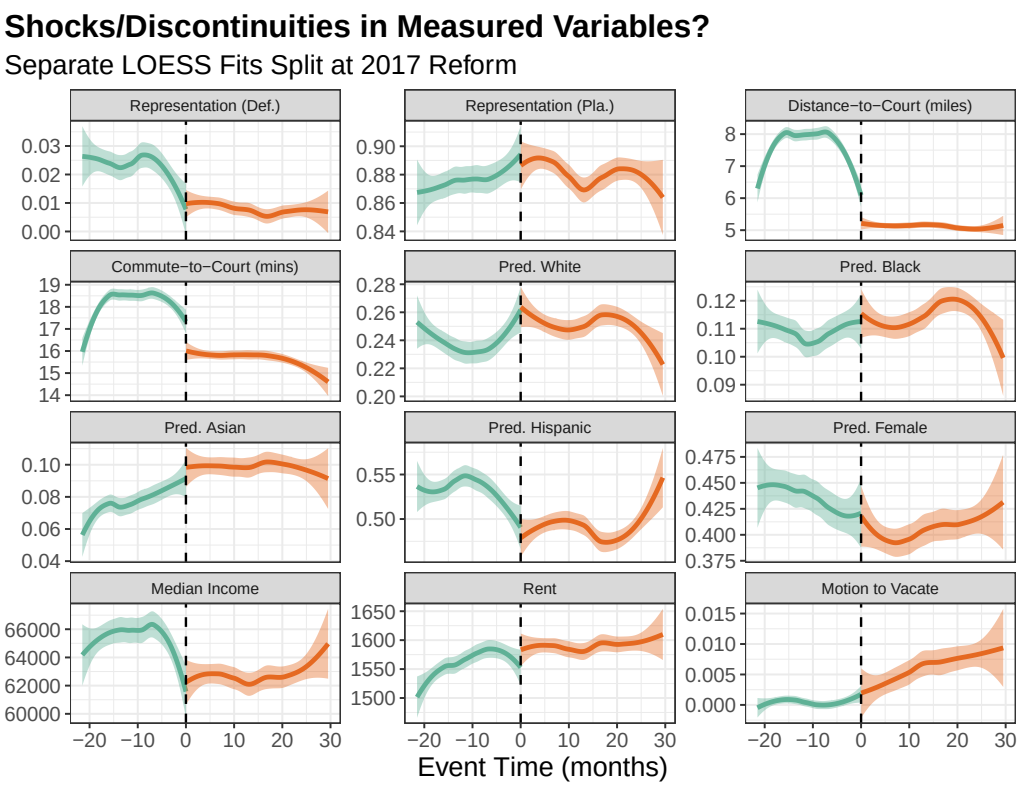

Figure 1.3: Checking for Discontinuous Shocks near Policy Reform

Note: Each covariate panel estimates the average among all observed eviction records relative to the 2017 court expansion policy. The x-axis is in event time to account for the staggered policy treatment. Separate LOESS regressions are fit on each side of the treatment using data from 2016–2020 only.

Using the defined cohorts, I next turn to plotting the average number of default evictions over time for each cohort. In Figure 1.4, I plot estimates of the average number of defaults for the cohorts in different colors for each year with collected and available data. The average number of default evictions for the Increase-Treated cohort is shown in green, the average for the Decrease-Treated cohort is shown in orange, and the average for Controls is shown in black. The court expansion policy shock is shown as the vertical dashed line in August 2017.

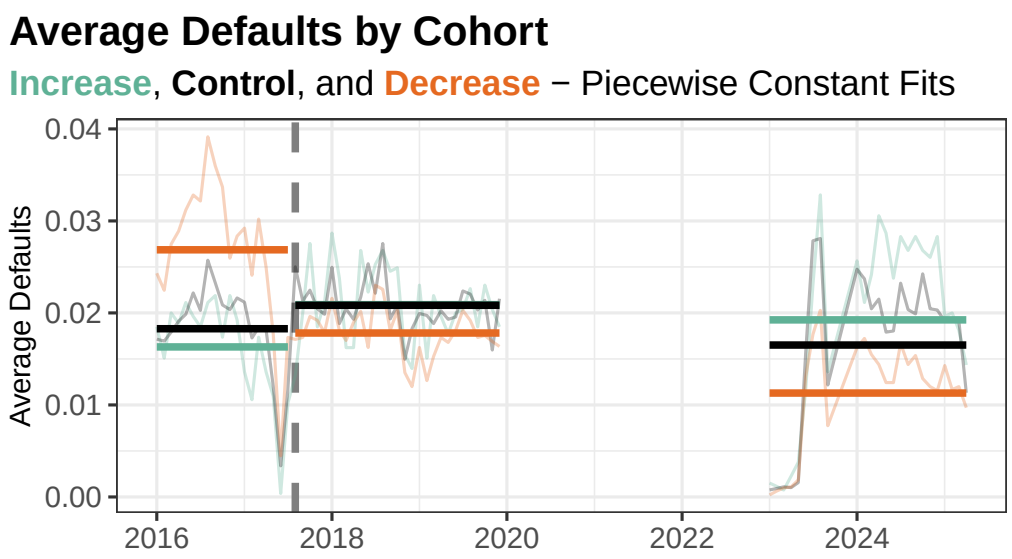

Figure 1.4: Cohort Averages: Piecewise Fits

Note: The plot shows the average number of default evictions for the three cohorts: Increase-Treated (green), Decrease-Treated (orange), and Controls (black). The raw averages are plotted transparently for time periods with available data. The piecewise constant fits are overlaid in the same colors. The dashed vertical line is plotted at the beginning of the policy shock (August 2017). Note that this plot shows averages for cases labeled with "default" status only, not all observed evictions.

Following this shock, the estimated average for the Increase-Treated cohort is higher than the Decrease-Treated cohort and about the same as the Control cohort, suggesting such tenants experienced the highest average defaults post-reform. The Increase-Treated group average post-reform is also higher than the pre-reform average for the Increase-Treated cohort. This observation is consistent with the tenant cost theory of default, since an increase in distance-to-court is an increased litigation cost which is expected, on average, to increase defaults. By contrast, the Decrease-Treated cohort saw the average defaults fall following the August 2017 policy change, which is also consistent with the tenant cost theory: lower distance-to-court is expected to lower defaults. But note that, for instance, average defaults were decreasing for the Decrease-Treated cohort pre-August 2017, so we cannot attribute causality on the basis of lower average defaults alone. Moreover, the control group also saw an increase in average defaults even though it experienced no change in distance-to-court costs. Therefore, in order to assess causality, we need to use a causal design (a counterfactual model) that models counterfactual outcomes for our cohorts.

Difference-in-Differences (DID) Design

One such counterfactual design, which is widespread in economics, is the difference-in-differences (DID) design. This strategy assumes that counterfactual outcomes for the treated group would have evolved "in parallel" to the control group (the *parallel trends* assumption). The DID strategy accounts for differences in baseline outcomes (pre-treatment differences) while controlling for time factors unrelated to the treatment. The results in this section rely on popular methods for DID designs.

In the standard causal inference setting, the researcher first specifies potential outcomes in two states: the treated and untreated (or control) states. The "fundamental problem of causal inference" (Holland, 1986) is that we cannot observe individual outcomes in both states simultaneously. However, under randomization, estimating the average treatment effect is straightforward: simply take the difference in average outcomes between the treated and control groups.

Unfortunately, not every causal question can be (or is) addressed by running a randomized experiment due to ethical, legal, financial, or other practical constraints. Researchers therefore turn to observational studies and methodologies to isolate causal effects. The most popular observational methodology in economics is the difference-in-differences (DID) approach: more than 30% of applied microeconomics working papers mentioned DID or similar designs in 2024.[39]

In the DID design, it is similarly assumed that the researcher has access to data on two groups (the "treated" and "control" groups) over time. At some point in time, the treated group receives a treatment, whereas the control group does not. The data available includes an outcome variable from before and after the treatment date for both groups. The goal is to estimate the causal effect of treatment on the treated group after the treatment date, which is called the average treatment effect on the treated (ATT). The "fundamental problem of causal inference" (Holland, 1986) in this setting is that we do not observe untreated outcomes for treated units after the treatment date. The DID method assumes that the average untreated outcome for units that actually receive treatment would have evolved in parallel to the control group average outcome.

In Appendix Figure A.6, I show how this assumption implies a counterfactual average outcome for the treated group. In this example, the treated group (solid blue line) receives treatment in the year 1950 and the control group (solid red line)

[39] Figure 5 of Goldsmith-Pinkham (2026).

never receives treatment. Naively taking differences post-1950 between the treated group average outcome and the control group average outcome would result in an over-estimate of the true treatment effect because the treated and control groups start (and stay) at different baselines in the pre-treatment periods. Indeed, in Appendix Figure A.6, the treated group average outcome at the beginning of the study period (year 1900) is 5 whereas the control group average outcome is 0.

Instead, the DID method assumes the counterfactual outcome for the treated group would've evolved "in parallel" to the control group average. This results in the dashed blue line for the treated group counterfactual. The true average treatment effect on the treated (ATT) is the difference between the observed outcome (solid blue line) and the counterfactual outcome (dashed blue line): it is the average difference in where treated units are versus where they would have been without treatment. This is illustrated in Appendix Figure A.6 by the double-sided black arrow labeled "Treatment Effect." Note that it is not a requirement of the methodology that the average treatment effect on the treated be constant in all post-treatment time periods.

Because this is a staggered rollout, with changes happening in August 2017 (Pomona to West Covina) and September 2017 (addition of Chatsworth, Compton, and Inglewood), we apply the staggered DID methodology of Callaway and Sant'Anna (2021). Formally, let $Y_{it}(1)$ be the potential number of defaults at address i in period t if the address has a change in distance-to-court, and $Y_{it}(0)$ the corresponding potential outcome under no change in distance-to-court. The observed number of defaults at address i in time t is given by Y_{it}. Define ΔD_i as the change in assigned distance-to-court at the time of treatment, and let G_i denote the period for each address i in which this change first occurs, with $G_i = \infty$ for never-treated controls as in Callaway and Sant'Anna (2021).

The causal estimand of interest for addresses whose assigned distance-to-court increased is defined as the average treatment effect on the treated at event time e by

$$ATT_{increase}(e) = \mathbb{E}[Y_{it}(1) - Y_{it}(0)|G_i = g, t - g = e, \Delta D_i > 0]$$

where again G_i denotes the period in which the distance-to-court first increases. Recall that we define things with g cohorts because of the staggered rollout. Similarly, define the estimand for addresses whose distance-to-court decreased as:

$$ATT_{decrease}(e) = \mathbb{E}[Y_{it}(1) - Y_{it}(0)|G_i = g, t - g = e, \Delta D_i < 0].$$

These parameters are ATTs because expectations are taken over addresses with a change in distance-to-court, with never-treated control units being used to iden-

tify counterfactual untreated trends for such addresses. Under the parallel trends assumption, these event-time ATTs measure how the number of default evictions evolves relative to the counterfactual untreated trend for the Increase and Decrease treated cohorts, respectively.

Baseline Results: Applying the DID Strategy

Here, I apply the DID methodology (Callaway & Sant'Anna, 2021) to the court expansion policy. I assume that units with changes in distance-to-court would've experienced parallel outcomes to units with zero changes in distance-to-court. Recall that I consider three cohorts of units: the Control units (zero change in distance-to-court), the Increase-Treated units (positive change in distance-to-court), and the Decrease-Treated units (negative change in distance-to-court). The treatment effect of interest is the average treatment effect on the treated (ATT), which is the difference in expected or average outcomes from treatment versus control for units that actually receive treatment. Event-study dynamic DID estimates are shown below in Figure 1.5.

Consider the first set of DID results comparing Increase-Treated units (i.e., units that experienced an increase in distance-to-court) with control units as the comparison group. This means we are using control units to impute counterfactual outcomes for the Increase-Treated units in the counterfactual world where these units do not experience an increase in distance-to-court. In the left panels of Figure 1.5, I plot the ATT estimates from this comparison, with and without additional covariates.

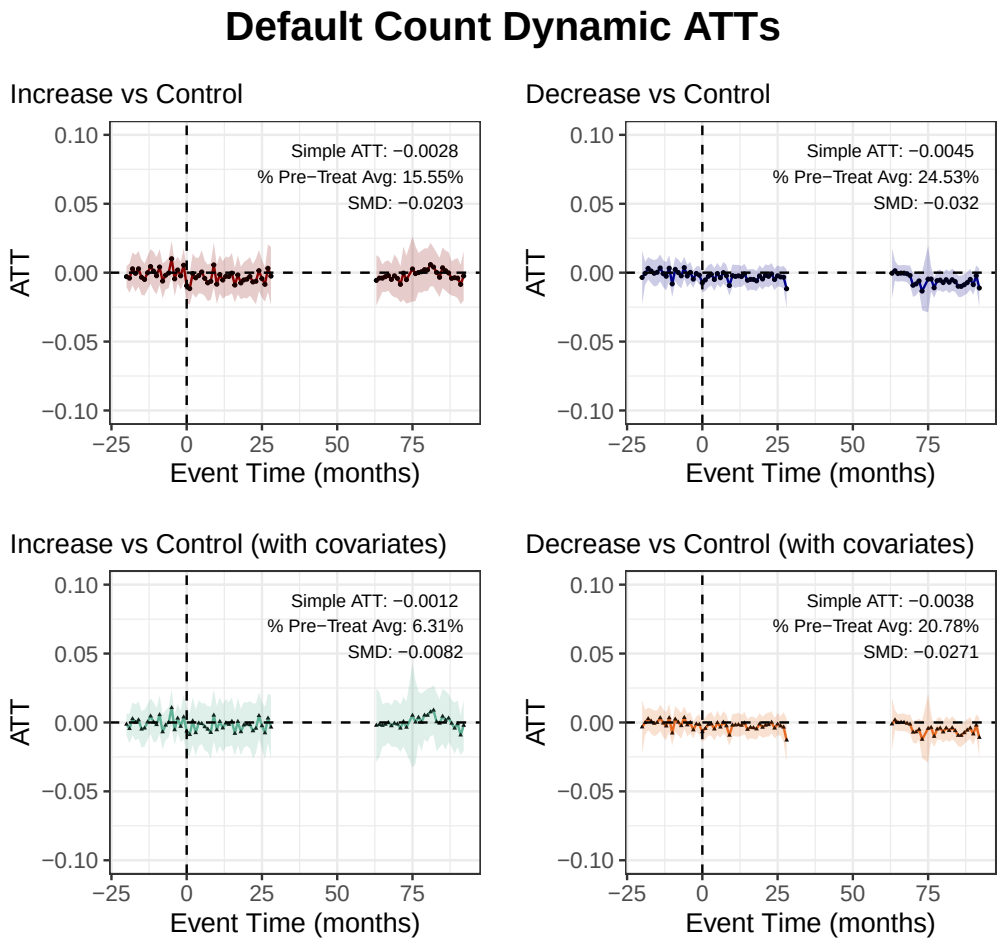

Figure 1.5: Dynamic DID Estimates: Comparisons with Controls

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations. The covariate models include gender, race, and income variables.

The point estimates after the court expansion policy (vertical dashed line) are mostly negative and all are insignificant. This means that units with an increased distance-to-court experienced a lower number of defaults on average than units with zero change in distance-to-court following the court expansion policy. However, note that the standard errors include the zero effect for all time periods (dashed horizontal line at 0) and the standardized mean difference (SMD = −0.0203) implies a small effect size. Finally, note that the aggregate effect (simple ATT = −0.0028) does not align with the expected direction, although it too is insignificant.

The right panels of Figure 1.5 repeat this exercise comparing units that are Decrease-Treated with Control units. The post-expansion ATT estimates for the effect on the Decrease-Treated units relative to the Control units are mostly negative and the overall aggregate ATT estimate (a weighted average across post-treatment periods) is negative and significant. The size of the aggregate ATT estimate can be assessed multiple ways: it is 20-24% of the Decrease cohort average in the pre-treatment period and is small relative to the pre-treatment variation (SMD ≈ -0.03). Note, however, that the estimates here too reflect a high degree of uncertainty: most of the post-treatment ATTs are not statistically significant, although the aggregate effect in the no covariate specification is significant and negative (in accord with tenant cost theory expectations).

In sum, I find some evidence that court expansion significantly impacted default evictions for early 2017 units using all Increase-Treated and Decrease-Treated units versus the Control group. Although the aggregate estimate signs sometimes align with expectations and are significant in one model, I assess below the sensitivity of the estimates to model specification and find model uncertainty. In addition, the event-study dynamic effects are almost always insignificant and in some cases do not align with tenant cost expectations (esp. for the Increase vs Control comparisons).

These comparisons use all units that experienced an increase or decrease in distance-to-court as a treatment cohort. What about units which experienced "large" changes in distance-to-court? I repeat the exercise comparing treated units (those with a change in distance-to-court) versus control units (no change in distance-to-court) by considering only treated units with above-median changes in distance-to-court. The results for this comparison are shown in Figure 1.6.

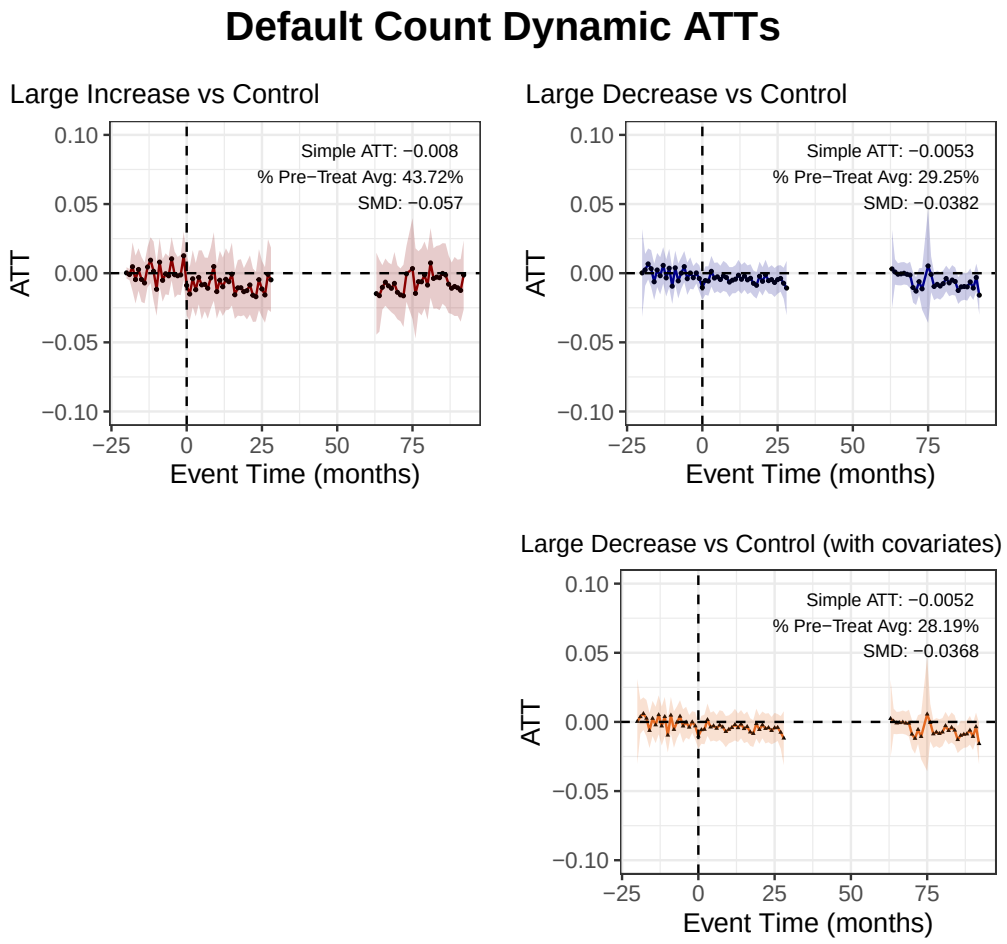

Figure 1.6: Dynamic DID Estimates: Above-Median Comparisons with Controls

Note: The panels show estimates of the average treatment effect on the treated for Above-Median Increase-Treated (left panel) and Above-Median Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations. The large increase sample size is too small to include the covariate specification, but the large decrease specification includes gender, race, and income variables.

In this case, the results are subject to increased uncertainty because they utilize a smaller sample.[40] The simple ATT estimate for units with an above-median increase in distance-to-court is negative and implies treatment (the August 2017 policy) decreased defaults. But note that this point estimate is insignificant and, therefore, I cannot rule out a true positive effect for units with a large increase in tenant costs (as measured by distance-to-court). By contrast, the aggregate estimate ("simple

[40]This is reflected in the larger confidence intervals on the ATT estimates.

ATT") for units with an above-median decrease in distance-to-court is negative and significant (with or without covariates), implying treatment (the August 2017 policy) decreased defaults but the effect size is small (SMD ≈ -0.04). In other words, there were fewer defaults for tenants with a large changes in distance-to-court but only the above-median decrease model is significant. Note, however, that the post-treatment period ATT point estimates and aggregate estimates (simple ATTs) are, in most periods, not statistically significant.

Model Specification & Robustness

In this subsection, I discuss several robustness checks. The checks test how the findings change using additional data, measurements, covariates, outcomes, or criteria to define the population of interest.

Tenant Cost Measurement. First, I test whether the results are sensitive to our distance-to-court measure of tenant costs. Rather than using distance-to-court, I measure tenant costs by changes in commute time to court. Because distance-to-court and commute times are highly correlated (0.84), this change in how tenant costs are measured should not make a large difference in the estimates. In Appendix A.4, I show that the dynamic event-study often remain insignificant, do not always align with expectations under the tenant cost theory, and are small in magnitude (as measured by standardized mean differences). Aggregate estimates are very similar to our baseline distance-to-court estimates: only the Decrease cohort model without covariates aligns with expectations and is significant.

Address Population. Next, I consider alternative ways to define the address population-of-interest. The baseline model uses the set of all observed addresses to form a panel dataset tracking defaults at those addresses over time. There are two potential problems with this: first, the analysis omits many addresses at-risk of eviction which experience zero defaults and are therefore not in the panel dataset; and, second, using all observed addresses may fail to account for possible selection and composition effects if certain addresses appear post-reform with defaults that did not have defaults pre-reform.

To address the first possible problem, I use new data from the LA City Controller office. Because a 2023 change in LA City law required all landlords to submit eviction notices to the City Controller, I can construct a set of addresses "at-risk" of eviction. Using this set of addresses only, I repeat the staggered DID estimation and plot results in Appendix Figure A.8 and Appendix Figure A.9. Pre-trend

differences for Increased vs Control cohorts appear smaller, although the same issues remain: insignificant, small effect estimates which (in some cases) do not align with expectations.

To address the second potential selection problem, I also analyze the effects of the reform only considering addresses that appeared in the default eviction records prior to August 2017. In Appendix Figure A.12 and Appendix Figure A.13, I find that the uncertainty increases (smaller samples) but the size of the standardized point estimates is slightly larger. Additionally, the direction of the aggregate estimates from this model ("Pre-Aug 2017" below) align with expectations but are insignificant.

Covariates. Next, I test whether covariates make a meaningful difference to the significance of the aggregate ATT estimates. Because not all covariates are available for every address (e.g., not all addresses have predicted race), I test the sensitivity of aggregate ATT estimates to the inclusion of many possible covariate combinations. Specifically, I use the all eviction panel dataset, all possible four covariate combinations, and compute aggregate ATT estimates. Results are shown in Appendix Figure A.20 and Appendix Figure A.21 for both the Decreased vs Control and Increased vs Control cohort comparisons. Darker colors (dark red or dark blue) indicate significant aggregate estimates, whereas lighter colors (light red or light blue) indicate insignificant estimates. Positive and negative aggregate estimates are shown in blue and red, respectively.

The specification curves suggest that the significance and sign of the Decreased vs Control aggregate estimates are sensitive to the covariate specification. By contrast, the Increased vs Control estimates are all insignificant, although the direction of the effect is sensitive to model specification. In general, the specification curve estimates show that aggregate ATT estimates are sensitive to model specification, generally insignificant, and therefore cannot be sign-identified. In both the Decreased vs Control and Increased vs Control cases, the majority of the aggregate estimates across specifications are positive and insignificant.

All Observed Outcomes. I also examine whether the reform had effects when we use all eviction outcomes, rather than only outcomes marked with "default" status. Using all observable evictions as the outcome-of-interest, I repeat the modeling exercise as in our baseline model and re-estimate the event-study dynamic plots in Appendix A.6. Aggregate estimates are all insignificant but the sign of the aggregate ATTs in the no covariate models do align with the tenant cost theory expectations.

Relaxing Parallel Trends. What about relaxing the parallel trends assumption underlying the DID analysis? Given that the parallel trends-based estimates are already insignificant, relaxing the assumption is unlikely to yield sign-identification. Nevertheless, I show in Appendix A.5 that a Manski-style partial identification approach adapted to our setting may be employed (C. Manski & Pepper, 2018; C. F. Manski, 2007). The method allows for parallel trend violations provided counterfactual cohort averages remain "sufficiently close" to observed cohort averages. Applying this partial identification strategy does not yield sign-informative bounds for most post-Aug 2017 time periods. Bounds on the standardized ATTs are, moreover, suggestive of small effect sizes under rule-of-thumb interpretations.

Summary. Finally, I plot the aggregate estimates from each of the previously models discussed models. In Figure 1.7, aggregate ATT estimates from different model specifications are shown for both Increase vs Control (green estimates) and Decrease vs Control (orange estimates) cohort comparisons. Additional models with covariates are included where sample size permits, which are shown with triangle (rather than circle) estimates.

The models show that, in most cases, aggregate ATT estimates are insignificant, do not align with tenant cost expectations, or both. Additionally, all aggregate estimates are of small effect sizes (SMD less than 0.20). Some models with restrictions on the time period of study (e.g., "pre-2020" panels only) are also estimated to account for possibility of post-moratoria changes in eviction practice affecting estimates, but this does not make a noticeable difference on the aggregate ATTs shown below.

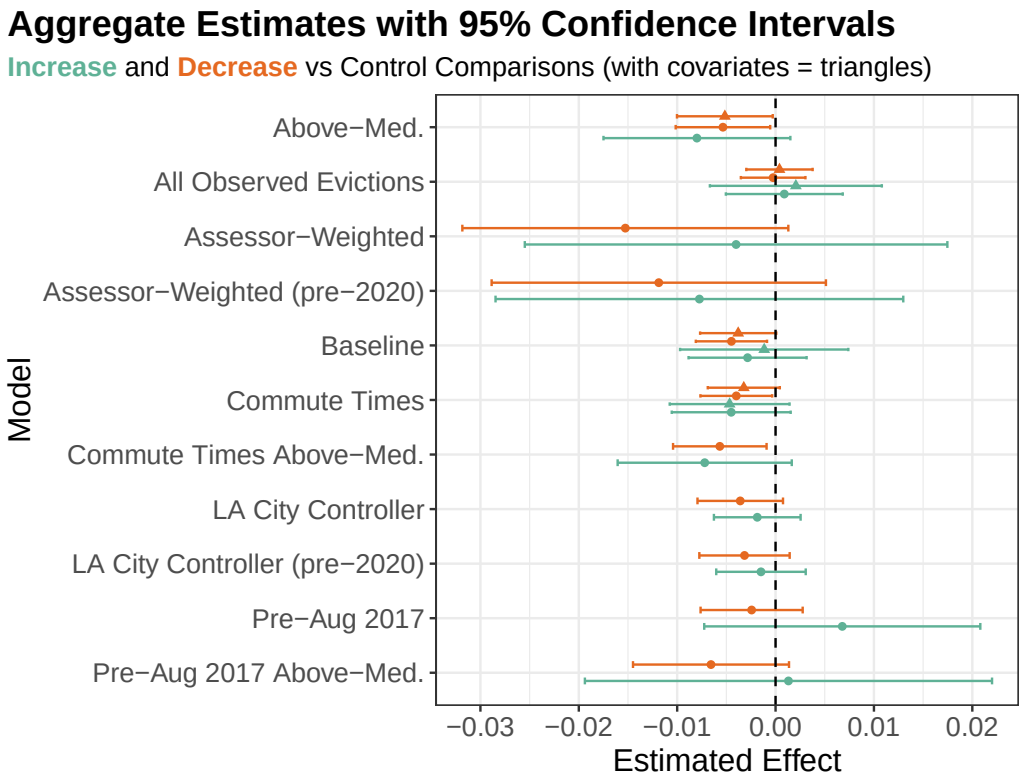

Figure 1.7: Model Summary: Aggregate DID Estimates

Note: Model names listed on the y-axis are explained above and further elaborated in the appendices. Green estimates are Increase vs Control comparisons and orange estimates are Decrease vs Control comparisons. In models with sufficient sample size, covariate specifications including gender, race, and income variables are also estimated (results shown with triangles). The standardized mean differences (SMD) across all models and comparisons range from −0.1086 to 0.0265, which suggest small effect sizes under standard rule-of-thumb interpretations.

1.6 Discussion & Policy Implications

What lessons to draw from the empirical exercises highlighted in Section 1.5? Because of uncertainty in the estimates (stemming from relatively small treated cohorts), I first caution against over-generalization. The point estimates are consistent with the tenant cost theory under some specifications, but we often cannot rule out zero effects in most post-treatment time periods or in the aggregate. If legally meaningful effects are limited only to a subset of the population (e.g., those with the highest distance-to-court values), then a larger dataset would be needed to discern the magnitude and direction of the effect.

On the other hand, the results suggest that expanding access via increasing the number of operating courts is not a panacea, at least not if expansion is a modest three court increase in a sprawling metropolitan region like Los Angeles County.

However, the results do not imply that distance-to-court is unimportant or that larger reforms would not find large effects: the results show that this particular 2017 reform (from 8 to 11 courts) did not significantly reduce defaults. In some filing courthouses, the average change in the distance-to-court may have been insufficient (see Appendix Figure A.3) to have meaningful effects. Below, I take a comparative institutional perspective and examine how the LA County filing rules might be meaningfully different from rules in other California counties.

This study accords with some of the procedural intervention literature discussed in Section 1.2, showing that relatively limited interventions are often of limited efficacy.[41] Larger changes to the eviction system may prove meaningful, but work remains to find cost-effective measures to optimize the legal details of the eviction system. I turn now to discussing some possibilities, including additional work focused on LA County and future research possibilities.

Comparative Institutional Analysis: California Eviction Filing Rules

To facilitate extrapolation and position our results within the context of other California eviction systems, this section estimates distance-to-court distributions for other large California counties. To best knowledge, this is the first attempt to empirically estimate and compare tenant costs of getting to court in California.

Using Eviction Lab data from 2010-2018, I approximate distance-to-court densities for 14 California counties plus LA County pre- and post-reform. First, I collected eviction filing rules for these 14 counties from court websites and court procedural rules. Most of the counties use a single courthouse to hear eviction cases, although other rules include: a nearest-court rule, a zip code rule (like LA County), and a complex neighborhood-case type rule. See Appendix A.8 for further details on the filing rules in these counties. Next, using Eviction Lab data on census tract eviction filings from 2010–2018, I approximate the distance-to-court densities by computing the distance from each census tract centroid to the assigned courthouse, and then weighting by the number of filings in that census tract. Densities are approximated for all California counties with more than 100 census tracts in the Eviction Lab data, as counties with too few census tracts may not well-approximate the true underlying density (e.g., census tracts may be too large). In addition to combating measurement error in the distance variable, much smaller counties are unlikely to be comparable to LA County in relevant respects.

[41] Some interventions, moreover, may backfire. *See supra* discussion of regressive effects of expanding small claims amounts.

The resulting approximations are shown in Figure 1.8. The approximated densities suggest that LA County pre-reform was comparable to other nearby southern California counties (e.g., Orange and Riverside Counties) in terms of the average distance-to-court and distribution shape. Post-reform, LA County is among the lowest distance-to-court counties here examined, with a reduction of 2.7 miles in the average distance-to-court. As discussed above, the results here imply that reduction may not have had a large practical impact for LA eviction outcomes. However, it does not imply that tenant cost concerns may not loom large in other settings. Indeed, the estimated densities in Figure 1.8 show there are at least 4 other California counties where distance-to-court is, on average, more than 10 miles away. Many of these counties have only one court hearing eviction cases, whereas LA County had 8 courthouses before the late 2017 court expansion. Further study using docket data or renter surveys may shed light on whether the cost of getting to court merits court expansion in high distance-to-court counties. The same estimation using commute times instead of distance-to-court proxy is repeated in Appendix Figure A.23.

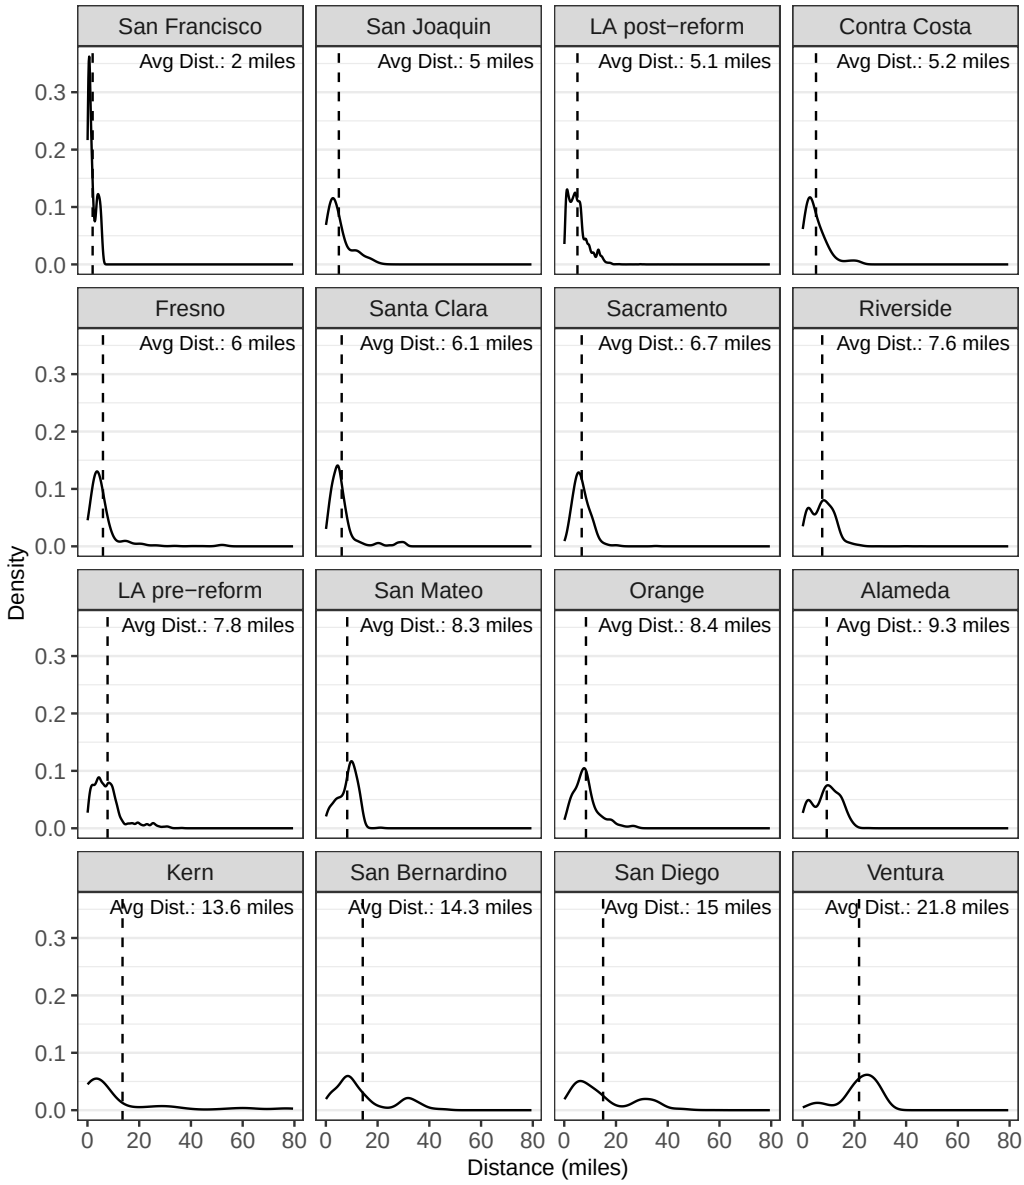

Figure 1.8: California County Distance-to-Court Densities

Note: Except for LA County, the densities are approximated using Eviction Lab data on census tract filings from 2010-2018. The centroid of each census tract is used to compute distance-to-court densities weighted by the number of Eviction Lab filings reported. The average distance-to-court is shown as a vertical dashed line in each county panel. The LA County pre- and post-reform density estimates are constructed using the collected micro-data discussed above.

LA County Localized Differences: Regression Discontinuity (Estes & Nelson, 2026)

One possible strategy is explored in Estes and Nelson (2026), where we examine the boundaries of the spatial map in LA County to test whether how the map is drawn impacts eviction outcomes. The idea is that, unlike the Fedex or Amazon delivery system, legal assignment rules are not designed, tested, and regularly optimized. Although the designers of legal rules and institutions may intend to improve welfare (or at least do no harm), legal rules often need to be fashioned without empirical evidence and are, therefore, non-optimal. Identifying areas for improvement with empirical study can assist legal and policy decision-makers to better fashion alternative rules, especially when it comes to the important task of drawing and/or redrawing maps.

This companion paper studies the spatial aspect the case assignment mechanism and suggests regions where alternative maps may be most effective. Using spatial variation stemming from the discontinuous assignment of cases to eviction courthouses, the paper estimates the effects of eviction procedures on tenant outcomes locally at courthouse boundaries.[42] The main contribution of the paper is to test whether, at the boundary of several eviction courthouse districts, there are significant differences in the default probability for tenants on one side of the boundary versus the other due to court assignment.

Formally, this type of study is a spatial regression discontinuity design, which takes advantage a spatial policy that creates borders between regions. A local causal effect is estimated by comparing cases along a policy border that receive different treatments. The intuition is that cases located nearby but across some quasi-random policy boundary have "similar" observed and unobserved characteristics. Under this assumption, differences between nearby units are therefore attributable to the difference in the spatial "treatment" variable because units are similar close to the boundary in expectation. Further tests of the distance mechanism specifically are carried out under additional causal assumptions.

In those regions where large local effects appear persistent, "re-optimizing" the court assignment map may help reduce default evictions. Although, as noted by Goldsmith and Vermeule (2002), lawmakers must often "make decisions in the short term under conditions of empirical uncertainty" thereby trading off "accuracy

[42] *Cf.* Collinson, Humphries, Kestelman, et al. (2024), which uses spatial comparisons from nearby zip codes.

for relevance and timeliness," court policy can be improved in the long-term by empirical assessment. Indeed, further investigation of adaptive court assignment policies—with periodic assessment, analysis, and adjustment—may be warranted to alleviate access-to-justice problems.

Transportation & Structural Reforms

Another possibility for further study is the impact of investment in transportation infrastructure on eviction outcomes. Although I am unaware of comprehensive panel datasets on changes to the transportation network in LA County, future work might look at how transportation network costs impact tenants. Additional survey research on tenant commuting patterns, in particular, could be helpful, including studies on: how tenants get to court, how far tenant jobs are from courts, and what methods of transportation tenants use in various cities.[43]

Finally, I conclude by noting that further research into long-term strategies to reduce evictions is needed. Trying to solve the eviction problem only after tenants receive an eviction notice may be too late. The underlying issue of high rent is left untreated if a broader, long-term strategy to reform housing markets (esp. to increase housing supply) is not undertaken. Advocates, voters, and policymakers seem increasingly willing to embrace a housing supply agenda, so researchers will have new data to monitor the impact of such reforms on rental markets. As such policy changes unfold, eviction researchers should monitor how eviction outcomes respond to these structural reforms.

References

Abramson, B. (2024). The Equilibrium Effects of Eviction Policies. https://boazabramson.github.io/assets/evictions_Jan2024.pdf

Armstrong, C. C., & Ryan Jr., C. J. (2024). Rural Renting: An Empirical Portrait of Eviction. *University of Cincinnati Law Review*, *93*, 1–39. https://scholarship.law.uc.edu/uclr/vol93/iss1/1

Baum-Snow, N., & Duranton, G. (2025). *Housing Supply and Housing Affordability* (tech. rep.). National Bureau of Economic Research. https://www.nber.org/papers/w33694.pdf

Bell, A., & Parchomovsky, G. (2004). A Theory of Property. *Cornell Law Review*, *90*, 531–616. http://scholarship.law.cornell.edu/clr/vol90/iss3/1

[43] Cell-tracking data, if available, might help pinpoint exact travel distances or times to court from tenant last location (e.g., job).

Bieretz, B., Burrowes, K., & Bramhall, E. (2020). Getting Landlords and Tenants to Talk. *Urban Institute*. https://www.urban.org/sites/default/files/publication/101991/getting-landlords-and-tenants-to-talk_3.pdf

Bratu, C., Harjunen, O., & Saarimaa, T. (2023). JUE Insight: City-wide Effects of New Housing Supply: Evidence from Moving Chains. *Journal of Urban Economics*, *133*, 103528. https://doi.org/https://doi.org/10.1016/j.jue.2022.103528

Buhler, S., & MacLean, I. (2024). Gendered Eviction in Saskatchewan. *Queen's University Law Journal*, *50*, 1. https://journal.queenslaw.ca/sites/qljwww/files/Issues/Vol%2050%20i2/1.%20Buhler.pdf

California Courts. (2024). *File Your Answer with the Court*. https://selfhelp.courts.ca.gov/eviction-tenant/respond-file

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/https://doi.org/10.1016/j.jeconom.2020.12.001

Caspi, A., & Rafkin, C. (2023). Legal Assistance for Evictions: Impacts, Mechanisms, and Demand. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4621983

Cassidy, M., & Currie, J. (2023). The Effects of Legal Representation on Tenant Uutcomes in Housing Court: Evidence from New York City's Universal Access Program. *Journal of Public Economics*, *222*, 104844. https://doi.org/https://doi.org/10.1016/j.jpubeco.2023.104844

Clair, M., Orozco, J., & Zhang, I. H. (2025). Spatial Burdens of State Institutions: The Case of Criminal Courthouses. *Social Service Review*, *99*(2). https://doi.org/https://doi.org/10.1086/734463

Collinson, R., Humphries, J. E., Kestelman, S., Nelson, S., van Dijk, W., & Waldinger, D. (2024). Equilibrium Effects of Eviction Protections: The Case of Legal Assistance. *Working Paper*. https://robcollinson.github.io/RobWebsite/rtc.pdf

Collinson, R., Humphries, J. E., Mader, N., Reed, D., Tannenbaum, D., & Van Dijk, W. (2024). Eviction and Poverty in American Cities. *Quarterly Journal of Economics*, *139*(1), 57–120. https://doi.org/https://doi.org/10.1093/qje/qjad042

Desmond, M. (2012). Eviction and the Reproduction of Urban Poverty. *American Journal of Sociology*, *118*(1), 88–133. https://doi.org/https://doi.org/10.1086/666082

Desmond, M. (2017). *Evicted*. Penguin Books.

Desmond, M., & Kimbro, R. T. (2015). Eviction's Fallout: Housing, Hardship, and Health. *Social Forces*, *94*(1), 295–324. https://doi.org/https://doi.org/10.1093/sf/sov044

Dougherty, C. (2024). America's Affordable Housing Crisis. *The New York Times*. https://www.nytimes.com/2024/03/27/briefing/affordable-housing-crisis.html

Elliott, R. (2025). The Eviction Crisis is Rampaging: Time to Implement Diversion Programs. *Journal of Dispute Resolution*, 77. https://scholarship.law.missouri.edu/jdr/vol2025/iss1/8

Elmendorf, C. S., Nall, C., & Oklobdzija, S. (2024). Do Housing Supply Skeptics Learn? Evidence from Economics and Advocacy Treatments. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4955033

Estes, M., & Nelson, K. (2026). Justice Divided, Justice Denied? The Effects of Court Rules on Eviction Outcomes in Los Angeles County. https://matt-estes.com/assets/papers/EstesNelson2026_JusticeDivided.pdf

Goldsmith, J., & Vermeule, A. (2002). Empirical Methodology and Legal Scholarship. *The University of Chicago Law Review*, *69*(1), 153–167. https://chicagounbound.uchicago.edu/uclrev/vol69/iss1/3

Goldsmith-Pinkham, P. (2026). Tracking the Credibility Revolution Across Fields. *NBER*, *Working Paper 35051*. https://www.nber.org/papers/w35051

Greiner, D. J., & Matthews, A. (2015). The Problem of Default, Part I. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2622140

Greiner, D. J., Pattanayak, C. W., & Hennessy, J. (2013). The Limits of Unbundled Legal Assistance: A Randomized Study in a Massachusetts District Court and Prospects for the Future. *Harvard Law Review*, *126*, 901. https://harvardlawreview.org/print/vol-126/the-limits-of-unbundled-legal-assistance-a-randomized-study-in-a-massachusetts-district-court-and-prospects-for-the-future/

Greiner, D. J., Pattanayak, C. W., & Hennessy, J. P. (2012). How Effective are Limited Legal Assistance Programs? A Randomized Experiment in a Massachusetts Housing Court. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1880078

Gyourko, J., & McCulloch, S. E. (2024). The Distaste for Housing Density. *NBER*, *Working Paper 33078*. https://www.nber.org/papers/w33078

Hoffman, D. A., & Strezhnev, A. (2023). Longer Trips to Court Cause Evictions. *Proceedings of the National Academy of Sciences*, *120*(2), e2210467120. https://doi.org/https://doi.org/10.1073/pnas.2210467120

Holland, P. W. (1986). Statistics and Causal Inference. *Journal of the American Statistical Association*, *81*(396), 945–960. https://doi.org/10.1080/01621459.1986.10478354

Khanna, K., Bertelsen, B., Olivella, S., Rossell Hayes, A., & Imai, K. (2024). *wru: Who are You? Bayesian Prediction of Racial Category Using Surname, First Name, Middle Name, and Geolocation* [R package version 3.0.3]. https://CRAN.R-project.org/package=wru

Kim, M., & Lee, H. (2025). Upzoning and Gentrification: Heterogeneous Impacts of Neighbourhood-level Upzoning in New York City. *Urban Studies*, *62*(10), 2009–2028. https://doi.org/https://doi.org/10.1177/00420980241298199

Klein, E., & Thompson, D. (2025). *Abundance*. Simon & Schuster.

Lanzas, S. (2024). Stay Housed Los Angeles: Safeguarding Tenants' Rights Beyond Rent Control. *Arizona Journal of International and Comparative Law*, *41*, 489. http://hdl.handle.net/10150/677067

Larson, E. (2006). Case Characteristics and Defendant Tenant Default in a Housing Court. *Journal of Empirical Legal Studies*, *3*, 121–144. https://doi.org/https://doi.org/10.1111/j.1740-1461.2006.00065.x

Lens, M. C., Nelson, K., & Gromis, A. (2020). The Neighborhood Context of Eviction in Southern California. *City and Community*, *19*(4), 912–932. https://doi.org/https://doi.org/10.1111/cico.12487

Los Angeles City Controller's Office. (2025). Eviction Notices. https://controller.lacity.gov/landings/evictions

Manski, C., & Pepper, J. (2018). How Do Right-to-Carry Laws Affect Crime Rates? Coping with Ambiguity Using Bounded-Variation Assumptions. *The Review of Economics and Statistics*, *100*(2), 232–244. https://doi.org/https://doi.org/10.1162/REST_a_00689

Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.

Mast, E. (2023). JUE Insight: The Effect of New Market-Rate Housing Construction on the Low-Income Housing Market. *Journal of Urban Economics*, *133*, 103383. https://doi.org/https://doi.org/10.1016/j.jue.2021.103383

Meyer, B. D., Wyse, A., & Logani, I. (2025). Life and Death at the Margins of Society: The Mortality of the US Homeless Population. *Review of Economics and Statistics*, 1–46. https://doi.org/https://doi.org/10.1162/rest.a.265

Mleczko, M., & Desmond, M. (2023). Using Natural Language Processing to Construct a National Zoning and Land Use Database. *Urban Studies*, *60*(13), 2564–2584. https://doi.org/https://doi.org/10.1177/00420980231156352

Monkkonen, P., Lens, M., O'Neill, M., Elmendorf, C., Preston, G., & Robichaud, R. (2024). Do Land Use Plans Affirmatively Further Fair Housing? Measuring Progress. *Journal of the American Planning Association*, *90*(2), 247–260. https://doi.org/https://doi.org/10.1080/01944363.2023.2213214

Mullen, L. (2021). *gender: Predict Gender from Names Using Historical Data* [R package version 0.6.0]. https://github.com/lmullen/gender

Nelson, K. (2023). The Political Determinants of Access to Justice. *Working Paper*.

Nelson, K., Gromis, A., Kuai, Y., & Lens, M. C. (2021). Spatial Concentration and Spillover: Eviction Dynamics in Neighborhoods of Los Angeles, California, 2005-2015. *Housing Policy Debate*, *31*(3-5), 670–695. https://doi.org/https://doi.org/10.1080/10511482.2020.1847163

Niblett, A., & Yoon, A. H. (2017). Unintended Consequences: The Regressive Effects of Increased Access to Courts. *Journal of Empirical Legal Studies*, *14*(1), 5–30. https://doi.org/https://doi.org/10.1111/jels.12140

NYC Gov. (2024). *NYC Office of Civil Justice FY24 Annual Report*. https://www.nyc.gov/assets/hra/downloads/pdf/services/civiljustice/OCJ_Annual_Report_2024.pdf

Office of Governor Gavin Newsom. (2024, September). Governor Newsom Signs Bipartisan Housing Package and Launches Prop 1 Homekey Initiative. https://www.gov.ca.gov/2024/09/19/governor-newsom-signs-bipartisan-housing-package-and-launches-prop-1-homekey-initiative/

Office of Mayor Karen Bass. (2022, December). Mayor Karen Bass Declares a State of Emergency on Homelessness. https://mayor.lacity.gov/news/mayor-karen-bass-declares-state-emergency-homelessness

Pew Research Center. (2022, January). A Growing Share of Americans Say Affordable Housing Is a Major Problem Where They Live. https://www.pewresearch.org/short-reads/2022/01/18/a-growing-share-of-americans-say-affordable-housing-is-a-major-problem-where-they-live/

Pew Research Center. (2024, September). Economic Ratings and Concerns. https://www.pewresearch.org/politics/2024/09/09/economic-ratings-and-concerns/#top-economic-concerns-food-and-consumer-prices-housing-costs

Pollock, J. (2024). Right to Counsel for Tenants Facing Eviction: Justification, History, and Future. *Fordham Urban Law Journal*, *51*, 1439. https://ir.lawnet.fordham.edu/ulj/vol51/iss5/4

Prescott, J. (2017). Improving Access to Justice in State Courts with Platform Technology. *Vanderbilt Law Review*, *70*(6), 1993–2050. https://scholarship.law.vanderbilt.edu/vlr/vol70/iss6/15

Rollet, V. (2025). Zoning and the Dynamics of Urban Redevelopment. *Working Paper*. https://vrollet.github.io/files/city_structure.pdf

Scherer, A. (2022). The Case Against Summary Eviction Proceedings: Process as Racism and Oppression. *Seton Hall Law Review*, *53*, 1. https://digitalcommons.nyls.edu/fac_articles_chapters/1537

Scherer, A. (2023). Stop the Violence: A Taxonomy of Measures to Abolish Evictions. *Fordham Urban Law Journal*, *51*, 1329. https://ir.lawnet.fordham.edu/ulj/vol51/iss5/2

Simshaw, D. (2022). Access to AI Justice: Avoiding an Inequitable Two-Tiered System of Legal Services. *Yale Journal of Law and Technology*, *24*, 150. https://yjolt.org/sites/default/files/simshaw_-_access_to_a.i._justice.pdf

Summers, N. (2020). The Limits of Good Law: A Study of Housing Court Outcomes. *University of Chicago Law Review*, *87*, 145. https://chicagounbound.uchicago.edu/uclrev/vol87/iss1/3

Summers, N., & Steil, J. (2025). Pathways to Eviction. *Law & Social Inquiry*, *50*(1), 129–169. https://doi.org/doi:10.1017/lsi.2024.23

Teresa, B. F., Howell, K. L., Suen, I.-S., Robinson, A., & Sabo, R. (2025). Moving from Crisis to Stability? The Success and Limits of an Eviction Prevention Program. *Housing Policy Debate*, *35*(3), 452–469. https://doi.org/https://doi.org/10.1080/10511482.2024.2368133

U.S. Department of Housing and Urban Development. (2024, July). HUD ACS Cost Burden Measure. https://www.census.gov/newsroom/press-releases/2024/renter-households-cost-burdened-race.html

Woodward, A. M. (2025). A Constitutional Right to Counsel in Eviction Proceedings: Solutions to Ohio's Housing Stability Crisis. *University of Cincinnati Law Review*, *93*, 920. https://scholarship.law.uc.edu/uclr/vol93/iss3/10

Chapter 2

JUSTICE DIVIDED, JUSTICE DENIED? THE EFFECTS OF COURT RULES ON EVICTION OUTCOMES IN LOS ANGELES COUNTY

2.1 Introduction

Landlords file more than 40,000 eviction cases against renter households each year in Los Angeles County (Dillon, Castleman, & Esquivel, 2023). The leading reason is non-payment of rent. In most LA neighborhoods, more than 90% of eviction notices are issued for non-payment of rent. Once filed, most cases end in eviction via a default judgment, where the court orders the tenant(s) to vacate the property and, if relevant, pay the landlord any unpaid rent. For tenants, the repercussions are increased housing instability, reduced earnings, reduced credit access, and increased hospital visits (Collinson et al., 2024).

With the expiration of COVID-19 eviction protections, state and local policymakers are increasingly exploring interventions to decrease eviction filings and keep tenants housed. California Gov. Gavin Newsom, for example, signed housing reforms in 2023 aimed at decreasing the number of evictions.[1] Local jurisdictions across the state have gradually begun to adopt more stringent tenant protections too. In the City of Los Angeles, Mayor Karen Bass has extended rent relief programs, stating: "We will do all we can to ensure that a wave of evictions does not hit our City as we continue confronting the homelessness crisis."[2]

Despite policymaker interest, the causes of eviction are not well-understood. Part of the problem is that data on California evictions is scarce and difficult to obtain, which has prevented researchers from fully identifying institutional mechanisms linking eviction filings with eviction case outcomes. This paper assembles the most comprehensive data on LA County evictions to-date. Our primary dataset for this paper consists of over 60,000 individual case docket records, which we use to better understand individual tenant outcomes, link to additional administrative data (e.g., assessor data), and isolate descriptive trends in LA County eviction cases. The docket records are used, for instance, to extract institutional details on the practice

[1] For media coverage, see Wiley (2023).

[2] Mayor (2023).

of eviction across LA County courthouses, including information on time-to-default, monetary judgments, total case times, and dozens of binary procedural variables (extracted via regular expressions). Additional institutional context (hearing times, parking fees, COVID-19 era moratoria, case misfiling) is also documented in the Appendices and incorporated into the analysis where possible.[3]

Using this novel data, we next study court assignment as an institutional determinant of eviction via default judgment in LA County. Default evictions occur when tenants fail to appear at court and lose *by default*. We quantify how the court assignment mechanism impacts the eviction default probability. Because LA County eviction cases are assigned to courthouses based on a unique spatial assignment rule, we are able to estimate the causal effect of courthouse assignment on default probability using a regression discontinuity design. The discontinuity design allows us to compare cases located across a courthouse boundary that are likely to be similar along unobservable dimensions.

For renters near a courthouse boundary, the paper shows that court assignment affects the eviction default probability. Our estimates for the local average treatment effect (LATE) of courthouse assignment range from 0.05–13.3 percentage points for seven different courthouse pairs. For example, for renters near the Santa Monica & Stanley Mosk courthouse boundary, the average treatment effect of being assigned to the Santa Monica courthouse is a 2.89 percentage point increase in the probability of default versus the Stanley Mosk courthouse. Implementing a variety of robustness checks and estimating additional outcome treatment effects at the boundary, however, we find both robustness and sensitivity to certain specifications.

To further isolate the causes of eviction defaults, we specify and estimate a spatial boundary mechanism model to determine how two mechanisms impact default probability. Specifically, we test how tenant costs[4] of getting to court and court procedures impact eviction default outcomes using boundary segment variation in local mechanism and outcome discontinuities. The impact of various design and estimation decisions in the "garden of forking paths" (Gelman & Loken, 2013) are explored, including varying: outcome measures, procedure measures, tenant

[3] For example, we use case time and money judgment to test for robustness as additional outcomes in our RDD design. Institutional details on parking fees, case hearing times, etc. are used to obtain a new measure of the tenant cost of getting to court in monetary terms (i.e., "generalized costs").

[4] As Prescott (2017) writes: "[T]he inability to access justice is rooted in something more physical, more mundane: the many and varied costs of getting to and physically using a brick-and-mortar courthouse."

cost measures, first-stage estimation procedure, bandwidth selection, the number of boundary segments, and other modeling choices.

Using many model specifications, we show that mechanism estimates are subject to specification, measurement, and estimation uncertainty. Nevertheless, we find model average and model median estimates of the effect of tenant costs (proxied by distance-to-court or measured in general monetary costs) to be small in standardized terms, although the sign of the effect is not consistent across models. Estimates of procedural effects are even more uncertain (wider model estimate interquartile ranges). In some cases average and median effect estimates align with the extant literature, however, the average and median effect sizes are small in general. Model estimates for each courthouse pair under various specification settings are provided and discussed.

Further studies or additional data may more precisely estimate the most important causal mechanisms determining eviction outcomes. Our work highlights some limitations, however, with using observational data to isolate causal mechanisms, showing that such studies face many specification choices that can change the magnitude and direction of mechanism effect estimates. The spatial boundary mechanism model here, nevertheless, may be useful in exploring how variation in outcomes and mechanisms along different boundary segments may identify mechanism coefficient effects in some settings under certain causal assumptions.

Our analysis informs ongoing policy debates. Typically, interest groups propose eviction reforms to give tenants more resources. For example, some tenant advocates endorse civil *Gideon*,[5] a right-to-counsel intervention (Baird, 2021). But providing tenants with legal representation in eviction court is costly. By contrast, this paper is concerned with institutional policy reforms. We show how courthouse assignment is an important institution with a disparate impact on tenants across courthouse boundaries. Because LA County has 23 active courthouses (including six[6] that previously heard eviction cases post-2010), policymakers can modify the courthouse assignment rule to minimize the cost of getting to court.[7]

[5]The proposal is a reference to *Gideon v. Wainwright* (1963), the landmark Supreme Court decision requiring states to provide defense counsel to indigent criminal defendants.

[6]Previous Eviction Courthouses (with year last heard eviction cases): Alhambra (2013), Beverly Hills (2012), Burbank (2013), Glendale (2013), Pomona (2017), Torrance (2013).

[7]We leave the optimal eviction courthouse assignment policy for future work.

Prior Literature

This paper is related to a growing eviction literature, that includes work in fields like sociology, law, urban planning, and public policy. Research from the Eviction Lab (Desmond, 2017), for example, discusses the demographic correlates of eviction, suggesting that defaults are randomly distributed.[8] In LA County, by contrast, Lens, Nelson, and Gromis (2020) and Nelson et al. (2021) demonstrate the importance of neighborhoods, housing markets, and spatial clustering to understanding eviction dynamics (the "geography of risk") in metropolitan areas. Larson (2006) similarly shows that case characteristics and neighborhood-level covariates are important determinants of a prevalent eviction outcome: default. This literature informs our research design. Intuitively, the regression discontinuity design disentangles correlation from causation by comparing cases located nearby so that observable and unobservable characteristics are similar across the courthouse boundary.

We also contribute to a nascent but growing literature on the causal determinants of eviction outcomes. Our work is related to Hoffman and Strezhnev (2023), who study the spatial determinants of Philadelphia evictions. Studying over 200,000 eviction cases in Philadelphia from 2005–2021, the authors use a selection-on-observables strategy with a variety of fixed effects to show that "a one hour increase in estimated travel time increases the probability of default by between 3.9 to 8.6 percentage points across different model specifications." In New York City and Chicago, Collinson et al. (2024) study the consequences of eviction using detailed administrative data. Using a judge-leniency design ("Judge IV"), the authors show that eviction increases homelessness, reduces earnings, and negatively impacts credit access.

Our work adds to this causal determinants of eviction literature by incorporating a regression discontinuity design in a sprawling Western city. This is important for two reasons. First, the quasi-experimental regression discontinuity design is said to have high degrees of internal validity (Lee & Lemieux, 2010). Second, it is important to study Western cities that differ in important ways from Philadelphia, NYC, and Chicago to test the generalizability of prior work. LA County may better represent Sunbelt cities in terms of spatial distribution of cases, transit infrastructure, commute time, etc. than other large cities. Moreover, the assignment process in LA County differs in significant respects from assignment in these three best-studied cities: for example, cases in Philadelphia are assigned to a single courthouse located at the intersection of the city's two primary public transit lines.[9] By contrast, LA

[8] See Desmond (2017) p. 358, n. 5.

[9] Cook County cases are heard at six different courthouses, although Chicago cases are heard

County cases are assigned to 1 of 11 courthouses by a unique neighborhood-zip code scheme. Like Hoffman and Strezhnev (2023), we show that distance does matter in LA County, although there is important geographic heterogeneity.

Finally, we contribute to the law and economics literature on proximity or access to court. Prescott (2017), for example, studies how the introduction of dispute resolution platforms impacted case times and case default rates. Like Prescott (2017), we show physical space affects case default rates.

The rest of the paper is organized as follows. Section 2.2 presents the institutional details necessary to understand eviction in Los Angeles County.[10] Section 2.3 describes the collected data, defines the outcome-of-interest, and documents descriptive findings. Section 2.4 describes the regression discontinuity design and results, including: identification, estimation, and local average treatment effects (LATEs). Section 2.5 specifies a spatial boundary segment mechanism model and estimates the effect of two potential causal mechanisms affecting eviction outcomes: tenant costs (proxied by distance-to-court or generalized monetary costs) and court procedures. Section 2.6 concludes. Further details can be found in the Appendices on data sources and collection, data cleaning, record linkage, descriptive findings, and other details.

2.2 Institutional Background

Although previous studies on eviction have focused on major metropolitan areas such as NYC and Chicago, fewer academic works concern evictions in California cities. Despite the consistently high volume[11] of eviction filings and policymaker concern, it is difficult to address CA eviction problems because empirical work is hampered by data availability problems. California counties do not uniformly record or report the number of cases filed by different types (e.g., eviction, small claims, and unlimited civil jurisdiction). The statewide eviction record sealing law is another obstacle. Under CA law, many eviction records are sealed to the public

downtown at the Daley Center. The suburban cases are assigned to one of the five suburban courthouse districts. The cities assigned to each district are listed at https://www.cookcountyclerkofcourt.org/divisions/suburban-districts-0. In New York City, eviction cases occur in seven courthouses determined either by County (Bronx, Kings, New York, Queens, Richmond) or two special eviction courthouses (Harlem and Red Hook).

[10] Certain institutional details cannot be treated fully here. We describe only the eviction details necessary to contextualize our identification strategy and interpret the results. For a more complete treatment of LA County eviction, including historical details, see Nelson (2023).

[11] See the Eviction Lab estimates in the top left panel of Appendix Figure B.13. The Eviction Lab-modeled estimates also show that LA filings constitute a roughly constant proportion of total statewide filings from 2000–2018.

unless certain conditions are met.[12] Simplifying, eviction records are sealed unless the landlord prevails at trial or the case ends in a default judgment against the tenant.

Due to these reporting and sealing issues, we study a particular eviction outcome: eviction via default judgment. Defaults can occur at two time points in an eviction lawsuit: first, if tenants do not file an Answer before the court-mandated deadline, and second, if tenants do not appear at an assigned court date. While the latter type is relatively rare, both types of default result in landlords winning and tenants losing their cases, oftentimes before their first court date. Because defaults are difficult to reverse, filing an Answer and attending eviction court dates are imperative.

There exist a myriad of obstacles, however, to getting to court. As argued most relevantly by tenants-plaintiffs in *Miles v. Wesley*, 801 F.3d 1060 (9th Cir. 2015):

> "[R]educing the number of courthouses handling unlawful detainer cases disproportionately impacts poor, disabled, and minority residents. ... [B]ecause individuals with disabilities and minorities are disproportionately renters who rely on public transportation, the closure of these courtrooms would have a disparate impact on these communities. [T]he importance of neighborhood court access is heightened in light of the expedited timeline of unlawful detainer actions, the fact that most low-income tenants are not represented by counsel, and the prospect that a default judgment could render a tenant homeless."

In response to *Miles*, the LA Superior Court system expanded the number of regional courthouses that process eviction cases, but many of the distance concerns cited in *Miles* continue to exist. Hoffman and Strezhnev (2023) confirmed anecdotal evidence in *Miles*, showing that longer travel times made tenants in Philadelphia, PA and Harris County, TX more likely to default than tenants with shorter travel times. The authors find that this "transit effect" disappeared when courts made virtual accommodations in eviction proceedings during the COVID-19 pandemic.

This growing literature focuses primarily on structural factors that shape the likelihood of tenants defaulting. To the extent that this research identifies institutional determinants of default, it focuses on COVID-19 era adaptations, rather than eviction case policies in less exceptional times. By contrast, our paper is the first to empirically assess how the idiosyncratic institutional process of court assignment

[12] See California Code of Civil Procedure §§ 1161.2, 1167.1.

(specifically, assigning eviction cases to regional courtrooms within a broader countywide court system) affects the likelihood of tenant default.

To do so, we collected comprehensive data on monthly eviction filings in LA County.[13] The top right panel of Appendix Figure B.13 shows the LA County filing volume data we collected via California Public Records Act (CPRA) requests to the LA Superior Courts (LASC). Our filing data augments the Eviction Lab Data (Appendix Figure B.13 top left panel) with exact filing volumes since 2000 and emphasizes the inclusion of 2019–2023 filings. Although eviction filings were low in 2020–2022 due to temporary eviction protections at the local, state, and federal levels, the number of eviction filings in 2023 exceeded the number of pre-pandemic filings in the 2017–2019 years.

During the same period, the time to complete an eviction case has been increasing. In the bottom panel of Appendix Figure B.13, the percentage of CA eviction cases disposed in 30 or 45 days decreased from 2014–2023. Eviction is a summary proceeding, but only 35% of the California evictions filed in 2023 were completed within 45 days. In Appendix B.3 and Appendix B.4, we show that LA County also experiences longer disposition and total case times.

In sum, LA County evictions are increasing in volume, take longer to complete, and are difficult to address due to lack of reliable empirical information for policymakers, legal officials, and researchers. This paper is the first to look at the relationship between institutional processes and default outcomes and represents the most comprehensive empirical description and causal analysis of how the LA County eviction process impacts eviction outcomes to-date. The remainder of Section 2.2 describes the LA County eviction system, including the eviction process and how eviction cases are assigned to different eviction courthouses.

Eviction Process

The eviction process begins when a landlord serves an eviction notice on a tenant. The vast majority of eviction notices are for unpaid rent: for example, in LA City, over 96% of the notices in 2023 were for non-payment of rent. Eviction notices typically give tenants 3 days (91% of eviction notices in LA City) to "cure" their lease breach by paying unpaid rent to the landlord. The amount of unpaid rent and number of eviction notices vary at the zip-code and building levels.[14]

[13] Trends plotted in Appendix B.6.

[14] See Nelson et al. (2021) for discussion of the spatial autocorrelation in eviction variables.

Next, the landlord initiates an eviction proceeding by filing an eviction lawsuit in the LA Superior Court system, which costs $240-$385.[15] After being notified of an eviction proceeding initiated against them,[16] tenants must file an Answer within 5 days (pre-2025)[17] with the court.[18] The filing fee for the tenant's Answer has varied over time, but the 2024 fee in eviction actions where the contested amount of rent is less than $12,500 is $225.[19]

If tenants do not file an Answer by the court-mandated deadline, landlords may petition the court to enter a default judgment against the tenant. Otherwise, after an Answer is filed, the court sets a trial date. Tenants who do not appear at trial will also receive a default judgment, but if they appear there is a judgment on the merits.[20] The judgment typically awards landlords any past due rent.

Following the court-issued judgment, the landlord may enforce the judgment by obtaining a writ of execution. The writ gives the sheriff permission to lock the tenant out of the premises. After obtaining the writ, the sheriff will serve the tenant a Notice to Vacate, which gives the tenant five days to move out. Five days after receipt of the Notice to Vacate, the sheriff will change the locks, forcing the tenant out of the residence. The typical process is represented graphically below in Figure 2.1.

[15] See nos. 11 and 14 in LASC (2024).

[16] Tenants are considered notified after being served the Summons and Complaint forms.

[17] The five days do not include weekends or holidays. Additionally, tenants may have longer to respond if they are improperly served. But see Assembly Bill 2347, which as of Jan 2025 gives tenants 10 days to respond.

[18] The tenant is not supposed to mail the Answer, as they will default if the Answer does not arrive. The official self-help page for California Courts strongly recommends against mailing the Answer, and instead says you should show up to file the Answer at the relevant court. See https://selfhelp.courts.ca.gov/eviction-tenant/respond-file.

[19] See no. 15 in LASC (2024). But note that there are fee waiver applications available.

[20] As explained in Engstrom et al. (Apr. 3, 2025), "tenants face high costs to participation: short response timelines, the need to make last-minute requests for time off work, high answer fees, and scarcity of low-fee legal help services. These costs influence whether a tenant chooses to participate in their case—even if they disagree on the merits." Nevertheless, benefits to court attendance may include reduced judgment amounts, additional time to move-out, or further bargaining gains. High tenant costs and low win probability, moreover, do not necessarily produce an eviction court paradox: as in the voting paradox literature (see, e.g., Riker and Ordeshook (1968)), probability-independent (i.e., non-instrumental) benefits from access, participation, and engaging in the legal process may influence tenant attendance decisions.

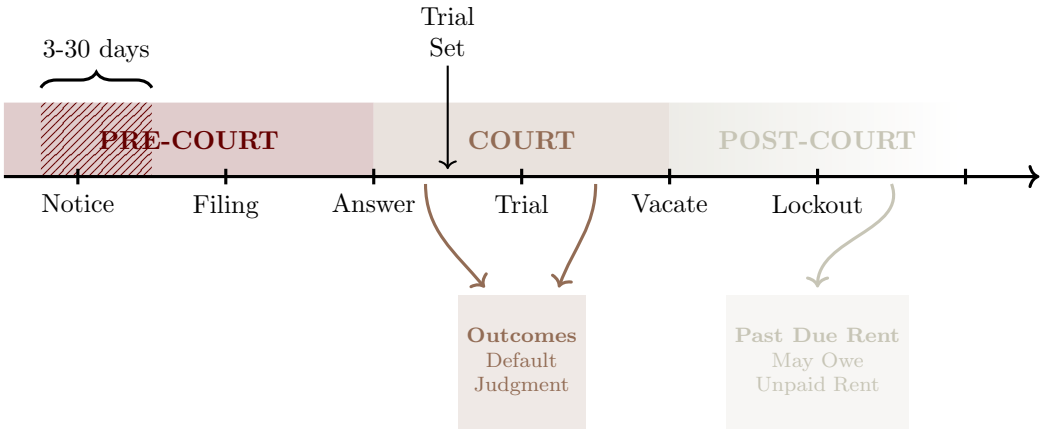

Figure 2.1: Eviction Timeline (Simplified)

This is a stylized timeline of events, which may not reflect the exact process for every eviction.[21] In particular, eviction moratoria and other COVID-19 protections altered the 2020–2022 process, e.g., giving tenants more time or limiting the landlord's ability to force the tenant to vacate. We document, moreover, in Appendix B.3 and Appendix B.4 that LA County time-to-disposition and total case times have recently increased, similarly to the statewide trends in Appendix Figure B.13.

Courthouse Assignment Rule: Map & History

According to LASC Local Rule 2.3(a)(2), eviction cases are assigned to courthouses based on a unique spatial mechanism. Under that rule, unlawful detainer actions:

> "must be filed in the courthouse serving the location and proper United States Postal Service zip code of the property in dispute using the Zip Code Table for Unlawful Detainer cases."

The eleven courthouses where eviction cases are filed include: Chatsworth, Compton, Governor George Deukmejian (Long Beach), Inglewood, Michael Antonovich Antelope Valley, Norwalk, Pasadena, Santa Monica, Stanley Mosk, Van Nuys East, and West Covina. Table 2.1 below illustrates how the courthouse assignment procedure works for the first few zip codes in Los Angeles County.

[21] The eviction timeline outlined above omits, for instance, any special motions that may arise.

Table 2.1: Zip Code Table for Unlawful Detainer Cases

Zip Code	City/Neighborhood	Modifier	Courthouse
90001	FLORENCE		Stanley Mosk
90001	HUNTINGTON PARK		Norwalk
90001	LOS ANGELES		Stanley Mosk
90002	FLORENCE		Compton
90002	LOS ANGELES		Compton
90002	LYNWOOD		Norwalk
90002	WATTS		Compton
90003	LOS ANGELES	North of Manchester	Stanley Mosk
90003	LOS ANGELES	South of Manchester	Compton
90004	LOS ANGELES		Stanley Mosk

Generally, the city-zip code pairs completely determine the assigned courthouse. However, in some zip codes the assignment is further determined relative to a particular street. For example, in Table 2.1 above, eviction cases arising in the 90003 zip code are assigned to the Stanley Mosk or Compton Courthouse if the tenant's address is north or south of Manchester Avenue, respectively. The full assignment map for LA County is given below in the left panel of Figure 2.2.

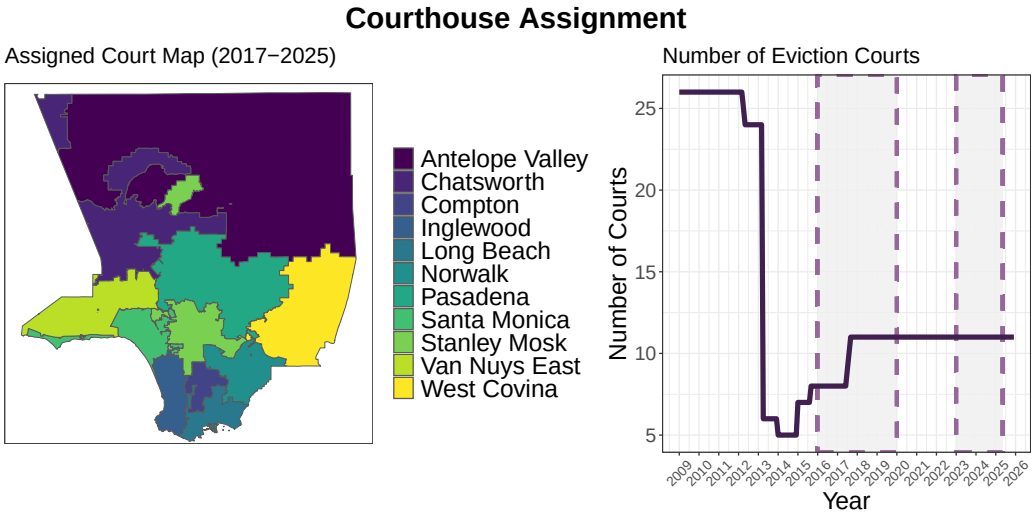

Figure 2.2: Courthouse Assignment

Note: The left panel shows the courthouse districts from 2017-2023, found by aggregating the units from the Simplified Zip Code Table (Appendix Table B.2). The right panel is the number of eviction courts from 2009-2024. The Catalina Courthouse is omitted from the left panel map and the analysis. In the right panel, LASC data has been collected for the shaded gray years within the dashed purple boxes. Note that 2020–2022 eviction data is unobserved due to local, state, and federal eviction moratoria.

Over time, the number of courthouses that hear eviction cases in LA County has fluctuated. As seen in the right panel of Figure 2.2 above, the number of courthouses that hear eviction cases has ranged between a low of five courthouses (2014) and a high of twenty-six courthouses (2009-2011). This paper focuses on the full years with 11 courthouses (2018–2024), which is the long horizontal line in the right panel of Figure 2.2.[22] See Appendix B.1 for more details on assignment.

2.3 Data and Outcomes

This section describes the data collection, cleaning, and linking (inc. geocoding) processes. We also define the main outcome-of-interest (default eviction) and briefly note how the data can be used to document descriptive trends in eviction cases.

Docket Records: Geocoding & Assessor Linkage

The first and most important dataset we collected for this paper includes tens of thousands of individual docket records in eviction cases scraped from the LA Su-

[22] We focus on the primary 11 courthouses, although technically the Catalina Courthouse also hears eviction cases for Santa Catalina Island. Unless otherwise stated, we generally do not consider this court because it has no observed evictions. We therefore omit it from the Figure 2.2 plots.

perior Courts from 2016-2023.[23] From the individual-level text files containing all publicly-available docket information, we use regular expression (regex) to extract detailed information for each individual case, including the following covariates: landlords (plaintiffs), tenants (defendants), attorneys (if applicable), address information, monetary awards, judge, courthouse, and dynamic case timing information. Particularly important is the address information, which includes the full address (i.e., street number, street, city, state, zip code) for the majority of cases across years.

Using the individual-level case address information, we then geocoded each address using the Mapbox API. The Mapbox API returns probabilistic matches with a matched address, latitude/longitude, and the match "confidence" score. In the analyses that follow, we consider only cases for which the Mapbox API returns a "High" or "Exact" match for the address we extracted from the individual docket records.

Using the geocoded addresses, we then linked our eviction dataset to information from the LA County Assessors Office. This gives us building-level covariates for each address and includes, for example, the age of the building, the number of units, and number of bathrooms. Many of these covariates serve as important proxies for tenant income or tenant rent, which are generally unobservable. Note that, for each year, we use assessor data on the number of units in a building to infer the number of units in that building for which we do *not* observe default for one of our outcome measures (imputed unit defaults outcome).

See Appendix B.2 for more details on data collection, including information on the case availability, the FOIA data on the total number of evictions in a courthouse each month, sampling rates in Appendix Figure B.1, geocoding information in Appendix Figure B.2, and Appendix B.8 for discussion of misfiling.

Outcome: Eviction Default

The main outcome of interest in this paper is eviction default. Using the docket and assessment data, we measure the number of eviction defaults each year at each address. The primary outcome is a binary variable that reflects whether we observe

[23] Unless otherwise noted, we exclude 2020-2022 from all analysis and discussion because of local, state, and national eviction moratoria. Moratoria are further described in Appendix B.2.

a rental unit i experience an eviction default:[24]

$$\text{eviction}_i = \mathbf{1}\{\text{unit } i \text{ is evicted via default}\}.$$

We impute non-default status to apartment units in a building each year $t \in 2018, 2019, 2023, 2024$ without an observed default and pool across years to create our analysis dataset. In the robustness checks, we discuss an alternative way of using the docket record data that does not involve imputing zeroes to unobserved apartment units.

In the analysis that follows, we also make additional restrictions on the docket records used. To avoid biases from geocoding, we first restrict the data to records with "high" or "exact" confidence matches from the Mapbox API. We also restrict analysis to residential case types with "default" status in the years 2018, 2019, 2023, and 2024. As discussed above, the year restriction is to only use full years with scraped cases under the current filing rule (which went into effect in Late 2017). Finally, note also that we focus on the sub-population of buildings that have at least one observed eviction (2018–2024). Other buildings that appear in the LA County assessor data but do not match an address in the scraped eviction data are excluded: assessed properties for which we do not observe an eviction may differ from properties with evictions in important but unobservable respects, including being unoccupied or unavailable to rent.

Docket Records: Descriptive Findings

The docket records are useful for empirically assessing eviction trends. For example, the data can be used to assess eviction case time trends (time-to-default or total case time), money judgment trends, or case misfiling rates. In Appendix B.3 and Appendix B.4, we show time-to-default and total case time are higher post-pandemic. This is consistent with the statewide trend toward longer case timelines, shown in the bottom panel of Appendix Figure B.13. We also document misfilings, where cases are heard in a courthouse different from the assignment rule courthouse. We do not find evidence that this is attributable to strategic landlord behavior in Appendix B.8.

[24] Note that the binary nature of this outcome will allow us to interpret the conditional expectation $\mathbb{E}[\text{eviction}_i | X_i = x]$ as the yearly probability of default eviction for units i that are x distance from the courthouse district boundary conditional on being an address with at least one rental unit with eviction default.

2.4 Regression Discontinuity Design: Methods & Results

This section explains our regression discontinuity design (RDD), which exploits variation stemming from the discontinuous assignment of cases to courthouses based on spatial location. We present results estimating the local average treatment effect (LATE) for seven different courthouse pairs across LA County.

Continuity-Based Sharp Regression Discontinuity Design

In what follows, we consider the "sharp" regression discontinuity design under continuity assumptions.

Setup & Identification

To begin, we assume a potential outcomes framework. Let Y_{0i} and Y_{1i} be the counterfactual outcomes for an individual i without and with treatment, respectively. Let the "forcing variable" be X_i be the distance to the boundary C. The binary treatment W_i is determined as follows:

$$W_i = \mathbf{1}\{X_i \geq C\}.$$

The observed outcome is:

$$Y_i = (1 - W_i)Y_{0i} + W_i Y_{1i}.$$

Define further the counterfactual conditional expectation functions (CEF):

$$\mu_g(x) = \mathbb{E}[Y_g | X = x]$$

for $g = 0, 1$. In the regression discontinuity design, the potential outcome CEFs are assumed continuous in the forcing variable X. Because W_i is a deterministic function of X_i note that ignorability (conditional mean independence) holds:

$$\mathbb{E}[Y_g | X, W] = \mathbb{E}[Y_g | X]$$

for $g = 0, 1$.

The estimand of interest in a sharp RDD is the treatment effect at the cutoff, or the local average treatment effect (LATE), defined as:

$$\tau_C = \mathbb{E}[Y_1 - Y_0 | X = C].$$

In our empirical setting, we have a treatment effect for each pair of courthouse districts that share a boundary C. That is, for each pair of courthouse districts, we

define treatment $W_i = 1$ if unit i is in the first courthouse district and $W_i = 0$ if unit i lies in the other courthouse district. For example, when comparing buildings near the Pasadena and West Covina courthouse boundary, we set $W_i = 1$ if unit i is assigned to the Pasadena courthouse and $W_i = 0$ if assigned to the West Covina courthouse.

Under these assumptions, the LATE is identified by the limits of the observable CEFs:

$$\tau_C = \lim_{x \to C^+} \mathbb{E}[Y|X = x] - \lim_{x \to C^-} \mathbb{E}[Y|X = x].$$

Estimation & Inference

Estimating the limits that identify the causal estimands above is usually accomplished non-parametrically. Below, we use the local linear estimator with uniform kernel as in Imbens and Lemieux (2008). These estimates are given by obtaining linear fits locally, i.e., within some bandwidth h_x on either side of the cutoff C:

$$\min_{\alpha_l, \beta_l} \sum_{i: X_i \in (C - h_x, C)} (Y_i - \alpha_l - \beta_l (X_i - C))^2 \quad \text{and} \quad \min_{\alpha_r, \beta_r} \sum_{i: X_i \in [C, C + h_x)} (Y_i - \alpha_r - \beta_r (X_i - C))^2$$

For each courthouse pair, we then estimate τ_C as:

$$\hat{\tau}_C = \hat{\alpha}_r - \hat{\alpha}_l.$$

An important consideration here is the bandwidth selection procedure. We present the LATE estimates for many possible bandwidth choices up to some distance from the boundary (e.g., up to 5km from the boundary). We also give the robust point estimates and CIs from the bandwidth selection procedure used by the `rdrobust` package (Calonico et al., 2023).

Estimating the LATEs

For each courthouse pair, we begin by considering only cases close to the boundary. To estimate the LATE at the courthouse boundaries, we consider only cases within a 5km buffer zone around the boundary. The cases we consider for Pasadena and West Covina, for instance, are shown in Figure 2.3. Buildings with observed defaults within 5km of the boundary are plotted in the left panel, with Pasadena defaults in red and West Covina defaults in blue. The location of the courthouses are also labeled as large points. The Pasadena courthouse is further from near-boundary evictions than the West Covina courthouse, which is visually shown in the left panel

map and confirmed in the right panel histograms. The right panel histograms show the distance-to-court that cases within the 5km boundary buffer face in getting to court.

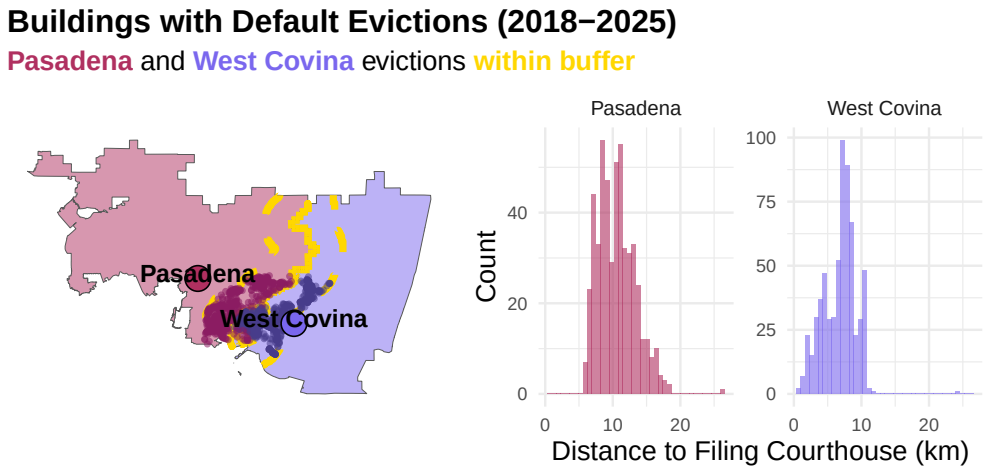

Figure 2.3: Pasadena & West Covina Map

Note: The left panel plots the boundary between Pasadena (red region) and West Covina (blue region) eviction districts. The eviction boundary is the solid yellow line, and cases within a 5km buffer (dashed yellow lines) are plotted as points: Pasadena evictions shown in red and West Covina evictions are shown in blue. The large labeled points are the Pasadena and West Covina courthouses. The right panel shows distance-to-court histograms for these eviction cases, with the Pasadena histogram in red and the West Covina histogram in blue.

Using these cases, we estimate local average treatment effects (LATEs)

$$\tau_C = \mathbb{E}(Y_1 - Y_0 | X = C)$$

at different boundaries C for the courthouse pairs below. In Figure 2.4, we report the LATE estimates from the robust bandwidth selection procedure in Calonico, Cattaneo, and Titiunik (2014) using companion software Calonico, Cattaneo, and Titiunik (2015) in the lower-right table. The courthouse pair plots show default probabilities as a function of distance to the boundary, with two lines of fit on each side of the boundary: a linear fit (dashed lines) and quartic fit (solid lines).

The estimated fits (within 5km of each boundary) are plotted by courthouse pair and `rdrobust` estimates are provided in the lower-right table in Figure 2.4. The absolute robust LATE estimates range from 0.05–13.3 percentage points. In two courthouse pairs, the robust LATE estimate is not significant: Chatsworth & Van Nuys and Pasadena & Stanley Mosk. In the remaining pairs, the robust LATE estimate is significant but varies in size. The smallest significant effect is 2.70 percentage points (Compton & Stanley Mosk), whereas the largest significant effect is -13.34 percentage points (Pasadena & Stanley Mosk).

Because eviction records are incomplete legal documents, they do not contain all relevant individual information on landlords and tenants. Nevertheless, using a variety of building-level, neighborhood-level, and individual-level data we assess the plausibility of the causal identifying continuity assumption on pre-determined covariates near the boundary of each courthouse pair. Building-level data (total value and taxable value) comes from LA County Assessor property records, demographic covariates (gender, race, median income, rent) are imputed to each record using census tract Census/ACS variables, and additional demographic (predicted race, predicted gender) and representation covariates are estimated using the docket records and regular expressions.

We show covariate continuity tests (estimated by `rdrobust`) for each courthouse pair boundary in Appendix B.9. Some courthouse pairs satisfy most of the continuity checks, with no significant treatment effects in many pre-determined covariates (robust p-values included in final columns of each courthouse pair table). In the Pasadena and West Covina pair, for example, we find no demographic (ACS-imputed or predicted), imputed income/rent, or representation continuity violations, although the Assessor property values are discontinuous. Some continuity test violations across the entire courthouse boundary may be spatial in nature, given that we are using cases across the entire courthouse pair boundary. Further tests localized to spatial boundary segments may be performed where more data are available, and we incorporate covariate balance conditions into the mechanism model that follows where possible. In other words, we include models in the mechanism estimation section that condition on covariate balance at spatial segments along the boundary, which may be expected to improve counterfactual comparisons. Further details are contained in Appendix B.9.

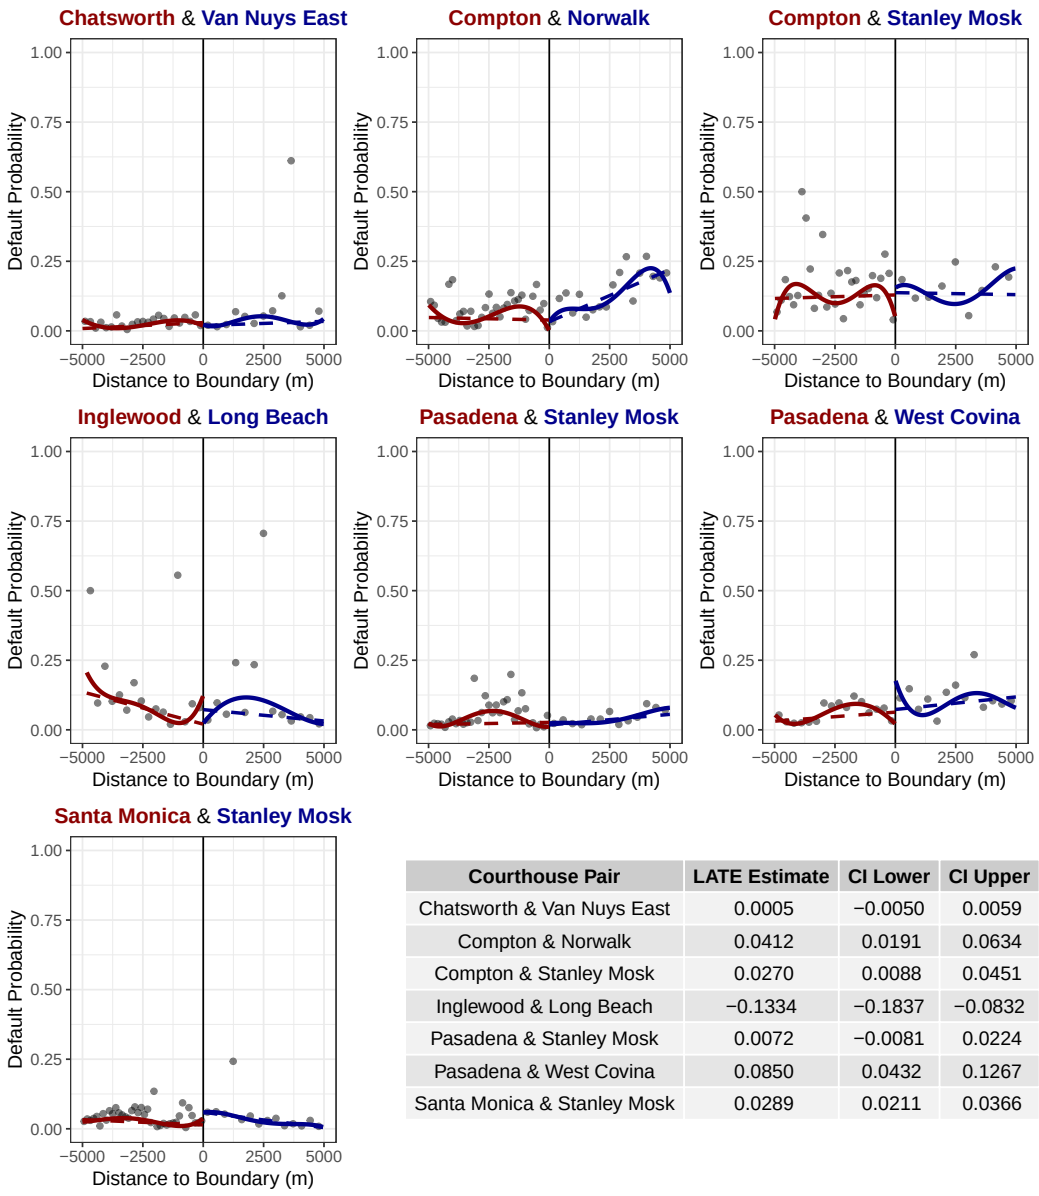

Courthouse Pair	LATE Estimate	CI Lower	CI Upper
Chatsworth & Van Nuys East	0.0005	−0.0050	0.0059
Compton & Norwalk	0.0412	0.0191	0.0634
Compton & Stanley Mosk	0.0270	0.0088	0.0451
Inglewood & Long Beach	−0.1334	−0.1837	−0.0832
Pasadena & Stanley Mosk	0.0072	−0.0081	0.0224
Pasadena & West Covina	0.0850	0.0432	0.1267
Santa Monica & Stanley Mosk	0.0289	0.0211	0.0366

Figure 2.4: LATE Estimates by Courthouse Pair (`rdrobust`)

Note: The LATE estimates ($\hat{\tau}_C$) at each courthouse pair boundary using the optimal bandwidth selection procedure in the `rdrobust` package (Calonico et al., 2023). The default quartic polynomial (solid line) and linear (dashed line) fits are plotted for each courthouse separately. The gray points are evenly-spaced binned means using the `rdplot()` function. The robust point estimates are reported (with robust CIs) in the table for each courthouse pair.

Robustness Checks

Next, we test the sensitivity of the default LATE estimates to various specification, design, and encoding decisions. The estimates from alternative specifications discussed here are included in Table 2.2.

Baseline. The first column in the robustness check Table 2.2 are the baseline specification results reported above, which are reproduced here for comparison purposes. We augment those estimates here with p-values and Bonferroni corrected p-values to account for testing seven courthouse pairs.

Non-Answers Only. The second column of Table 2.2 considers docket records without a recorded answer only. Defaults with and without a recorded answer may differ in important ways, yet many of the LATE treatment effect estimates are similar whether we consider only non-answer defaults or all defaults. For example, estimates from the Compton and Norwalk, Inglewood and Long Beach, and Santa Monica and Stanley Mosk comparisons remain significant (even with a Bonferroni correction) and of similar magnitude as the Baseline specification in Column (1).

Uniform Zip Codes Only. The third column of Table 2.2 subsets the data to zip codes with one assigned courthouse. As discussed in Appendix B.1, the filing rule encoding process from shapefiles and neighborhood boundaries may not be exact. Accordingly, in this specification we rely only on observations from zip codes where there is one unambiguous filing courthouse for all eviction cases within that zip code. In some cases, this does not appear to qualitatively change the LATE estimates: for example, the Compton and Norwalk and Santa Monica and Stanley Mosk point estimates from Column (3) are similar to the baseline specifications in Column (1). However, the zip code specification changes some estimates quite drastically, including one where the estimates lie outside the probability bounds (Inglewood and Long Beach) and some specifications where the sign of the estimate changes (e.g., Pasadena and Stanley Mosk comparison).

Assessor Match Quality. The fourth column of Table 2.2 subsets data to rows with higher quality Assessor match scores. Because we use the number of units in each apartment building and impute zeroes to units without observed defaults, this could bias estimates if the number of apartment units at an address is incorrectly matched from LA County Assessor data. The estimates are obtained from excluding cases where the Jaro-Winkler (JW) string similarity score, which measures differences between the observed address and the assessor address, falls below a threshold. The distribution of JW string similarity scores is shown in Appendix Figure B.20.

Here, this results in some changes to the estimates. One estimate changes sign (Pasadena and Stanley Mosk) and one is no longer significant (Inglewood and Long Beach). On the other hand, certain courthouse pair estimates are similar directionally and in magnitude to the Baseline estimates, and they remain significant: for example, Compton and Norwalk, Pasadena and West Covina, and Santa Monica and Stanley Mosk estimates are similar in magnitude and remain significant.

Non-Imputed Outcomes/Case Default Flag. The fifth column tests whether non-default imputation to units in observed apartment buildings is meaningful. Instead of using all units in observed buildings, this specification uses all observed eviction docket records and flags cases with "default" status using regular expressions. Then, using only these records and no others, the records with default status are encoded as 1 and those without are encoded as 0, which does not require imputing zeroes to unobserved apartments in observed apartment buildings. It also changes the interpretation of the local treatment effect: the RDD estimates the effect of court assignment on default status among eviction docket records.

The estimates from these specifications are included in Column (5) of Table 2.2. Most estimates are not significant (five pairs) or change sign from the Baseline specification (one pair). The remaining pair (Copmton and Norwalk) is significant and directionally consistent, but suggests a larger effect: assignment to Compton courthouse increases default status by 12.6pp over assignment to Norwalk courthouse, versus the 4.1pp estimate from the imputed units Baseline specification.

Bandwidth Sensitivity. Finally, we test the sensitivity of the LATE estimates to the bandwidth. In Appendix Figure B.19, we plot the local linear LATE estimates as a function of the specification bandwidth h_x. In a few cases, we see bandwidth-dependence in the magnitude and direction of the effect. For example, the magnitude and sign for the Inglewood and Long Beach LATE estimate as we decrease the bandwidth from 5000 meters to 0 meters changes sign a few times and magnitude near the `rdrobust` optimal bandwidth (vertical dashed black line). On the other hand, some of the LATE estimates are robust to the `rdrobust` selected bandwidth: for example, the Santa Monica and Stanley Mosk and Chatsworth and Van Nuys East LATE estimates are positive and roughly constant across bandwidths h_x.

	(1) Baseline	(2) Non-Answers	(3) Zips	(4) Assessor Match	(5) All Outcomes
Chatsworth and Van Nuys East					
Treatment Effect	0.000	0.001	-0.097	0.000	0.103
95% CI	(-0.005, 0.006)	(-0.003, 0.006)	(-0.137, -0.057)	(-0.008, 0.009)	(-0.004, 0.209)
p-value	0.864	0.602	0.000	0.972	0.058
Bonferroni p-value	1.000	1.000	0.000	1.000	0.408
Compton and Norwalk					
Treatment Effect	0.041	0.037	0.052	0.057	0.126
95% CI	(0.019, 0.063)	(0.015, 0.059)	(0.033, 0.072)	(0.029, 0.085)	(0.001, 0.250)
p-value	0.000	0.001	0.000	0.000	0.048
Bonferroni p-value	0.002	0.007	0.000	0.000	0.339
Compton and Stanley Mosk					
Treatment Effect	0.027	0.030	0.163	0.128	-0.020
95% CI	(0.009, 0.045)	(0.007, 0.054)	(0.147, 0.178)	(0.117, 0.140)	(-0.168, 0.127)
p-value	0.004	0.013	0.000	0.000	0.787
Bonferroni p-value	0.025	0.088	0.000	0.000	1.000
Inglewood and Long Beach					
Treatment Effect	-0.133	-0.133	1.592	-0.107	-0.191
95% CI	(-0.184, -0.083)	(-0.184, -0.083)	(0.643, 2.541)	(-0.529, 0.316)	(-0.421, 0.039)
p-value	0.000	0.000	0.001	0.621	0.104
Bonferroni p-value	0.000	0.000	0.007	1.000	0.728
Pasadena and Stanley Mosk					
Treatment Effect	0.007	-0.044	-0.040	-0.040	0.017
95% CI	(-0.008, 0.022)	(-0.054, -0.033)	(-0.052, -0.029)	(-0.050, -0.030)	(-0.064, 0.099)
p-value	0.356	0.000	0.000	0.000	0.676
Bonferroni p-value	1.000	0.000	0.000	0.000	1.000
Pasadena and West Covina					
Treatment Effect	0.085	0.048	0.303	0.117	-0.005
95% CI	(0.043, 0.127)	(0.011, 0.085)	(0.152, 0.453)	(0.077, 0.157)	(-0.209, 0.200)
p-value	0.000	0.011	0.000	0.000	0.963
Bonferroni p-value	0.000	0.075	0.001	0.000	1.000
Santa Monica and Stanley Mosk					
Treatment Effect	0.029	0.031	0.058	0.030	-0.103
95% CI	(0.021, 0.037)	(0.024, 0.038)	(0.040, 0.075)	(0.023, 0.036)	(-0.181, -0.026)
p-value	0.000	0.000	0.000	0.000	0.009
Bonferroni p-value	0.000	0.000	0.000	0.000	0.061

Table 2.2: RDD LATE Robustness Checks

Note: Column (1) is our baseline specification described in the previous section. Column (2) LATE estimates are for default outcomes without answers only. Column (3) LATE estimates consider only zip codes with unambiguous LASC filing rules. Column (4) LATE estimates are from subsetting the data to rows with high Assessor data match probabilities. Column (5) LATE estimates are from not imputing non-default outcomes to unobserved apartment units.

Additional Outcome Variables and LATE Estimates

Next, we include additional LATE estimates using further outcome variables. The effect of court assignment on the following case outcomes is estimated here: total case times, time-to-default (measured two ways, see), and total money judgment amounts. Unlike defaults, we do not have to rely on imputation to other apartment units to obtain variation in these outcome variables. Consequently, we are able to use only the observed docket records to estimate LATEs for these four outcomes.

Following the same estimation procedure, we obtain estimates of the effect of court assignment on these four outcomes in Appendix Table B.14. For each of the seven courthouse pairs and four outcomes of interest, Appendix Table B.14 gives the mean outcome for cases within the `rdrobust` selected bandwidth around the boundary. Means are measured in days for Columns (1)-(3) and dollars for Column (4). Robust treatment effect estimates, 95% confidence intervals, and p-values are also provided for each outcome and courthouse pair. Appendix Table B.15 provide additional details and specifications, including: incorporating covariates (gender, race, income); Bonferroni corrected p-values; and other design information for each specification (e.g., bandwidths, number of observations).

In general, estimates for the additional outcome variables are insignificant. Although this could be due to smaller sample sizes around the boundaries, we find, for example, only two significant estimates of court assignment on time-to-default, in the Compton and Norwalk and Santa Monica and Stanley Mosk pairings. However, this appears to be non-robust to the time-to-default measure: that is, the estimates in columns (2) and (3) in Appendix Table B.14 are not significant using both time-to-default measures (discussed in Appendix B.3). The only remaining significant effect is the money judgment amount in the Compton and Norwalk pair, with higher average money judgments for cases assigned to the Compton courthouse. Otherwise, many outcome averages appear similar across courthouses near the boundaries: for instance, time-to-default averages across courthouse pairs are similar, with only a few days difference in time-to-default. Some total case time averages differ (e.g., 270 days in Chatsworth vs 175 days in Van Nuys East), but the differences are not statistically significant. Ongoing differences in average case times may be further investigated, to explore whether observable case, demographic, and neighborhood factors explain total case times in parts of LA County.

2.5 Testing Mechanisms: Distance-to-Court and Procedures

This section further investigates why default gaps emerge at courthouse boundaries. Because assignment to court is a large bundle of treatments, isolating the precise mechanism determining outcomes requires additional causal assumptions and counterfactual modeling. Here, we test the impact of two particular mechanisms which operate discontinuously at courthouse boundaries as part of a bundle of assigned court treatments: tenant costs (proxied by tenant distance-to-court or measured as general monetary costs) and court procedures (measured by the first principal component of a procedure dimension reduction analysis).[25] Using of mechanism measures, we specify a spatial mechanism model and use variation in outcome and mechanism gaps across boundary segments to estimate common tenant cost and procedure effects. The model allows for spatial heterogeneity across boundary segments in trends and levels, but assumes local continuity at each segment apart from the two studied mechanisms.

Although there is a literature on the impact of tenant costs on defaults and a literature on the importance of eviction court procedures (see Section 3.1 above), this paper quantifies the relative impact of these two mechanisms. Of course, tenant costs and court procedures may impact litigants in different ways at different case phases. Although it may be unlikely that procedural differences influence default outcomes at early stages, known differences in courthouse procedures may nevertheless affect tenant expectations and, thereby, cause defaults. Similarly, distance-to-court may affect certain stages of litigation more acutely and may shape tenant expectations. The model studied here allows for the possibility of procedural gaps or distance-to-court gaps affecting default gaps near courthouse boundaries.

Procedural Index and Principal Component Analysis (PCA)

The first step in determining the relative impact of procedure vs tenant costs on default outcomes is obtaining a measure of court procedures. To do so, we extract dozens of procedural terms using regular expressions ("regex") from docket records and encode them as binary variables. For example, we encode that a "Demand for Jury Trial" is requested if that term appears in the docket records. The list of binary procedural variables we extracted is recorded in Appendix Table B.3.

There are differences in procedures across courthouses. In Appendix Table B.4, we show the averages for our binary procedure variables across the 11 courthouses.

[25] The procedure and tenant cost mechanisms correspond roughly to the "cost-focused reforms" and "compliance-focused reforms" in Engstrom et al. (Apr. 3, 2025).

The table shows that the appearance of certain procedural terms differs across courthouses. Whether the difference in means across courthouses is significant is tested by likelihood ratio (LR) tests from logistic regression of procedure variables on courthouse indicators, with (gender, race, income) and without controls. The null hypothesis that procedure means are equal across the courthouses is rejected (final column, red-colored p-values) for several procedural variables.

To reduce the dimensionality of the procedural measure, we next turn to a principal component analysis ("PCA"). Because the procedures are correlated[26] and too numerous for causal mechanism analysis, we combine them using PCA into a single procedural index. The idea is to obtain a procedural index that captures the largest share of variation across cases. The resulting first component (Procedure Index) is rescaled from 0 to 100, with higher values representing more procedural activity or procedural intensity.

Three PCA specifications using different sets of procedural variables are used in the mechanism specifications. Histograms of the procedure index by courthouse are included for each procedural variable measure in Appendix Figure B.14, Appendix Figure B.15, and Appendix Figure B.16. The main components of each index vary across specification, but the shape of the procedure index distributions does not vary significantly across courthouse in the three specifications. This suggests that the Procedure Index, i.e., the first principal component which explains the largest share of procedural variation, is similar across courthouses. Nevertheless, the procedure index may differ in cases near the courthouse boundaries and, as a mechanism, cause the local gaps in default outcomes discussed above.

See Appendix B.7 for further discussion of institutional details and analysis of court procedures, including discussion of hearing start times, parking fees, and further detailing the 39 binary procedural variables analyzed from the docket records.

Mechanism Model

Next, we specify a spatial boundary mechanism model. The model allows for local discontinuities in outcomes and the two mechanisms at each spatial boundary segment, but otherwise assumes local continuity at each boundary segment. Then, using variation across boundary segments in outcome and mechanism discontinuities, the effect of tenant costs and procedure is identified assuming a common coefficient on each mechanism across boundary segments for each courthouse pair. See Appendix

[26] The correlation may result, e.g., from the intensity or quality of litigation activity.

B.11 for a full simulation example further illustrating the model, including specifying an example data generating process, identification, and estimation details using the segment specific discontinuities.[27]

Denote the local discontinuity in default outcomes Y at the boundary segment s as follows:

$$\Delta Y(s) = \lim_{x \to c^+} \mathbb{E}[Y|X = x, S = s] - \lim_{x \to c^-} \mathbb{E}[Y|X = x, S = s]$$

and similar jumps or discontinuities in the two mechanisms, distance-to-court D and procedure P, at each segment s which we denote:

$$\Delta D(s) = \lim_{x \to c^+} \mathbb{E}[D|X = x, S = s] - \lim_{x \to c^-} \mathbb{E}[D|X = x, S = s]$$
$$\Delta P(s) = \lim_{x \to c^+} \mathbb{E}[P|X = x, S = s] - \lim_{x \to c^-} \mathbb{E}[P|X = x, S = s]$$

We assume, locally around the boundary, a linear structural form of default potential outcomes:

$$Y(X, S, D, P) = \mu(X, S) + \beta_D D + \beta_P P + \epsilon$$

where β_D, β_P are assumed to be the same across boundary segments s and $\mu(X, S)$ is a continuous function of X at each boundary segment. This model allows baseline levels and smooth spatial trends in outcomes to vary flexibly across boundary segments. But it rules out any discontinuous direct effects at the boundary segment s other than via the two mechanisms D and P and assumes common or shared coefficients on D and P across segments.

Using these assumptions, we can obtain for each boundary segment s a differences equation by taking limits from the two sides of the boundary within a segment s:

$$\Delta Y(s) = \beta_D \Delta D(s) + \beta_P \Delta P(s)$$

where we used the identifying restriction that $\lim_{x \to c^+} \mu(x, s) = \lim_{x \to c^-} \mu(x, s)$ for each segment s. See Appendix B.11 for additional details.

Then, with many segments $s = 1, 2, \ldots, \bar{S}$, we can identify β_D, β_P with cross-segment variation in the mechanism discontinuities via a two-stage procedure. First, for each segment s we estimate discontinuities or jumps in outcomes and mechanisms

[27] The example incorporates spatial heterogeneity (across boundary segments), differences in levels and trends at each boundary segment (provided trends other than the mechanisms are continuous), and discontinuities in mechanisms along the boundary provided there is a shared coefficient on each mechanism term across boundary segments.

using local comparisons: $\widehat{\Delta Y(s)}, \widehat{\Delta D(s)}, \widehat{\Delta P(s)}$. We estimate the local jumps at each segment two ways in what follows, local linear regression and local differences in means. Second, we estimate the regression:

$$\widehat{\Delta Y(s)} = \beta_D \widehat{\Delta D(s)} + \beta_P \widehat{\Delta P(s)} + u_s$$

across segments. This regression is estimated without an intercept, reflecting the identifying assumption that there is no remaining direct boundary discontinuity once the mechanisms (distance-to-court and procedure) are included. Bootstrapped standard errors are used and specifications varying parameters (e.g., bandwidth) or incorporating certain robustness checks (e.g., covariate controls) are estimated. Please see Appendix B.11 for specification settings, which are also discussed below.

Mechanism Results

Using the procedure detailed above, we estimate many specifications of the effects of distance-to-court and procedure on default outcomes. Because of the variety of researcher-driven measurement, modeling, and estimation degrees of freedom (Gelman & Loken, 2013; Leamer, 2010; Sims, 2010; Stevenson, 2026), we present distributions of results below across many specifications. Here, the specification settings include the following:

1. Outcome Measures: two outcome measures (the no-imputation case default flag and the imputed unit defaults).
2. Procedure Measures: three procedure index specifications (discussed above).
3. Tenant Cost Measures: two tenant cost measures (distance-to-court proxy and generalized tenant costs, discussed below).
4. First-Stage Estimation: two first-stage jump estimation methods (local linear regression and local difference in means).
5. Bandwidths: various bandwidths around each boundary segment (0.5km, 1km, and 2km).
6. Number Boundary Points: varying number of boundary points (10, 20, 50).
7. Weighted vs Unweighted Estimation: weighted vs unweighted cross-segment boundary regressions.

Appendix B.11 also displays average estimates across specifications broken up by each courthouse pair and various settings.

The main results using distance-to-court as our measure of tenant costs of getting to court are shown below. In Figure 2.5, we plot histograms of beta coefficient estimates on the two mechanisms (distance-to-court and procedure index) for two outcome measures (the no-imputation case default flag and the imputed unit defaults). In each panel, the average and median estimates, along with interquartile range (IQR) are also labeled. The results here are coefficient estimates in non-standardized terms: for example, a negative estimate of the distance-to-court coefficient (β_D) implies that an increase in the distance-to-court discontinuity decreases the default status discontinuity. Under the identification assumptions, the estimates of the beta coefficients also have a structural interpretation with implications for the marginal effects of the mechanisms on default outcomes. For example, a positive estimate of β_D implies that an increase in distance-to-court increases default status holding procedure fixed.

The results across specifications are shown in Figure 2.5. The histograms represent specification coefficient values from varying different settings (e.g., bandwidth), measures, or estimation choices. Four panels are shown, with results by the two outcome measures for each of the mechanism coefficients (procedure and distance-to-court). The average, median, and interquartile range across models are displayed in each panel, with the average also plotted as a vertical black dashed line. Recall that the case default flag outcome involves no outcome imputation, whereas the imputed unit defaults outcome involves imputing zeroes to apartment units without observed default evictions.

The results in Figure 2.5 suggest specification uncertainty. In each panel, the IQRs do not identify the sign of the effect. Moreover, the direction of the effects differ based on outcome measure (procedure beta coefficient) or whether the mean or median is used (distance-to-court beta coefficient, imputed unit defaults outcome). The negative median distance-to-court beta estimates imply that higher distance-to-court decreases defaults, which does not align with the extant eviction literature but may result from the spatial distribution of cases across LA County. In particular, because the largest number of defaults occur in downtown close to courthouses, this may imply a small marginal effect of distance-to-court on defaults. If the effect were small, more data may be needed to overcome the specification uncertainty. And, indeed, standardized estimates included in Appendix Figure B.25 suggest that

model average and model median effect sizes are relatively small, especially on the distance-to-court (tenant cost) variable. The IQRs imply relatively small effect sizes in three out of four panels, although the wide IQR of procedure index betas using the case default flag outcome contains possibly large effects. Estimates from unweighted regressions are also included in Appendix Figure B.28 and Appendix Figure B.29.

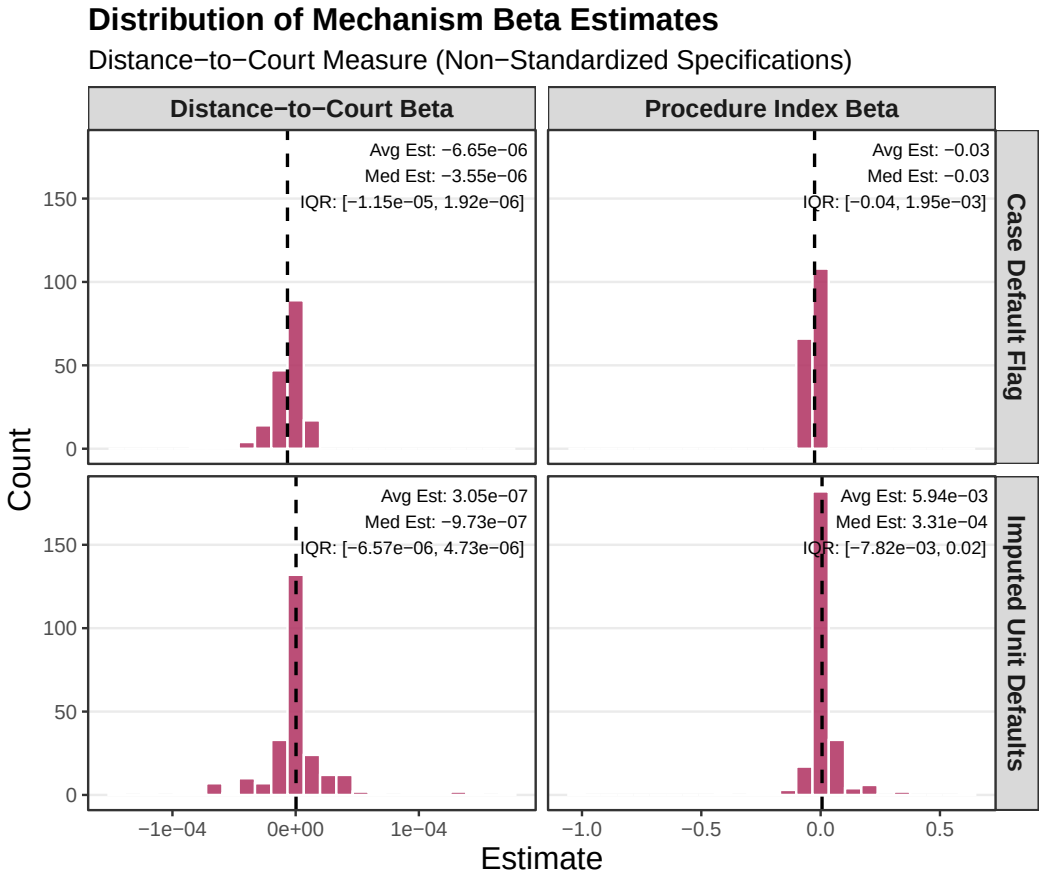

Figure 2.5: Distribution of Mechanism Beta Estimates (Distance, Non-Standardized)

Note: Each panel reports non-standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the distance-to-court variable.

Robustness: Generalized Tenant Costs

We also estimate specifications where the distance-to-court measure of tenant costs is replaced with a generalized tenant cost metric. Specifically, we use a generalized

cost measure that reflects both the direct costs of getting to court (e.g., distance-to-court, parking fees) and time/opportunity costs of appearing for a court hearing. For tenant i appearing at court c, we define i's generalized cost of attending court c as follows:

$$GC_{ic} = \underbrace{M_{ic}}_{\text{direct costs}} + \underbrace{v_i T_{ic}}_{\text{indirect costs}}$$

where M_{ic} is the direct monetary cost of traveling to court c, T_{ic} is the time it takes to attend court, and v_i is i's hourly time value (opportunity costs). Due to inflation in the 2018–2024 period, we use the FRED database CPI-U series (CPIAUCSL) to inflation-adjust all nominal amounts, with December 2024 as the base period.

First, we measure the direct costs as:

$$M_{ic} = 2 \cdot \text{distance-to-court}_{ic} \cdot \frac{\text{gas price}_{t(i)}}{25} + \text{parking fee}_c.$$

This is the monetary cost of i traveling some distance-to-court (roundtrip), assuming a fuel economy of 25mpg and using weekly LA County regular gasoline price data from the U.S. Energy Information Administration. The variable $\text{gas price}_{t(i)}$ is the weekly price of gas in the week $t(i)$ nearest i's filing date (from docket records). Parking fees vary by courthouse and are further explained in Appendix B.7.

Next, we impute indirect costs, which are the monetary opportunity costs from having to spend time attending court. We assume the time Tic it takes to attend court is 3.5 hours (either 8:30am–12:00pm or 1:30pm–5:00pm sessions)[28] plus the roundtrip commute time:

$$T_{ic} = 3.5 + 2 \cdot \text{commute-to-court}_{ic}$$

with v_i equal to i's census tract median income converted to an hourly wage (i.e., median income divided by 2080 annual working hours).

The resulting inflation-adjusted generalized cost measure is expressed in December 2024 terms. It varies over time (Appendix Figure B.24) and is positively correlated with distance-to-court (Appendix Figure B.23).

The full mechanism model is re-estimated using the generalized cost measure and the procedure index across specification settings. As above, common mechanism coefficient estimates under the mechanism model are shown in Appendix Figure B.26 (non-standardized estimates) and Appendix Figure B.27 (standardized estimates)

[28] See Appendix B.7 for further information.

specification histograms. The estimated average, median, and IQR are shown in each panel, with the average coefficients across specifications plotted as a vertical black dashed line. Estimates from unweighted regressions using the generalized cost measure are also included in Appendix Figure B.30 and Appendix Figure B.31.

The across-specification histograms again imply high uncertainty, with the IQRs failing to identify the sign of the effect. The histograms are approximately normal distributed, with average and median effects directionally similar to the distance-to-court specifications. In standardized terms, the effect sizes are also similar: the median and average effects are small for three out of four panels. The IQRs also imply relatively small effect sizes for three out of four panels, but do not pin down the direction of the coefficients. Overall, the results suggest that the generalized cost measure and distance-to-court measure both produce similar distributions of mechanism coefficient estimates, similar averages and medians across specifications, and similar IQRs containing both positive and negative mechanism coefficient estimates.

2.6 Conclusion

This paper extends research on the relationship between individual-level characteristics and eviction outcomes by exploring the mundane, yet complex institutional-level processes shaping the likelihood of tenant default in the Los Angeles Superior Court system. The institutional geography of eviction case processing in Los Angeles involves assigning cases into regional courthouses across a large metropolitan county. Since a default judgment occurs at the beginning of the eviction process, some tenants lose the cases against them before they have an opportunity to defend themselves in court. Thus, for tenants facing eviction, the likelihood of losing their case and home, and potentially experiencing homelessness, is a function of institutional factors rather than the individual demographic and cultural variables tested in much of the literature.

We find that the institutionalized and automated process of relying on the drawing of certain boundaries (case assignment) may disadvantage tenants living along the margins of jurisdictional boundaries. Indeed, in our regression discontinuity design, we show there are significant gaps in the local average eviction default probability across the boundaries for different courthouse regions. The results are, however, robust to some specification choices and sensitive to others, which we have described in the main text and detailed more extensively in the Appendices. Although some econometric uncertainty is unavoidable, these details matter because

they show where further research and additional data may produce more precise estimates, shedding light on how and why local gaps emerge. Additional court data transparency, data access, and data reporting requirements (esp. additional disaggregated or micro-data availability), in particular, may facilitate investigating the causes and consequences of eviction institutions (Nelson, 2026).

The result is that some tenants living in certain areas are likelier than other tenants to default. The spatial boundary mechanism model offered here further tests how and why local default gaps may emerge. Using data on two mechanisms (tenant costs and court procedures), we specify and estimate a mechanism model with common mechanism coefficients across boundary segments and gaps identified by variation in mechanism gaps along the boundary. The results imply a high degree of econometric uncertainty, with the model largely failing to pin down the direction of the mechanism coefficients (positive or negative). The average and median specification estimates, moreover, often imply small effect sizes for tenant costs, although the data can sustain a wide range of estimates (i.e., wide IQRs).

The multi-faceted nature of eviction institutions makes welfare-improving legal reform difficult, particularly in cities like Los Angeles. The impact of different levers in the eviction machine—assignment maps, court waiting times, legal defense resources, and procedures for dealing with court non-attendance—on tenants, landlords, and the rental market is not fully understood. Determining how to optimally pull these levers should be done *ceteris paribus* and in a data-driven manner. The models here imply that certain mechanisms may be more important than tenant costs of getting to court and courthouse procedural variation. For example, market factors (e.g., housing supply) largely outside judicial control may be larger determinants of aggregate eviction rates.

By focusing on institutional mechanisms affecting eviction default,[29] this paper suggests further study, experimentation, or reform may pay dividends in eviction prevention. Seemingly small interventions to redraw and optimize jurisdictional maps could decrease defaults and, by extension, procedural inequality in the civil justice system. The court assignment lever can also complement other policies, like creating tenant protections to keep tenants housed and expanding tenants' access to attorneys. This multi-pronged policy approach toward eviction cases could help mitigate the consequences of LA County's affordable housing and homelessness problems.

[29] See also Estes, 2025.

References

Baird, A. (2021). *Right to Counsel in Eviction Cases: A California Case Study*. https://sites.bu.edu/dome/2021/08/27/right-to-counsel-in-eviction-cases-a-california-case-study/

Calonico, S., Cattaneo, M. D., Farrell, M. H., & Titiunik, R. (2023). *rdrobust: Robust Data-Driven Statistical Inference in Regression-Discontinuity Designs* [R package version 2.2]. https://CRAN.R-project.org/package=rdrobust

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, *82*(6), 2295–2326. https://doi.org/https://doi.org/10.3982/ECTA11757

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs. *R Journal*, *7*(1), 38–51. https://rdpackages.github.io/references/Calonico-Cattaneo-Titiunik_2015_R.pdf

Collinson, R., Humphries, J. E., Mader, N., Reed, D., Tannenbaum, D., & Van Dijk, W. (2024). Eviction and Poverty in American Cities. *Quarterly Journal of Economics*, *139*(1), 57–120. https://doi.org/https://doi.org/10.1093/qje/qjad042

Desmond, M. (2017). *Evicted*. Penguin Books.

Dillon, L., Castleman, T., & Esquivel, P. (2023). Over 40,000 Eviction Notices Have Gone Out in LA This Year, Many to Upscale Apartments. *Los Angeles Times*. https://www.latimes.com/california/story/2023-09-07/over-40-000-eviction-notices-have-gone-out-in-la-this-year-many-to-upscale-apartments

Engstrom, D. F., Hagan, M., Bernal, D., Caspi, A., Knowlton, N., & Sela, A. (Apr. 3, 2025). A Blueprint for Expanding Access to Justice in Los Angeles Superior Court's Eviction Docket [Policy Report]. *Stanford Law School*. https://law.stanford.edu/publications/a-blueprint-for-expanding-access-to-justice-in-los-angeles-superior-courts-eviction-docket/

Estes, M. (2025). Eviction Reduction Policies: Do More Courts Reduce Defaults in LA County? *Working Paper*. https://matt-estes.com/assets/papers/Estes2025_EvictionReduction.pdf

Gelman, A., & Loken, E. (2013). The Garden of Forking Paths: Why Multiple Comparisons Can Be a Problem, Even When There Is No "Fishing Expedition" or "p-hacking" and the Research Hypothesis Was Posited Ahead of Time. *Department of Statistics, Columbia University*, *348*(1-17), 3. https://sites.stat.columbia.edu/gelman/research/unpublished/forking.pdf

Hoffman, D. A., & Strezhnev, A. (2023). Longer Trips to Court Cause Evictions. *Proceedings of the National Academy of Sciences*, *120*(2), e2210467120. https://doi.org/https://doi.org/10.1073/pnas.2210467120

Imbens, G. W., & Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*, *142*(2), 615–635. https://doi.org/https://doi.org/10.1016/j.jeconom.2007.05.001

Larson, E. (2006). Case Characteristics and Defendant Tenant Default in a Housing Court. *Journal of Empirical Legal Studies*, *3*, 121–144. https://doi.org/https://doi.org/10.1111/j.1740-1461.2006.00065.x

LASC. (2024). *Civil Fee Schedule (Public)*. https://lascpubstorage.blob.core.windows.net/cpw/LIBSVCExecutiveSupport-265-2024FeeSchedule010124.pdf

Leamer, E. E. (2010). Tantalus on the Road to Asymptopia. *Journal of Economic Perspectives*, *24*(2), 31–46. https://doi.org/DOI:10.1257/jep.24.2.31

Lee, D. S., & Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, *48*, 281–355. https://doi.org/10.1257/jel.48.2.281

Lens, M. C., Nelson, K., & Gromis, A. (2020). The Neighborhood Context of Eviction in Southern California. *City and Community*, *19*(4), 912–932. https://doi.org/https://doi.org/10.1111/cico.12487

Mayor, L. (2023). *Mayor Bass: Press Release on Sep. 19, 2023*. https://mayor.lacity.gov/news/apply-now-mayor-bass-councilmember-raman-and-los-angeles-housing-department-announce-rental

Nelson, K. (2023). The Political Determinants of Access to Justice. *Working Paper*.

Nelson, K. (2026). A Civil Injustice: The State of Eviction in California. 2010–2024. https://civilinjustice.net/report

Nelson, K., Gromis, A., Kuai, Y., & Lens, M. C. (2021). Spatial Concentration and Spillover: Eviction Dynamics in Neighborhoods of Los Angeles, California, 2005-2015. *Housing Policy Debate*, *31*(3-5), 670–695. https://doi.org/https://doi.org/10.1080/10511482.2020.1847163

Prescott, J. (2017). Improving Access to Justice in State Courts with Platform Technology. *Vanderbilt Law Review*, *70*(6), 1993–2050. https://scholarship.law.vanderbilt.edu/vlr/vol70/iss6/15

Riker, W. H., & Ordeshook, P. C. (1968). A Theory of the Calculus of Voting. *The American Political Science Review*, *62*(1), 25–42. https://doi.org/doi:10.2307/1953324

Sims, C. A. (2010). But Economics Is Not an Experimental Science. *Journal of Economic Perspectives*, *24*(2), 59–68. https://doi.org/DOI:10.1257/jep.24.2.59

Stevenson, M. T. (2026). Humans in the Loop: The Next Frontier in the Credibility Revolution. *Virginia Public Law and Legal Theory Research Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6308160

Wiley, H. (2023). Newsom Signs Tenant Bills to Lower Security Deposits, Prevent More Evictions. *Los Angeles Times*. https://www.latimes.com/california/story/2023-10-11/newsom-tenant-protection-california-renters-eviction-security-deposits

Chapter 3

MONOTONE REGRESSION DISCONTINUITY: PARTIAL IDENTIFICATION, ESTIMATION, AND EXTRAPOLATION IN RDDS

3.1 Introduction

Regression discontinuity is one of the most popular observational causal inference designs in applied economics. Goldsmith-Pinkham (2026) estimates that more than 10% of all applied microeconomics NBER papers employ a regression discontinuity design ("RDD"). These studies generally rely on one of two untestable assumptions to identify an average treatment at the boundary or cutoff: a continuity assumption or a local randomization assumption. Both assumptions may fail to hold in a variety of settings. For example, in the most common "sharp" RDD, treatment is determined entirely by a forcing or running variable, but if agents manipulate the forcing variable this can violate the continuity assumption (McCrary, 2008). Consequently, the RDD will not in general identify a causal effect. Placebo tests identifying discontinuities at non-treatment running variable values are also said to invalidate the design (Imbens & Lemieux, 2008).

This paper considers an alternative identifying assumption compatible with non-standard data generating processes (DGPs). The paper shows that in a setting where the average potential outcome functions are monotone in the forcing variable and outcomes are bounded, the cutoff or boundary average treatment effect is partially identified. Sharp bounds on the cutoff average treatment effect are derived. The bounds are compatible with jumps or discontinuities in the potential outcome conditional expectation functions (CEFs) at the cutoff that determines treatment. Discontinuous but monotone data generating processes (e.g., discrete monotone processes) generally may have partially identified cutoff treatment effects provided potential outcomes or averages of potential outcomes are also bounded locally.

Next, the paper proposes two novel estimators for the monotone setting. Using the monotonicity of the CEFs, the paper defines *local extremum* and *local reverse kernel* estimators. The estimators are designed for robustness to potentially confounding features of data generating processes and researcher free parameters (bandwidth selection). For example, the estimators can be used in RDD studies

with: discontinuous DGPs, trend reversals, discrete running variables, measurement error, or missing data near the cutoff (e.g., "donut RDD"). Both estimators are also bandwidth-robust: they cover the true treatment effect bounds at a variety of selected bandwidths. Instead, bandwidth selection, which is the most important free-parameter in RDD estimation, becomes an issue of efficiency: certain bandwidths may achieve tighter treatment effect bounds or narrower partially-identified regions.

The paper illustrates the performance of the two estimators in comparison to the local linear estimator (the most popular RDD estimation strategy) via simulations. The local extremum and local reverse kernel estimators have good coverage properties in finite samples across a range of data generating processes. The local linear estimator often performs less well at estimating bounds, but may still cover the true treatment effect when there is a discontinuity at the cutoff in certain cases.

The identification power of monotonicity in other designs targeting other estimands is then demonstrated. First, results are derived to show that monotonicity and outcome boundedness partially identify extrapolated bounds away from the boundary or cutoff: assuming there is some window around the cutoff determining treatment where both CEFs are monotone, the average treatment effect (ATE) can be bounded in that entire window. Such bounds may be of more policy relevance than point identification of *merely local* cutoff average treatment effects. Second, the partial identification in fuzzy RDD settings is discussed, under the additional "monotone treatment response" assumption (Manski, 2007). The fuzziness of the assignment can tighten the bounds away from the boundary, which is demonstrated via an example.

Finally, the method is applied to the RDD study from Sejas-Portillo, Moro, and Stowasser (2025). Using discrete energy efficiency ratings, UK properties receive a simplified, low-information energy band rating (A to G). Prices are generally monotone increasing in the running variable, and Sejas-Portillo, Moro, and Stowasser (2025) document price discontinuities at the energy band cutoffs. The application here reproduces those estimates and applies the estimators to obtain bounds on cutoff effects and placebo effects at other values of the running variable. The results here show that the treatment effect estimates at the cutoffs hold, even under the different counterfactual assumption, alternative estimators, and different bandwidths. Placebo tests further suggest that treatment effects at the energy band cutoffs are large under various estimators.

Prior Literature

I contribute to different strands of the RDD literature. First, I contribute to the literature concerned with continuity failures near the boundary. For example, manipulation of the forcing variable can invalidate the continuity assumption (McCrary, 2008). Several ad hoc strategies—e.g., the so-called "donut hole" RDD (Noack & Rothe, 2023)—have been suggested, but monotone RDD may provide bounds that are informative in such settings. Second, I contribute to a growing literature concerned with continuity-type assumptions. Korting et al. (2023) show that evaluating CEFs for discontinuities can be difficult, even for expert evaluators. By contrast, the monotone RDD method advanced here relies on a monotonicity assumption, whose plausibility may be simpler for researchers to recognize and evaluate. In particular, the plausibility of the monotonicity assumption may be tested on the observable data. Additional work on estimation and inference in monotone regression discontinuity settings has also been performed for various estimators, including isotonic regression (Babii & Kumar, 2023). Third, this work similarly contributes to studying how certain estimators and specifications perform in various settings (Gelman & Imbens, 2018), with emphasize on bandwidth selection, discrete data generating processes, and other applied issues (e.g., missing data, running variable measurement error). Finally, I contribute to the vast literature on partial identification. For example, I extend Manski (2007) type monotone bounds to the RDD setting.

The rest of the paper is organized as follows. Section 3.2 reviews continuity-based point identification, states the main partial identification result, and discusses estimation strategies. Section 3.3 surveys simulation study results and evaluates small, finite sample coverage properties for the three types of estimators. Section 3.4 shows how to bound the ATE under global monotonicity and gives identification results in the fuzzy regression discontinuity design. Section 3.5 is the application to Sejas-Portillo, Moro, and Stowasser (2025). Further details and results can be found in the Appendices.

3.2 RDD Identification: Local Treatment Effect

This section explains a new regression discontinuity methodology, which exploits variation stemming from the discontinuous treatment assignment to partially identify the local (or cutoff) treatment effect in a sharp regression discontinuity design. We first consider the typical sharp RDD identification strategy and contrast it to our partial identification method. In our method, we give upper and lower bounds on the local treatment effect for a data generating process (DGP) that is discontinuous

but monotone at the treatment cutoff. Finally, we discuss estimation and inference strategies, including model specification and bandwidth selection.

Standard Point Identification: Identifying τ_C Under Continuity in the Forcing Variable

To begin, we assume a potential outcomes framework. Let Y_{0i} and Y_{1i} be the counterfactual outcomes for an individual i without and with treatment, respectively. The binary treatment W_i is determined from the "forcing variable" X_i and a known cutoff or threshold C by $W_i = \mathbf{1}\{X_i \geq C\}$. The observed outcome is given by the switching equation $Y_i = (1 - W_i)Y_{0i} + W_i Y_{1i}$. Define further the counterfactual conditional expectation functions (CEFs):

$$\mu_g(x) = \mathbb{E}[Y_g | X = x]$$

for $g = 0, 1$. In the usual RDD setting, the potential outcome CEFs are assumed to be continuous in x. Moreover, because W_i is a deterministic function of X_i note that ignorability (conditional mean independence) holds:

$$\mathbb{E}[Y_g | X, W] = \mathbb{E}[Y_g | X]$$

for $g = 0, 1$.

The usual estimand of interest in a sharp RDD is the local average treatment effect (or cutoff treatment effect), defined as:

$$\tau_C = \mathbb{E}[Y_1 - Y_0 | X = C].$$

Following Wooldridge (2010), we first relate the potential outcome CEFs to the observed CEF as follows. By continuity of $\mu_g(\cdot)$, ignorability, and the treatment assignment rule, we find that:

$$\underbrace{\mathbb{E}[Y_0 | X = C]}_{\equiv \mu_0(C)} = \lim_{x \to C^-} \mathbb{E}[Y_0 | X = x] = \lim_{x \to C^-} \mathbb{E}[Y_0 | W = 0, X = x] = \lim_{x \to C^-} \mathbb{E}[Y | X = x].$$

By a similar argument, we identify $\mu_1(C)$ with the limit from the right:

$$\mu_1(C) = \lim_{x \to C^+} \mathbb{E}[Y | X = x].$$

It follows that the local average treatment effect τ_C is identified by a difference in one-sided limits of the observed CEF:

$$\tau_C = \mu_1(C) - \mu_0(C) = \lim_{x \to C^+} \mathbb{E}[Y | X = x] - \lim_{x \to C^-} \mathbb{E}[Y | X = x].$$

In general, note that $\mathbb{E}[Y|X = x, W = 0]$ is estimable from the observable data for all values $x < C$ (and similarly for $W = 1$ with $x \geq C$). Therefore, we can estimate τ_C by estimating their limits. The most common strategy is local linear regression:[1] for observations within a small bandwidth $h > 0$, regress Y_i on $1, (X_i - C)$ for two subsets of data, one with $X_i \in (C - h, C)$ and the other with $X_i \in [C, C + h)$. The difference in the estimated intercepts gives $\hat{\tau}_C$.

An example of a sharp regression discontinuity is given below in Figure 3.1. The dashed lines are unobservable in the sharp design, whereas the solid lines are observable.[2] In the Figure 3.1 example, there is a constant treatment effect of 7 across the support of X, which also means that $\tau_C = ATE$.

[1] Local polynomial regression can also be used to estimate the limit of the conditional means. See Wooldridge (2010) or Imbens and Lemieux (2008) for further details, including information on bandwidth selection.

In general, to estimate the local treatment effect we do not use global polynomial fits to estimate the limits of the CEFs following results in Gelman and Imbens (2018). In Gelman and Imbens (2018), the authors find *inter alia* that (global) higher order polynomial estimation attaches noisy weights to observations and may not achieve nominal coverage rates.

[2] That is, the solid lines are, in principle, observable by virtue of the sharp treatment assignment. In practice, of course, the solid lines can only be (consistently) estimated from the data.

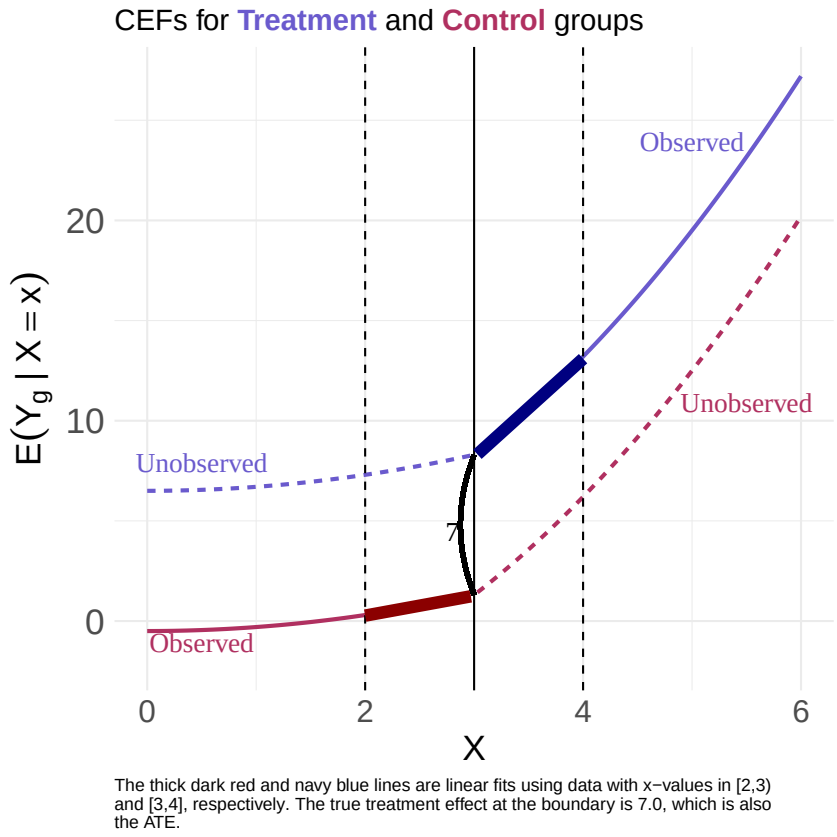

Figure 3.1: Sharp RDD: Continuous CEFs with Constant Treatment Effect

Partial Identification: Bounding τ_C Under Monotonicity in the Forcing Variable and Bounded Outcomes

The example in Figure 3.1 assumes continuity and a constant treatment effect across the entire support of forcing variable X, which may not hold in practice. For example, manipulation of the forcing variable X can invalidate the crucial CEF continuity assumption (McCrary, 2008). Here, I detail simple assumptions to obtain bounds on the cutoff treatment effect while allowing for discontinuities in the CEFs. Specifically, I assume monotonicity in the forcing variable and boundedness in the outcome-of-interest to partially identify the local treatment effect or the cutoff effect τ_C.

Again assume the standard (sharp) RDD setting, with treatment $W_i = \mathbf{1}\{X_i \geq C\} \in \{0, 1\}$ as above, for forcing variable X_i and cutoff C. Recall the CEFs were written

as:

$$\mu_1(x) = \mathbb{E}[Y_1|X = x]$$
$$\mu_0(x) = \mathbb{E}[Y_0|X = x]$$

Consider the following monotonicity[3] and boundedness[4] assumption on the CEFs:

Assumption 1 (Monotonicity)**.** *The conditional expectation functions μ_1, μ_0 are weakly monotone increasing in x for all $x \in supp(X)$.*

Assumption 2 (Boundedness)**.** *The conditional expectation functions μ_1, μ_0 are bounded above by M_1 and bounded below by M_0 for all $x \in supp(X)$.*

Some initial remarks on these assumptions and what is and is not assumed. First, the continuity of $\mu_1(x), \mu_0(x)$ in x is not assumed. Second, global monotonicity (Assumption 1) is strictly stronger than required for bounding τ_C: it is assume for exposition, but it is sufficient for bounding τ_C that there be a small bandwidth $h > 0$ s.t. in the interval $(C - h, C + h)$ around the cutoff C the CEFs μ_1, μ_0 are monotone. Third, it is assumed for exposition that $\mu_1(x), \mu_0(x)$ are weakly monotone increasing, but it is not required that μ_1, μ_0 are monotone *increasing*, or even that they are *both increasing* or *both decreasing*.[5] Fourth, in some settings, the upper bound on $\mu_0(C)$ may be given by $\mu_1(C)$, which corresponds to Manski's "monotone treatment response" (Manski, 2007).[6] Fifth, observe that ignorability (conditional mean independence) still holds because $W_i = \mathbf{1}\{X_i \geq C\}$ is a deterministic function of X_i, so that for $g = 0, 1$ we have $\mathbb{E}[Y_g|X, W] = \mathbb{E}[Y_g|X]$.

Under the assumptions, we can partially identify the usual estimand in an RD setting, the average treatment effect at the cutoff:

$$\tau_C = \mathbb{E}[Y_1 - Y_0|X = C] = \mu_1(C) - \mu_0(C).$$

[3]Technically, we only require monotonicity on the unobservable portion of the CEF. However, it would be odd to make an assumption on only the unobservable half of a CEF in the RDD design if, on the observable portion of the CEF, monotonicity does not hold near the cutoff. See discussion below about motivating this assumption in context and using the observable portion of the CEFs to test for monotonicity.

[4]Note that boundedness is only necessary to establish a lower bound in what follows. As with monotonicity, we also only need some form of local boundedness for the unobservable half of the CEFs in the sharp design. We assume boundedness over the entire support of the forcing variable X in the exposition though.

[5]That is, μ_1 can be monotone increasing, and μ_0 can be monotone decreasing.

[6]Otherwise, Assumption 2 may be particularly relevant when the outcomes of interest are rates that are bounded above by 1 and below by 0, or other outcomes that are guaranteed to be bounded above and below.

To discuss identification, it will be useful to write the CEFs as separate functions to the left and right of the cutoff C:

$$\mu_1(x) = \mathbf{1}\{x < C\}\mu_1^L(x) + \mathbf{1}\{x \geq C\}\mu_1^R(x)$$
$$\mu_0(x) = \mathbf{1}\{x < C\}\mu_0^L(x) + \mathbf{1}\{x \geq C\}\mu_0^R(x)$$

Discontinuity at the cutoff C in this notation means, for example, that $\mu_0^L(C) \neq \mu_0^R(C)$. The identification problem for τ_C is that $\mu_0(C) = \mu_0^R(C)$ is unobservable by the assignment rule ($W_i = \mathbf{1}\{X_i \geq C\}$).

Although we cannot point identify the average treatment effect at the cutoff, we can partially identify it. The (sharp) bounds are given by Theorem 1:

Theorem 1 (Local Treatment Effect (τ_C) Bounds)**.** *Let μ_1, μ_0 be the potential outcome CEFs defined as above, satisfying Assumption 1 and Assumption 2. Then bounds for the local treatment effect τ_C are $\tau_C \in [\underline{\tau_C}, \overline{\tau_C}]$, where:*

$$\overline{\tau_C} = \mu_1^R(C) - \lim_{x \to C^-} \mu_0^L(x)$$
$$\underline{\tau_C} = \mu_1^R(C) - M_1$$

Proof. By the sharp RDD treatment assignment rule, $\mu_1(C) = \mu_1^R(C)$ is observable but the identification problem is that $\mu_0(C) = \mu_0^R(C)$ is unobservable.[7] However, by monotonic increasing assumption, we have that $\mu_0(x) \leq \mu_0(C)$ for all $x < C$, which implies that

$$\mu_0(C) \geq \sup_{x<C} \mu_0(x) = \sup_{x<C} \mu_0^L(x) = \lim_{x \to C^-} \mu_0^L(x)$$

where the one-sided (left) limit is equal to the supremum, guaranteed to exist, and finite (by boundedness assumption).[8] The upper bound does not require Assumption

[7] Although $\mu_1(C) = \mu_1^R(C)$ can also bounded from above by $\lim_{x \to C^+} \mu_1^R(x)$, it is observable under the assignment rule and so does not pose an identification problem. That is, the identification and estimation tasks are kept conceptually distinct. Below, we generally use the estimator of $\mu_1^R(C)$ from local linear regression, even though this is in principle observable under the treatment assignment rule and can be consistently estimated by the sample mean for units with $X_i = C$. This is consistent with the standard continuity-based RDD practice. For identification, we technically only need the inequalities on $\mathbb{E}[Y_0|X = C]$.

[8] If $f(x)$ is a monotone increasing function on interval $(0, C)$ that is bounded above as assumed, then:

$$\lim_{x \to C^-} f(x) = \sup_{x \in (0,C)} f(x)$$

More generally, bounded monotone functions always have well-defined and finite one-sided limits on their domain.

2 but the lower bound does: to identify the lower bound $\underline{\tau_C}$, use $\mu_0(x)$ bounded above on $\text{supp}(X)$ by M_1 to obtain

$$\mu_0(C) \leq M_1$$

These inequalities together imply that the local treatment effect $\tau_C = \mu_1(C) - \mu_0(C)$ is bounded by:

$$\underbrace{\mu_1^R(C) - M_1}_{\underline{\tau_C}} \leq \tau_C \leq \underbrace{\mu_1^R(C) - \lim_{x \to C^-} \mu_0^L(x)}_{\overline{\tau_C}}$$

The bounds are sharp because, under the assumptions, $\mu_0(C) = \mu_0^R(C)$ can take any value in $[\lim_{x \to C^-} \mu_0(x), M_1]$ and remain compatible with the observed data. □

The plots in Figure 3.2 give an example DGP compatible with our assumptions. In the left plot, the true CEFs for the treatment and control groups are plotted. The solid lines are observable (by which we mean we can estimate them from data), whereas the dashed lines are unobservable because of the sharp treatment assignment. The typical RDD assumption that both CEFs $(\mu_1(x), \mu_0(x))$ are continuous in x **does not hold.** The functions for the potential outcome CEFs are given by:

$$\mu_1(x) = \mathbf{1}\{x < 3\}(0.2x^2 + 2) + \mathbf{1}\{x \geq 3\}(0.7x^2 + 2)$$
$$\mu_0(x) = \mathbf{1}\{x < 3\}(0.2x^2 - 2) + \mathbf{1}\{x \geq 3\}(0.7x^2 - 5)$$

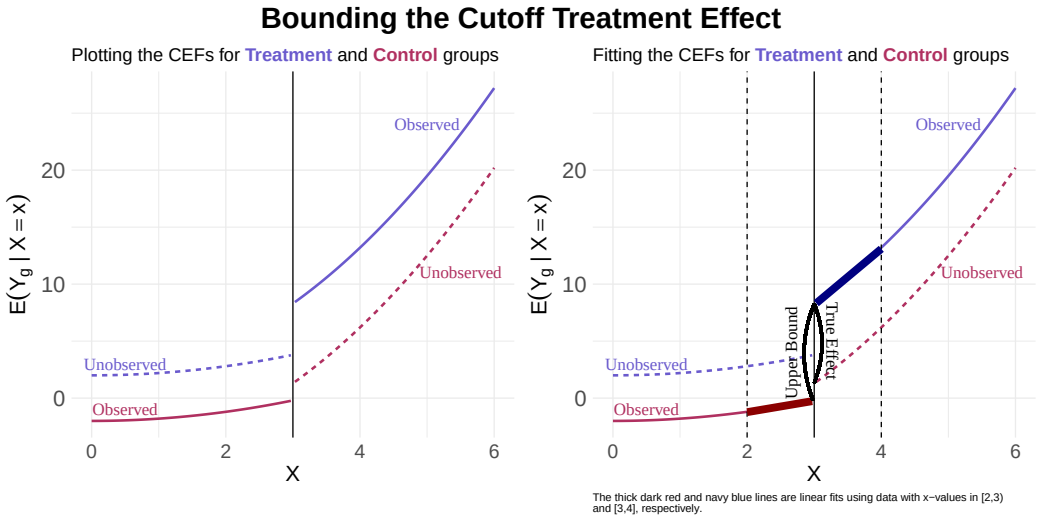

Figure 3.2: Bounding the Treatment Effect at the Cutoff $C = 3$

In the right panel of Figure 3.2, we show how the usual estimation procedure of obtaining a linear fit in a small bandwidth around the cutoff $C = 3$ estimates the upper bound $\overline{\tau_C}$ rather than τ_C. For this particular example, the true effect at the cutoff $C = 3$ is given by the difference between the solid blue line (observable μ_1) and the dashed red line (unobservable μ_0):

$$\tau_C = \mu_1(3) - \mu_0(3) = (0.7(3)^2 + 2) - (0.7(3)^2 - 5) = \boxed{7}$$

But the upper bound $\overline{\tau}_C$ at the cutoff is given by taking the difference between $\mu_1(3)$ and the lower bound for $\mu_0(3)$. In Figure 3.2 above, this results in the upper bound:

$$\overline{\tau_C} = (0.7(3)^2 + 2) - (0.2(3)^2 - 2) = 8.3 - (-0.2) = \boxed{8.5}$$

found from the known CEFs above.[9]

This example shows that for the class of DGPs considered here (monotone potential outcome CEFs with discontinuities at the cutoff) the standard RDD estimation is biased for τ_C. In this example, the standard estimation procedure for a sharp RDD will produce an estimated $\hat{\tau}_C$ that is 20% larger than the true effect although, in DGPs like the example, local linear regression still consistently estimates *the bounds* $(\overline{\tau_C}, \underline{\tau_C})$.[10]

Estimation & Inference

Monotonicity on the entire support of the forcing variable is a strong shape restriction. However, as discussed above, we only require a small bandwidth $h > 0$ s.t. μ_1, μ_0 are monotone in $(C - h, C + h)$. Although we leave the optimal estimation and inference procedures for future work, I study here two novel estimation and inference procedures: local extremum and local reverse kernel estimators. In Section 3.3, I investigate the finite-sample performance of these estimators in comparison to the widely used local linear estimator via simulations.

[9] In practice, we obtain estimates on $\hat{\mu_1}(3)$ and $\hat{\mu_0}(3)$ from the local linear regression in a small window around the cutoff, as in the rightmost plot. For the lower bound, $M_1 = 27.2$ bounds both μ_0, μ_1 from above on $[0, 6]$, so that a lower bound on τ_C is:

$$\underline{\tau_C} = \mu_1(C) - M_1 = (0.7(3)^2 + 2) - (0.7(6)^2 + 2) = 8.3 - 27.2 = \boxed{-18.9}$$

This bound is sharp if the researcher does not assume a stronger form of boundedness. Note that if we instead assume $\mu_0(C)$ is bounded above by $\mu_1(C)$ (e.g., "monotone treatment response"), then $M_1 = \mu_1(C)$ and the lower bound is zero. See Manski (2007).

[10] For example, local linear regression in a window around the cutoff will consistently estimate $\overline{\tau_C}$ when the potential outcome CEFs are monotone increasing with only one discontinuity at the cutoff and are otherwise smooth. For more general monotone DGPs with many discontinuities, a modified estimation procedure may be necessary and bandwidth selection will be much more important.

Recall that our estimands are given by:

$$\overline{\tau_C} = \mu_1^R(C) - \lim_{x \to C^-} \mu_0^L(x)$$
$$\underline{\tau_C} = \mu_1^R(C) - M_1$$

Assuming that M_1 is known to the researcher, the estimation task is twofold: we must estimate $\mu_1^R(C)$ and $\lim_{x \to C^-} \mu_0^L(x)$. If there are sufficiently many treated units with $X_i = C$, then $\mu_1^R(C)$ can be estimated with the sample average for units with $X_i = C$. Assuming this is not the case, we consider three ways to estimate these quantities: (1) local linear estimators, (2) local extremum estimators, and (3) reverse kernel estimators.

Local Linear Estimation

The local linear estimator is well-known and oft-applied in regression discontinuity designs. The estimation procedure is straightforward: obtain linear fits within a bandwidth $h > 0$ on either side of the cutoff C:

$$\min_{\alpha_L, \beta_L} \sum_{i: C-h < X_i < C} (Y_i - \alpha_L - \beta_L(X_i - C))^2 \quad \text{and} \quad \min_{\alpha_R, \beta_R} \sum_{i: C \leq X_i < C+h} (Y_i - \alpha_R - \beta_R(X_i - C))^2$$

Then estimate $\mu_1^R(C)$ and $\lim_{x \to C^-} \mu_0^L(x)$ as follows:

$$\widehat{\mu_1^R(C)} = \hat{\alpha_R}$$
$$\widehat{\lim_{x \to C^-} \mu_0^L(C)} = \hat{\alpha_L}$$

which implies that estimated upper and lower bounds are given by:

$$\widehat{\overline{\tau_C}} = \hat{\alpha_R} - \hat{\alpha_L}$$
$$\widehat{\underline{\tau_C}} = \hat{\alpha_R} - M_1$$

As noted in Imbens and Lemieux (2008), standard least squares methods of inference may be used. When the bandwidth h goes to zero at the appropriate rate as sample sizes increase, the estimates are asymptotically normal so that EHW standard errors may be used. As shown in Kolesar and Rothe (2018), the EHW standard errors have good coverage properties even in small samples where the running variable is discrete; by contrast, CIs based on clustering by the running variable ("CRV") may fail to achieve 95 percent coverage.

Local Extremum Estimation

Although local linear regression is the most commonly used RDD estimation strategy, for discontinuous but monotone DGPs (which are allowed by our identification strategy) local linear regression may fail to provide good estimates of the upper and lower bound. Although it is possible higher order polynomials may work, in practice the introduction of higher order polynomials may introduce specification bias.

To allow for discontinuous processes and alleviate specification bias, we consider a local extremum estimator of the estimands of interest. Specifically, our estimator is motivated by rewriting the left limit of $\mu_0^L(x)$ as follows:

$$\lim_{x \to C^-} \mu_0^L(x) = \lim_{h \to 0^+} \left(\inf_{0 < C - x < h} \mu_0^L(x) \right)$$

which suggests, e.g., taking a local minimum (when μ_0 is monotone increasing) over a shrinking window $(C - h, C)$ where h is a bandwidth. In particular, for μ_0 monotone increasing and bandwidth $h > 0$ we will have that:

$$\min_{x \in (C-h, C)} \mu_0(x) \leq \lim_{x \to C^-} \mu_0^L(x)$$

Accordingly, for μ_0, μ_1 both monotone increasing we can estimate, for bandwidth $h > 0$, the two quantities of interest as:

$$\widehat{\mu_1^R(C)} = \max_{x \in [C, C+h)} \frac{1}{n_{x,1}} \sum_{i: X_i = x, W_i = 1} Y_i$$

$$\widehat{\lim_{x \to C^-} \mu_0^L(C)} = \min_{x \in (C-h, C)} \frac{1}{n_{x,0}} \sum_{i: X_i = x, W_i = 0} Y_i$$

where $n_{x,w}$ is the number of units with $X_i = x$ and $W_i = w$. For inference, we use a simple bootstrap procedure to obtain standard errors and construct confidence intervals, which performs well in our Section 3.3 simulation study.

The local extremum estimator is illustrated below in Figure 3.3. In the left panel, the solid red line is the true control outcome CEF for DGP1 (discussed below) and each point is the population mean $\mathbb{E}[Y|X = x]$: the vertical gray and brown lines show the estimated minimums with bandwidths of $h = 1.5$ and $h = 1.0$, resp. In the right panel, we show the estimated minimum $\widehat{\lim_{x \to C^-} \mu_0^L}(C) = \min_{x \in (C-h, C)} \mathbb{E}[Y_0|X = x]$ varying h from 3 to 0: as the bandwidth h goes to 0, the local extremum estimator $\widehat{\lim_{x \to C^-} \mu_0^L}(C)$ approaches the true value $\lim_{x \to C^-} \mu_0^L(C)$.

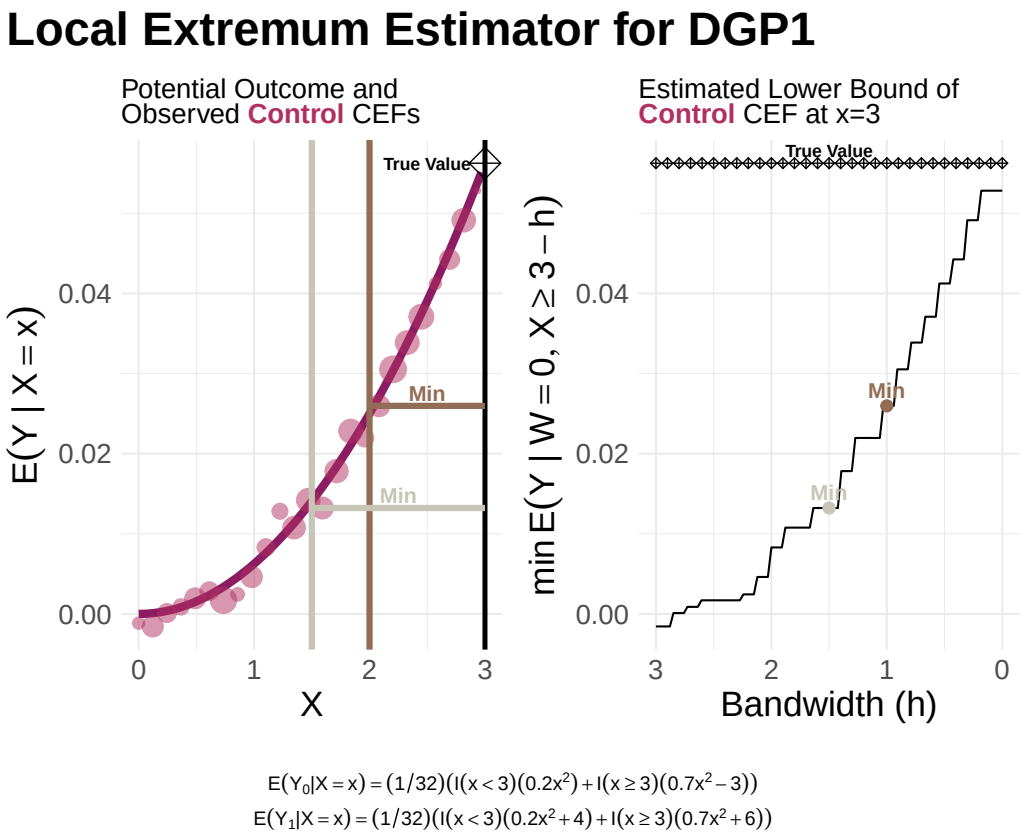

Figure 3.3: Local Extremum Estimation (DGP1 Example)

Note: In the left plot, the dark red, monotonic curve is the true potential outcome CEF, with the true limit value at x=3 represented by the diamond cross point. The points are the observed CEF with volume proportional to the share of observations with X=x. The vertical lines mark bandwidth windows (bandwidths h=1.5 and h=1.0) with the horizontal lines denoting the estimated minimum of the observed control CEF within each window. The right plot is the estimated minimum of the control CEF by bandwidth. The tan and gray points correspond to the minimum estimates from the left plot.

One advantage of the local extremum estimator is evident from Figure 3.3: the estimator $(\lim_{x \to C^-} \widehat{\mu_0^L}(C))$ is a lower bound for the true left-sided limit $(\lim_{x \to C^-} \mu_0^L(C))$ for every bandwidth value $h \in [0, 3]$. The local extremum estimator is, in this sense, robust to bandwidth selection although it may be overly conservative for certain bandwidths. For instance, the larger bandwidths (e.g., $h = 2.5$) in the right panel of Figure 3.3 give estimated lower bounds near zero, when the true left-sided limit value at $x = 3$ is near 0.056. Moreover, bandwidth selection is often the most important researcher decision in an RDD study, which recent work from Korting et al. (2023) shows is consequential for evaluating the validity of the regression discontinuity design.

Another thing to note is how this estimator uses the monotonicity assumption. For a globally monotone increasing function as in Figure 3.3, the estimator increases in a stepwise but discontinuous fashion as bandwidth $h \to 0$. Again, we do not require global monotonicity of the potential outcome CEFs, but the estimator will perform similarly provided there is a window around the cutoff C where the function is monotone. In a setting with a trend reversal, e.g., the step-like behavior of the local extremum estimator may only hold in a smaller window around the cutoff C. For that reason, it makes sense to show results for a variety of bandwidths, even if that is not strictly necessary for identification.

Local Reverse Kernel Estimation

The next estimator we consider is a local reverse kernel estimator. I motivate this estimator similarly: it is a conservative estimator that uses the monotonicity assumption and is robust to bandwidth selection. As is well-known, kernel estimators often behave poorly at point identification near the boundary.[11] Nevertheless, kernel estimators can *bound* or partially identify the boundary.

To begin, I therefore consider reverse kernel weighting functions. For example, unlike the traditional parabolic kernel which equals $K(u) = \frac{3}{4}(1 - u^2)$ for $|u| \leq 1$ and is zero elsewhere, we consider a *reverse* parabolic kernel defined by:

$$K(u) = \begin{cases} \frac{3}{2}u^2 & \text{if } |u| \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

This kernel places relatively more weight on further away observations. Using this kernel function $K(u)$, our estimator for a given bandwidth $h > 0$ is then given by:

$$\widehat{\lim_{x \to C^-} \mu_0^L(C)} = \frac{\sum_{i:W_i=0} K\left(\frac{C-X_i}{h}\right) Y_i}{\sum_{i=1}^n K\left(\frac{C-X_i}{h}\right)}$$

The estimator is analogous for estimating $\mu_1^R(C)$.

An example of how the local reverse kernel estimator works is given below in Figure 3.4. In the top-left of the left panel, the inset plot shows the parabolic and reverse parabolic kernel weights as a function of u ($K(u)$). The left panel, shows a noisy sine curve data generating process: the solid blue (resp. dashed green) curve is the reverse

[11] See, e.g., Imbens and Lemieux (2008) for a discussion.

parabolic (resp. parabolic) fitted kernel regression line with bandwidth $h = 0.5$. In the right panel is the estimated boundary value by bandwidth h: for example, the solid green and blue boxes are the estimates for the right boundary value, whereas the horizontal line is the true right boundary value. As the bandwidth $h \to 0$, both estimates generally serve as a good upper bound on the true right boundary value. However, the reverse parabolic kernel estimator remains an upper bound for almost all bandwidth choices: by contrast, the parabolic kernel estimator is not an upper bound for bandwidth $h < 0.25$.

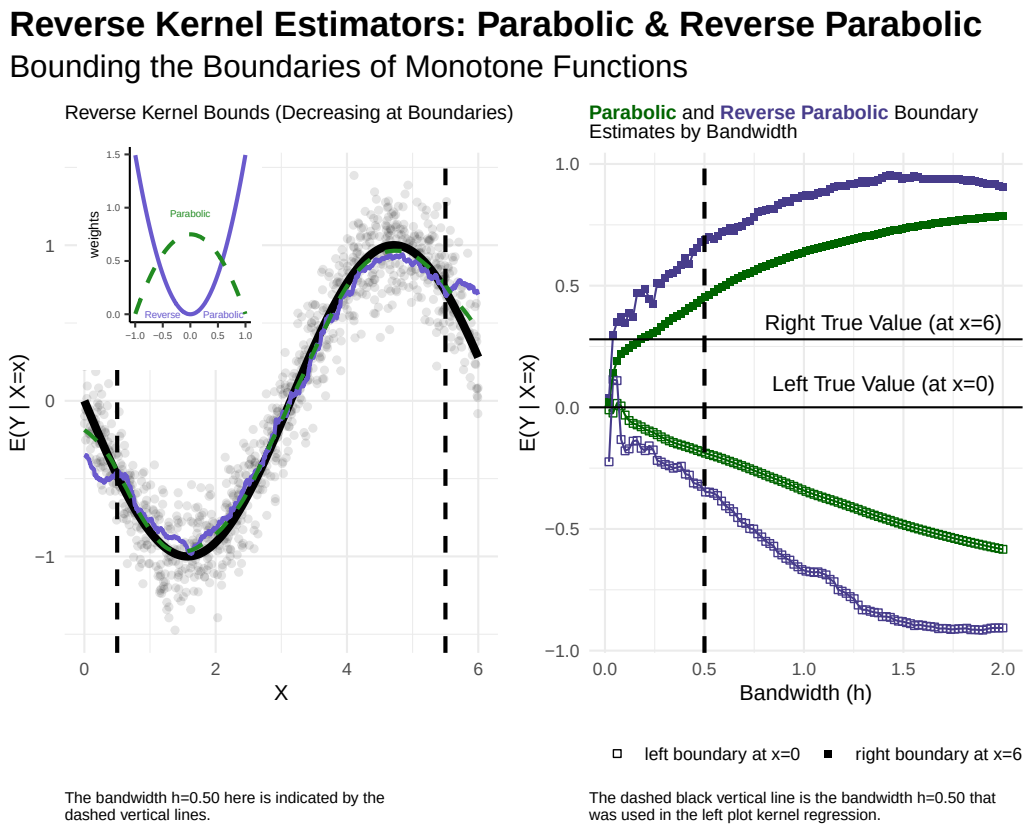

Figure 3.4: Kernel Estimation: Parabolic & Reverse Parabolic Kernels

Note: The left plot is the unknown function (solid black) with reverse parabolic (solid blue) and parabolic (dashed green) kernel estimates. The right plot are the kernel estimates by bandwidth for the same decreasing-at-boundaries DGP. The inset plot in the top-left of the left plot give the weights for the parabolic (green) and reverse parabolic (blue) kernels.

When the CEF is increasing near the left boundary, then the reverse kernel estimator serves as an upper bound on the true left boundary value. Conversely, when the CEF is decreasing near the left boundary (as it is in Figure 3.4), then the estimators serve as a lower bound on the true value at the left boundary. Similarly, when the CEF is

increasing (resp. decreasing) near the right boundary, the reverse kernel estimator is a lower (resp. upper) bound on the true right boundary value.

It should also be apparent that the local reverse kernel estimator will likely only work well for either the upper or lower bound. And whether it is a good upper or lower bound will depend on the monotonicity near the boundary. Indeed, we verify this in the Section 3.3 simulations below. The reason is, for example, that if the control CEF is decreasing near its right boundary, then the kernel estimate at the right boundary will be larger than the true left-sided limit $\lim_{x \to C^-} \mu_0^L(x)$. Accordingly, since we subtract that estimate, which is larger than the true left-sided limit value, to compute the upper bound, the estimated upper bound will be smaller the true upper bound. However, this will not affect the lower bound: subtracting a larger quantity will simply result in an estimated lower bound that is less than the true lower bound and, hence, it may still serve as a bound. Both the local extremum and local reverse kernel estimators are, in this sense conservative and may sacrifice some efficiency (in terms, e.g., of partially-identified region width) for higher coverage rates.[12]

Summary: Three Estimators

In the rest of the discussion, we will use the following notation for our three different estimators (local linear, local extremum, and local reverse kernel) of $\lim_{x \to C^-} \widehat{\mu_0^L}(C)$ and $\widehat{\mu_1^R(C)}$:

$$\widehat{\mu_{0,LL}^L(C)}(h) = \hat{\alpha_L} \qquad \widehat{\mu_{1,LL}^R(C)}(h) = \hat{\alpha_L}$$

$$\widehat{\mu_{0,LE}^L(C)}(h) = \min_{x \in (C-h, C)} \frac{1}{n_{x,0}} \sum_{i: X_i = x, W_i = 0} Y_i \qquad \widehat{\mu_{1,LE}^R(C)}(h) = \max_{x \in [C, C+h)} \frac{1}{n_{x,1}} \sum_{i: X_i = x, W_i = 1} Y_i$$

$$\widehat{\mu_{0,LRK}^L(C)}(h) = \frac{\sum_{i: W_i = 0} K\left(\frac{C - X_i}{h}\right) Y_i}{\sum_{i=1}^n K\left(\frac{C - X_i}{h}\right)} \qquad \widehat{\mu_{1,LRK}^R(C)}(h) = \frac{\sum_{i: W_i = 1} K\left(\frac{C - X_i}{h}\right) Y_i}{\sum_{i=1}^n K\left(\frac{C - X_i}{h}\right)}$$

where $K(\cdot)$ is a reverse kernel function, $n_{x,w}$ is the number of units with $X_i = x$ and $W_i = w$, $\hat{\alpha_L}$ comes from the minimization problem $\min_{\alpha_L, \beta_L} \sum_{i: C-h < X_i < C} (Y_i - \alpha_L - \beta_L (X_i - C))^2$, and $\hat{\alpha_R}$ comes from the minimization problem $\min_{\alpha_R, \beta_R} \sum_{i: C \leq X_i < C+h} (Y_i - \alpha_R - \beta_R (X_i - C))^2$. Note that each estimator is a function of the bandwidth h.

[12] Confidence intervals with coverage probability "at least the nominal level regardless of the size of the sample" are referred to as conservative. (Finkelstein, Tucker, & Veeh, 2000). In some cases, "traditional and conservative intervals provide nearly equal coverage probability" while remaining "competitive" in terms of the lengths of the conservative and traditional CIs. (Finkelstein, Tucker, & Veeh, 2000).

3.3 Simulation Results

In this section, we evaluate the three estimators (local linear, local extremum, and local reverse kernel) with simulations. We assess the coverage rates for each estimator as well as show how the two more conservative estimators (local extremum and local reverse kernel) give conservative bounds across more bandwidths and for relatively small sample sizes.

The simulation uses data generated from four data generating processes ("DGPs"). In all example DGPs, the true potential outcome CEFs have $M_1 = 1$ and $M_0 = 0$ upper and lower bounds, resp., near the cutoff of $C = 3$. In each DGP, treatment is determined by a sharp cutoff rule $W_i = \mathbf{1}\{X_i \geq 3\}$ and the forcing variable is always $X_i \in [0, 6]$. The potential outcome CEFs and the population data for each DGP are shown below in Figure 3.5. The treated outcome for control units are not observable for units with $X_i < 3$, whereas the control outcome for treated units is not observable for units with $X_i \geq 3$. There are 10,000 or 100,000 observations in each DGP, and the number of unique values in $\text{supp}(X)$ ranges from 30 to 1000. Please see Appendix C.3 for further plots and information on each DGP and associated population data.

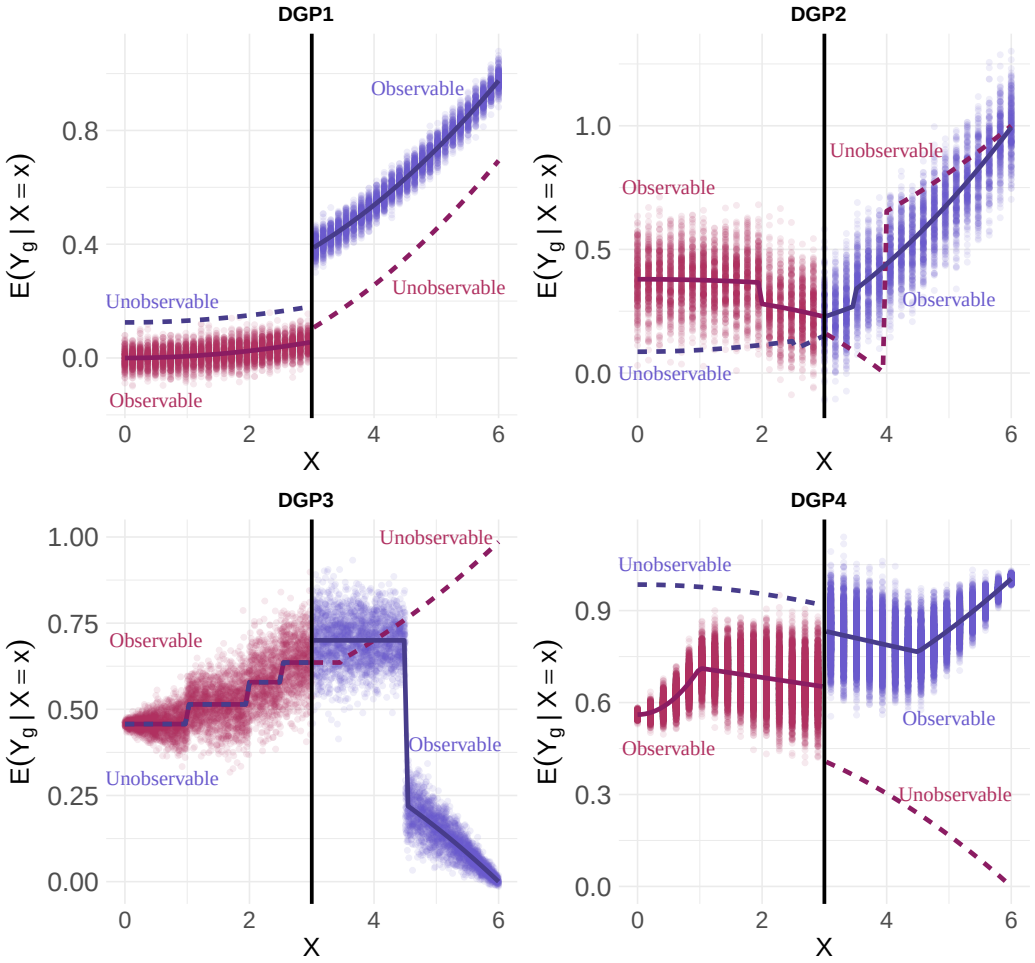

Figure 3.5: Data Generating Processes for Simulation (DGPs)

Note: The solid (dashed) lines are the observable (unobservable) potential outcome CEFs. The full population data with added noise is plotted as points.

For each population dataset, we take 1000 simulation samples of sample size $s \in \{100, 500, 1000\}$ with bandwidths $h \in \{0.5, 1.0, 1.5, 3.0\}$ for each DGP. In each simulation run, we estimate the local cutoff treatment effect bounds as discussed above. This involves, for example, estimating the left-sided limit $\lim_{x \to C^-} \mu_0^L(x) = \lim_{x \to C^-} \mathbb{E}[Y_0 | X = x]$ and $\mu_1^R(C)$.

For each each DGP, bandwidth, and sample size combination, I evaluate the distribution of estimates. For example, in Appendix Figure C.7 I show the estimated upper bounds on the local average treatment effect for DGP1. The estimates generally improve as bandwidth shrinks and sample size increases. For nearly all simulation

runs, the Local Kernel and Local Extremum cover the true upper bound on the local average treatment effect (the black vertical line) although the estimates are frequently upward biased. In this sense, the Local Kernel and Local Extremum estimators are "conservative." By contrast, the local linear estimator is less biased but has worse coverage properties. See Appendix C.5 for full local simulation results for each DGP plotted separately.

We next present information on each estimator across all simulations. In Table 3.1, I show results from the simulation study for each DGP, bandwidth, and sample size combination for the upper bound on the cutoff average treatment effect. I compute for each combination the finite sample bias, standard deviation, coverage of the upper bound, and coverage of the true treatment effect.

As seen in Table 3.1, the local kernel and local extremum estimators often achieve higher coverage of the upper bound across DGPs, bandwidths, and sample sizes. This is indicated by the green shading in the cells: for each row, the estimator(s) with green shaded cells in the coverage bound ("cov bound") column(s) achieve the highest coverage. As expected, the bias typically decreases as sample size increases and bandwidth decreases for each estimator and DGP. Note, however, that for some bandwidths and sample sizes the bias for the local kernel and local extremum estimators is larger, typically larger bandwidths and smaller sample sizes. Generally speaking, this corresponds to larger partially identified intervals and it is, in this sense, that the local kernel and local extremum estimators are conservative. The direction of the bias, moreover, is often in the expected direction. For example, because both DGP1 CEFs are monotone increasing, it is unsurprising that the upper bound estimates from the local extremum estimator are higher than the true upper bound and are decreasing as a function of the bandwidth.[13]

[13]That is, an estimate of the local minimum (at some bandwidth to the left of the cutoff) will result in larger upper bound than true if the estimated local minimum is smaller than $\lim_{x \to C^-} \mu_0(x)$, since it is being subtracted from the estimate of $\mu_1(C)$.

			Local Linear (LL)				Local Reverse Kernel (LRK)				Local Extremum (LE)			
dgp	**bandwidth**	**sample size**	**bias**	**sd**	**cov bound**	**cov true**	**bias**	**sd**	**cov bound**	**cov true**	**bias**	**sd**	**cov bound**	**cov true**
DGP1	0.5	100	0.041	0.019	0.996	1.000	0.064	0.018	1.000	1.000	0.122	0.014	1.000	1.000
		500	0.017	0.008	0.996	1.000	0.057	0.007	1.000	1.000	0.094	0.006	1.000	1.000
		1000	0.012	0.004	1.000	1.000	0.056	0.002	1.000	1.000	0.088	0.004	1.000	1.000
	1.0	100	0.036	0.015	0.992	1.000	0.119	0.011	1.000	1.000	0.233	0.019	1.000	1.000
		500	0.014	0.006	0.991	1.000	0.118	0.004	1.000	1.000	0.192	0.007	1.000	1.000
		1000	0.010	0.004	0.993	1.000	0.118	0.003	1.000	1.000	0.184	0.005	1.000	1.000
	1.5	100	0.030	0.015	0.982	1.000	0.184	0.011	1.000	1.000	0.344	0.022	1.000	1.000
		500	0.010	0.006	0.960	1.000	0.184	0.005	1.000	1.000	0.304	0.008	1.000	1.000
		1000	0.005	0.004	0.905	1.000	0.184	0.004	1.000	1.000	0.294	0.006	1.000	1.000
	3.0	100	0.014	0.014	0.837	1.000	0.419	0.025	1.000	1.000	0.724	0.032	1.000	1.000
		500	-0.008	0.006	0.112	1.000	0.419	0.011	1.000	1.000	0.714	0.015	1.000	1.000
		1000	-0.012	0.004	0.002	1.000	0.420	0.007	1.000	1.000	0.691	0.009	1.000	1.000
DGP2	0.5	100	-0.024	0.087	0.605	0.911	-0.061	0.118	0.599	0.651	0.019	0.094	0.647	0.999
		500	-0.030	0.087	0.740	0.797	-0.025	0.099	0.797	0.797	0.006	0.090	0.797	0.842
		1000	-0.021	0.075	0.743	0.868	-0.008	0.083	0.868	0.868	0.014	0.078	0.868	0.868
	1.0	100	-0.016	0.089	0.654	0.867	0.069	0.109	0.773	0.978	0.240	0.070	1.000	1.000
		500	-0.011	0.050	0.482	0.953	0.115	0.055	0.953	1.000	0.201	0.045	0.981	1.000
		1000	-0.009	0.026	0.223	0.989	0.124	0.028	0.989	1.000	0.198	0.025	0.989	1.000
	1.5	100	0.032	0.031	0.858	0.999	0.205	0.023	0.999	1.000	0.407	0.048	1.000	1.000
		500	0.003	0.012	0.595	1.000	0.207	0.009	1.000	1.000	0.337	0.017	1.000	1.000
		1000	-0.005	0.009	0.291	1.000	0.207	0.006	1.000	1.000	0.323	0.012	1.000	1.000
	3.0	100	0.016	0.038	0.711	0.993	0.484	0.045	1.000	1.000	0.863	0.070	1.000	1.000
		500	-0.015	0.013	0.111	1.000	0.487	0.016	1.000	1.000	0.827	0.028	1.000	1.000
		1000	-0.023	0.009	0.004	1.000	0.487	0.011	1.000	1.000	0.806	0.020	1.000	1.000
DGP3	0.5	100	0.334	0.304	0.992	0.992	0.280	0.314	0.751	0.751	0.530	0.239	1.000	1.000
		500	0.237	0.297	0.989	0.989	0.212	0.300	0.622	0.622	0.527	0.214	1.000	1.000
		1000	0.172	0.272	0.977	0.977	0.154	0.274	0.496	0.496	0.477	0.198	1.000	1.000
	1.0	100	0.067	0.060	0.982	0.982	0.042	0.052	0.990	0.990	0.345	0.058	1.000	1.000
		500	0.017	0.013	0.909	0.909	0.039	0.007	1.000	1.000	0.418	0.032	1.000	1.000
		1000	0.008	0.009	0.798	0.798	0.039	0.005	1.000	1.000	0.429	0.027	1.000	1.000
	1.5	100	0.056	0.030	0.965	0.965	0.086	0.016	1.000	1.000	0.382	0.044	1.000	1.000
		500	0.010	0.013	0.773	0.773	0.087	0.007	1.000	1.000	0.458	0.037	1.000	1.000
		1000	-0.001	0.009	0.452	0.452	0.087	0.005	1.000	1.000	0.482	0.032	1.000	1.000
	3.0	100	0.292	0.051	1.000	1.000	-0.300	0.039	0.000	0.000	0.374	0.040	1.000	1.000
		500	0.223	0.019	1.000	1.000	-0.303	0.017	0.000	0.000	0.442	0.036	1.000	1.000
		1000	0.207	0.013	1.000	1.000	-0.304	0.012	0.000	0.000	0.467	0.033	1.000	1.000
DGP4	0.5	100	-0.178	0.302	0.617	0.633	-0.250	0.317	0.139	0.633	-0.152	0.285	0.633	0.634
		500	-0.106	0.259	0.788	0.798	-0.142	0.264	0.025	0.798	-0.103	0.254	0.798	0.798
		1000	-0.064	0.216	0.866	0.873	-0.093	0.220	0.005	0.873	-0.065	0.213	0.873	0.873
	1.0	100	-0.055	0.234	0.816	0.840	-0.138	0.246	0.006	0.840	0.027	0.201	0.840	0.898
		500	0.005	0.107	0.970	0.973	-0.049	0.108	0.000	0.973	0.028	0.100	0.973	0.973
		1000	0.014	0.046	0.990	0.995	-0.035	0.047	0.000	0.995	0.026	0.045	0.995	0.995
	1.5	100	-0.015	0.188	0.890	0.906	-0.107	0.196	0.000	0.906	0.072	0.160	0.906	0.948
		500	0.020	0.023	0.990	0.999	-0.044	0.022	0.000	0.999	0.055	0.024	0.999	0.999
		1000	0.016	0.006	0.992	1.000	-0.043	0.003	0.000	1.000	0.036	0.009	1.000	1.000
	3.0	100	-0.692	0.050	0.001	0.003	-0.596	0.040	0.003	0.003	-0.347	0.042	0.003	0.998
		500	-0.745	0.015	0.000	0.000	-0.597	0.009	0.000	0.000	-0.366	0.017	0.000	1.000
		1000	-0.756	0.010	0.000	0.000	-0.597	0.006	0.000	0.000	-0.382	0.004	0.000	1.000

Table 3.1: Results of Simulation Study (Upper Bound)

Note: The table shows the finite-sample bias, standard deviation (sd), and coverage of Local Linear (LL), Local Reverse Kernel (LRK), and Local Extremum (LE) estimators for each DGP, bandwidth, and sample size 3-tuple. "cov bound" is the coverage rate for whether the true upper bound on the treatment effect is less than or equal to the estimated upper bound on the treatment effect, whereas "cov true" is coverage of whether the true treatment effect at the cutoff is less than or equal to the estimated upper bound on the treatment effect. For each row, if the estimator covers the bound with at least 0.95 probability, it is highlighted in green.

We also show the simulation results for the lower bound on the cutoff average treatment effect in Table 3.2. See Appendix C.6 for coverage information in plot formats for both the "cov bound" (coverage of the bound) and "cov true" (coverage of the true treatment effect).

			Local Linear (LL)				Local Reverse Kernel (LRK)				Local Extremum (LE)			
dgp	**bandwidth**	**sample size**	**bias**	**sd**	**cov bound**	**cov true**	**bias**	**sd**	**cov bound**	**cov true**	**bias**	**sd**	**cov bound**	**cov true**
DGP1	0.5	100	0.139	0.352	0.815	0.834	0.204	0.356	0.000	0.832	0.146	0.346	0.712	0.838
		500	0.008	0.126	0.970	0.982	0.062	0.127	0.000	0.982	0.018	0.125	0.336	0.982
		1000	-0.006	0.003	0.985	1.000	0.045	0.002	0.000	1.000	0.004	0.003	0.069	1.000
	1.0	100	0.008	0.162	0.963	0.970	0.125	0.165	0.000	0.970	0.009	0.156	0.878	0.985
		500	-0.011	0.004	0.998	1.000	0.096	0.003	0.000	1.000	-0.002	0.006	0.617	1.000
		1000	-0.008	0.003	1.000	1.000	0.096	0.002	0.000	1.000	0.001	0.004	0.354	1.000
	1.5	100	-0.019	0.086	0.987	0.992	0.161	0.086	0.000	0.992	-0.015	0.083	0.903	0.997
		500	-0.016	0.004	1.000	1.000	0.153	0.004	0.000	1.000	-0.005	0.007	0.744	1.000
		1000	-0.014	0.003	1.000	1.000	0.153	0.003	0.000	1.000	0.000	0.005	0.476	1.000
	3.0	100	-0.055	0.011	1.000	1.000	0.372	0.024	0.000	1.000	-0.016	0.021	0.755	1.000
		500	-0.043	0.005	1.000	1.000	0.372	0.011	0.000	1.000	-0.012	0.011	0.889	1.000
		1000	-0.041	0.003	1.000	1.000	0.372	0.007	0.000	1.000	-0.005	0.007	0.742	1.000
DGP2	0.5	100	-0.373	0.334	0.996	0.999	-0.262	0.359	0.646	0.997	-0.386	0.333	1.000	1.000
		500	-0.213	0.297	0.997	1.000	-0.149	0.302	0.508	1.000	-0.222	0.293	1.000	1.000
		1000	-0.148	0.252	1.000	1.000	-0.096	0.254	0.411	1.000	-0.156	0.249	1.000	1.000
	1.0	100	-0.301	0.296	1.000	1.000	-0.080	0.308	0.227	0.389	-0.408	0.258	1.000	1.000
		500	-0.115	0.157	1.000	1.000	0.057	0.156	0.047	0.056	-0.164	0.146	1.000	1.000
		1000	-0.073	0.079	1.000	1.000	0.084	0.077	0.011	0.011	-0.113	0.074	1.000	1.000
	1.5	100	-0.119	0.051	0.999	1.000	0.106	0.036	0.001	0.077	-0.405	0.075	1.000	1.000
		500	-0.051	0.020	0.996	1.000	0.109	0.013	0.000	0.002	-0.254	0.024	1.000	1.000
		1000	-0.035	0.014	0.994	1.000	0.109	0.009	0.000	0.000	-0.221	0.017	1.000	1.000
	3.0	100	-0.199	0.072	1.000	1.000	0.345	0.065	0.007	0.007	-0.477	0.089	1.000	1.000
		500	-0.124	0.020	1.000	1.000	0.349	0.018	0.000	0.000	-0.341	0.036	1.000	1.000
		1000	-0.109	0.014	1.000	1.000	0.350	0.012	0.000	0.000	-0.278	0.022	1.000	1.000
DGP3	0.5	100	0.102	0.167	0.546	1.000	0.156	0.178	0.308	0.882	-0.085	0.118	0.671	1.000
		500	0.098	0.166	0.650	1.000	0.119	0.171	0.450	0.930	-0.165	0.098	0.925	1.000
		1000	0.072	0.153	0.734	0.999	0.086	0.156	0.558	0.963	-0.170	0.099	0.901	1.000
	1.0	100	-0.041	0.038	0.967	1.000	0.000	0.034	0.563	0.993	-0.168	0.047	0.993	1.000
		500	-0.020	0.009	0.985	1.000	-0.002	0.005	0.675	1.000	-0.217	0.039	1.000	1.000
		1000	-0.014	0.007	0.985	1.000	-0.002	0.004	0.741	1.000	-0.231	0.038	1.000	1.000
	1.5	100	-0.042	0.022	0.973	1.000	0.000	0.011	0.509	1.000	-0.157	0.038	1.000	1.000
		500	-0.021	0.010	0.985	1.000	-0.001	0.005	0.545	1.000	-0.205	0.037	1.000	1.000
		1000	-0.015	0.006	0.989	1.000	0.000	0.003	0.533	1.000	-0.225	0.037	1.000	1.000
	3.0	100	0.100	0.036	0.001	1.000	-0.444	0.038	1.000	1.000	-0.694	0.011	1.000	1.000
		500	0.138	0.016	0.000	1.000	-0.447	0.017	1.000	1.000	-0.705	0.005	1.000	1.000
		1000	0.146	0.011	0.000	1.000	-0.448	0.012	1.000	1.000	-0.708	0.003	1.000	1.000
DGP4	0.5	100	-0.197	0.152	0.997	1.000	-0.145	0.163	0.940	1.000	-0.256	0.132	1.000	1.000
		500	-0.102	0.134	0.991	1.000	-0.088	0.136	0.999	1.000	-0.140	0.125	1.000	1.000
		1000	-0.068	0.112	0.995	1.000	-0.062	0.113	1.000	1.000	-0.101	0.107	1.000	1.000
	1.0	100	-0.127	0.117	0.995	1.000	-0.106	0.121	1.000	1.000	-0.336	0.086	1.000	1.000
		500	-0.042	0.056	0.997	1.000	-0.062	0.053	1.000	1.000	-0.172	0.046	1.000	1.000
		1000	-0.024	0.025	0.991	1.000	-0.055	0.024	1.000	1.000	-0.135	0.023	1.000	1.000
	1.5	100	-0.107	0.095	0.995	1.000	-0.104	0.094	1.000	1.000	-0.355	0.069	1.000	1.000
		500	-0.035	0.017	0.993	1.000	-0.074	0.012	1.000	1.000	-0.208	0.017	1.000	1.000
		1000	-0.023	0.009	0.997	1.000	-0.073	0.005	1.000	1.000	-0.168	0.010	1.000	1.000
	3.0	100	-0.478	0.033	1.000	1.000	-0.306	0.024	0.997	1.000	-0.519	0.034	1.000	1.000
		500	-0.450	0.011	1.000	1.000	-0.306	0.007	1.000	1.000	-0.501	0.022	1.000	1.000
		1000	-0.443	0.008	1.000	1.000	-0.306	0.005	1.000	1.000	-0.465	0.011	1.000	1.000

Table 3.2: Results of Simulation Study (Lower Bound)

Note: The table shows the finite-sample bias, standard deviation (sd), and coverage of Local Linear (LL), Local Reverse Kernel (LRK), and Local Extremum (LE) estimators for each DGP, bandwidth, and sample size 3-tuple. "cov bound" is the coverage rate for whether the true lower bound on the treatment effect is greater than or equal to the estimated lower bound on the treatment effect, whereas "cov true" is coverage of whether the true treatment effect at the cutoff is greater than or equal to the estimated lower bound on the treatment effect. For each row, if the estimator covers the bound with at least 0.95 probability, it is highlighted in green.

Further Simulations: More Bandwidths, Missing Data, Measurement Error

Next, I conduct additional simulations to estimate bounds on local treatment effects. Specifically, I conduct three sets of additional simulations: more bandwidths, missing data, and measurement error.

The first set of additional simulations use the same design as before but more bandwidths. Using the four DGPs, more bandwidths allows us to better see how coverage and efficiency (partial ID region widths) vary as the bandwidth varies. The local extremum and local reverse kernel estimators may have better coverage

properties but result in larger partial ID regions and are, in this sense, conservative estimation strategies. Indeed, in Appendix Figure C.19 and Appendix Figure C.20 we see that the local extremum and local reverse kernel estimators often obtain high coverage rates for one or both bounds, but often at the cost of efficiency for higher bandwidth selections. In settings where coverage or bandwidth-robustness are important, estimators adapted to bounding may be practically useful.

The second set of additional simulations uses the four DGPs specified above subject to missing data. Specifically, we conduct simulations where data near the cutoff (running variable values $|X - 3| \leq 0.5$) are missing. Such missingness may occur in practice or in an econometric exercise where data near a cutoff is possibly manipulated (donut RDD). Coverage and partial ID region widths are shown in Appendix Figure C.21 and Appendix Figure C.22. In general, the results suggest coverage rates for both bounds are low for each estimator, with some success, however, in covering either the upper or lower bound. In such settings, missingness near the boundary may result in severe extrapolation errors if one proceeds with the continuity-based approach and a local linear estimator (depending on the DGP). In such situations, proceeding in the alternative (i.e., assuming DGP is monotone increasing, and then assuming DGP is monotone decreasing) in the missing region, and reporting resulting bounds may be more credible than extrapolation.

The third set of additional simulations uses the original DGPs subject to running variable measurement error. The coverage and partial ID region widths are shown in Appendix Figure C.23 and Appendix Figure C.24. In the four DGPs, only the local extremum estimator achieves nominal 95% coverage rates at any bandwidth value. Upper and lower bounds alone may be estimated adequately, however, for certain DGPs using the local linear or local reverse kernel estimators. Region widths are generally much larger than in our baseline simulations (esp. for the local extremum estimator), suggesting that in situations which across-the-board measurement error, high coverage of partially-identified bounds comes at an efficiency cost.

3.4 Extensions: Extrapolation to ATE & Fuzzy Monotone RDD

Although regression discontinuity designs are frequently lauded for a high degree of internal validity, policy-makers are not only interested in the effect of a policy for individuals near the cutoff. But point-identifying common estimands like the average treatment effect (ATE) or average treatment effect on the treated (ATT) may require heroic extrapolation, which may be implausible in a sharp regression discontinuity

design where there is a complete failure of overlap. Existing methods to extrapolate treatment effects in an RD setting generally assume constant treatment effects,[14] polynomial functional forms on the potential outcome CEFs,[15] or otherwise restrict the heterogeneity of treatment effects across individual units.[16]

But under our two assumptions, we can partially identify the more general estimand ATE by extrapolation. Indeed, a partial identification approach more fully exploits the modeled information to get lower and upper bounds for larger sub-populations. In the case of global monotonicity, this is the entire population; but in the case of monotonicity on only some of the support of X, this corresponds to the sub-population with running variable x-values that are compatible with monotone potential outcome CEFs $(\mu_1(x), \mu_0(x))$.

I also briefly consider the fuzzy regression discontinuity design. In this case, further assumptions are needed to bound the local cutoff treatment effect (or to bound the ATE under global monotonicity assumptions). We consider bounds when there is monotone treatment selection ("MTS"), a partial identification assumption discussed in Manski (2007).

ATE Partial Identification

Under the assumptions of global monotonicity (μ_1, μ_0 both weakly increasing) and boundedness we can obtain bounds on ATE. Define μ_1, μ_0 as before:

$$\mu_1(x) = \mathbb{E}[Y_1|X = x] = \mathbf{1}\{x < C\}\mu_1^L(x) + \mathbf{1}\{x \geq C\}\mu_1^R(x)$$
$$\mu_0(x) = \mathbb{E}[Y_0|X = x] = \mathbf{1}\{x < C\}\mu_0^L(x) + \mathbf{1}\{x \geq C\}\mu_0^R(x)$$

and suppose that the forcing variable X is discrete.[17] Recall that, by the sharp assignment rule $W_i = \mathbf{1}\{X_i \geq C\}$, only $\mu_1^R(x)$ (for $x \geq C$) and $\mu_0^L(x)$ (for $x < C$) are observable.

Theorem 2 gives bounds on the ATE under the monotonicity and boundedness assumptions:

[14] See, e.g., Wooldridge (2010) discussing constant treatment effect.

[15] See, e.g., Angrist and Rokkanen (2012) for a good parametric extrapolation example, which considers parametric extrapolation by assuming a specific functional form of the CEF. As noted therein, extrapolation based on functional form restrictions on the CEF is unlikely to achieve point identification. By contrast, here we are focused on partial identification and more flexible shape restrictions.

[16] See Cattaneo, Idrobo, and Titiunik (2024) for further details.

[17] A similar approach works for continuously distributed forcing variables. Replace the summations by integrals in the obvious way and use the pdf of X. We use discrete X for expositional purposes and translation into practical settings.

Theorem 2 (ATE Bounds)**.** *Suppose $\mu_1(x), \mu_0(x)$ are monotone increasing for all $x \in supp(X)$. Let M_0, M_1 be global bounds as in Assumption 2. Then the average treatment effect (ATE) is bounded by:*

$$\underline{ATE} \leq ATE \leq \overline{ATE}$$

where:

$$\overline{ATE} = \sum_{x<C} Pr(X = x) \cdot \left(\mu_1^R(C) - \mu_0^L(x)\right) + \sum_{x \geq C} Pr(X = x) \cdot \left(\mu_1^R(x) - \lim_{x \to C^-} \mu_0^L(x)\right)$$

$$\underline{ATE} = \sum_{x<C} Pr(X = x) \cdot (M_0 - \mu_0^L(x) + \sum_{x \geq C} Pr(X = x) \cdot (\mu_1^R(x) - M_1)$$

The proof (in Appendix C.2) is a straightforward application of the LIE and the assumptions. The results for other cases (e.g., both μ_1, μ_0 monotone decreasing) are similar but omitted. Appendix C.2 also works through a specific example (with a visual example) on bounding the ATE under our assumptions.

Partial Identification of $ATE(\Omega)$ Under Local Monotonicity

Monotonicity on the entire support of the forcing variable is a strong shape restriction. However, there may be a region Ω for which a monotone CEF is consistent with the observable data. The goal, then, is to obtain an estimated region $\hat{\Omega} \subseteq \text{supp}(X)$ that is consistent with monotone CEFs from the observable data.

When μ_1, μ_0 are monotone for a proper subset $\Omega \subsetneq \text{supp}(X)$, we can partially identify the average treatment effect on that subset. Because each CEF in a sharp RDD is only observable on one side of the cutoff C, we also use a symmetry assumption on the monotonicity of each potential outcome CEF across the cutoff. We consider the following local monotonicity and symmetry assumption:

Assumption 3 (Local Symmetric Monotonicity)**.** *The conditional mean functions μ_1, μ_0 are (locally) weakly monotone increasing in x for all $x \in \Omega_1, \Omega_0 \subsetneq supp(X)$, respectively, where $\Omega_1 = [C - a, C + a]$ and $\Omega_0 = [C - b, C + b]$ for $a, b \in \mathbb{R}_+$.*

Under Assumption 3, we cannot partially identify the ATE, but we can partially identify the average treatment effect for units with $X \in \Omega = \Omega_1 \cap \Omega_0$:

$$ATE(\Omega) = \mathbb{E}[Y_1 - Y_0 | X \in \Omega]$$

The proof is analogous to the proof of Theorem 2.

If only local monotonicity holds, the econometrician must know in advance or estimate the monotone region Ω. By symmetry, we only need to find the a, b values for which the observable part of μ_1, μ_0 is compatible with monotonicity. Then we can partially identify $ATE(\Omega)$ for $\Omega = \Omega_1 \cap \Omega_0$. If Ω is estimated, the estimated monotone region $\hat{\Omega}$ is dependent on the model specification and research context. One possibility is to proceed conservatively and consider the smallest estimated interval $\hat{\Omega}$ that holds under several plausible models.[18]

Nonparametric Estimation & Inference: ATE Bounds

The quantities M_0, M_1 are assumed known. Estimating the bounds $\overline{ATE}$ and $\underline{ATE}$ can be done non-parametrically. To estimate the upper bound $\overline{ATE}$, for example, the econometrician needs to estimate $\mu_1^R(C)$, $\lim_{x \to C^-} \mu_0^L(x)$, $\mu_0^L(x)$ for $x < C$, and $\mu_1^R(x)$ for $x \geq C$. I considered estimation methods for the first two estimands above. For the last two estimands, a variety of estimation strategies are available. For instance, if the running variable X is discretely distributed and there are many observations for each unique x-value, one can estimate $\mu_1^R(x) = \mathbb{E}[Y_1 | W = 1, X = x]$ for each $x \geq C$ with the sample analogue and use bootstrap procedures to obtain CIs.

Fuzzy RDD

I also consider identification in a fuzzy RDD setting where the potential outcome CEFs are monotone. As before, we assume the running variable X is discrete for expositional simplicity. In the fuzzy RDD, there are treated and control observations on both sides of a cutoff C. Although the treatment probability does not change from zero to one at the cutoff, there is still a discontinuous change in the assignment probability at the cutoff. Importantly, this means that ignorability may not hold, i.e.:

$$\mathbb{E}[Y_g | W = w, X = x] \neq \mathbb{E}[Y_g | X = x]$$

[18] In Manski (2007), for example, monotonicity in the observable CEF is a testable implication of the monotone treatment response and monotone selection assumptions. There, the author suggests using the bootstrap to compute a 95 percent confidence band for the observed CEF. If the confidence band contains everywhere monotone functions, the data is compatible with a monotone CEF on its domain.

for $g = 0, 1$ and $w = 0, 1$. Because treatment W is not a deterministic function of X, there are therefore four CEFs of interest:

$$\mu_{00}(x) = \mathbb{E}[Y_0|W = 0, X = x]$$
$$\mu_{01}(x) = \mathbb{E}[Y_0|W = 1, X = x]$$
$$\mu_{10}(x) = \mathbb{E}[Y_1|W = 0, X = x]$$
$$\mu_{11}(x) = \mathbb{E}[Y_1|W = 1, X = x]$$

Of course, we can estimate (or observe for $x \in \text{supp}(X)$) $\mu_{00}(x), \mu_{11}(x)$ for x-values less than C and x-values greater than or equal to C, respectively. But $\mu_{01}(x), \mu_{10}(x)$ are not observed for any x-values: for instance, $\mu_{10}(x)$ is the conditional mean of the treated potential outcome among untreated units with $X = x$.

I consider the following assumption to partially identify the ATE in a fuzzy design. The assumption is a conditional form of monotone treatment selection ("MTS") as in Manski (2007):

Assumption 4 (Conditional Monotone Treatment Selection)**.** *For $g = 0, 1$ and $w = 0, 1$ and any $x \in supp(X)$, if $w \geq w'$, then*

$$\mathbb{E}[Y_g|W = w, X = x] \geq \mathbb{E}[Y_g|W = w', X = x].$$

This assumption relates the unobservable CEFs (μ_{01}, μ_{10}) to the observable ones. I further assume that μ_{11}, μ_{00} are monotonic increasing and that outcomes are uniformly bounded above and below by M_1, M_0, resp., as before. In this case, treatment effect bounds may be obtained, which are formally discussed in Appendix C.9.

Essentially, the fuzziness allows for "tightening" the bounds on either side of the cutoff. This is visualized with an example in Appendix Figure C.29. The two large dots are examples of treated units with $x < C$ (green dot) and control units with $x > C$ (purple dot). The fact that we can observe some treated units to the left of the cutoff and some control units to the right of the cutoff results in step functions that more closely approximate the unobservable CEFs. The bounds are, therefore, tighter in the fuzzy case than the sharp case because the bounds utilize this extra information. See Appendix C.9 for further discussion of this example.

3.5 Application: Energy Efficiency and Property Price Discontinuities

This section turns to an application on the impact of energy efficiency ratings on property prices. Specifically, this section estimates bounds on property prices as a

function of discontinuous energy efficiency rating bands (Sejas-Portillo, Moro, & Stowasser, 2025). Using data on UK residential properties, Sejas-Portillo, Moro, and Stowasser (2025) estimate the effects of simplified energy efficiency bands (ranging from A to G) which are themselves determined by underlying energy efficiency scores (1-100). The bands are "arbitrary and technically redundant," yet Sejas-Portillo, Moro, and Stowasser (2025) find "sizeable price discontinuities at rating band thresholds." As described in Sejas-Portillo, Moro, and Stowasser (2025), the result is consistent with a behavioral effect, where less informative rating bands induce price premiums that "are not warranted by actual gains in energy efficiency."

For our purposes, the application is attractive due to the CEF being plausibly a plausibly monotone function of a discrete running variable. Indeed, as visualized in Figure 3 of Sejas-Portillo, Moro, and Stowasser (2025) and reproduced in Appendix Figure C.25, the average property price per meter squared is increasing as a function of the running SAP (Standard Assessment Procedure) score. In Appendix Figure C.25, the multiple cutoffs as a function of the running SAP score are plotted as red dashed vertical lines and labeled for each energy band rating (A to G). Average prices are plotted at each SAP score in blue points, with larger points representing more available data.

The cutoff effect estimates in the baseline specification are reproduced in Table 3.3. The local treatment effect estimates τ at the energy band cutoffs, as well as estimation details, are provided for the first three energy band cutoffs.

	E–D	F–E	G–F
τ	0.019	0.029	0.023
Robust SE	(0.002)	(0.003)	(0.003)
RBC 95% CI	[0.017, 0.027]	[0.024, 0.041]	[0.018, 0.035]
RBC p-value	[0.000]	[0.000]	[0.000]
BW estimate	3.710|3.710	4.196|4.196	3.816|3.816
BW bias	5.470|5.470	6.807|6.807	6.996|6.996
N	1,325,863|3,413,478	299,568|1,325,863	65,293|299,568
Effective N	404,837|740,804	113,639|245,331	16,631|31,631

Table 3.3: Local Linear RDD Price Discontinuity Estimates (Sejas-Portillo, Moro, and Stowasser (2025) Table 3)

Next, using the data for these first three cutoffs (E–D, F–E, and G–F), I estimate

the bounds on treatment effects as a function of the bandwidth parameter. The results from each estimator and cutoff are shown in Figure 3.6. The horizontal black dashed line is the `rdrobust` (Calonico, Cattaneo, & Titiunik, 2014, 2015) estimate with default selections. As seen in the right-most G–F cutoff panel of Figure 3.6, for example, the estimated bounds from the local extremum and local reverse kernel estimators are higher than the `rdrobust` estimates and the estimates in Sejas-Portillo, Moro, and Stowasser (2025). This means that, under a running variable monotonicity assumption, the true treatment effect at the G–F cutoff may be even larger than the baseline estimate in Sejas-Portillo, Moro, and Stowasser (2025). Upper bound treatment effect estimates from the smallest bandwidths are similar (only slightly larger) to the rdrobust and Sejas-Portillo, Moro, and Stowasser (2025) estimates. In other cases, the upper bound estimates and the estimated effects from Sejas-Portillo, Moro, and Stowasser (2025) are similar, though the local linear upper bound estimates at many bandwidths are below the rdrobust estimates.

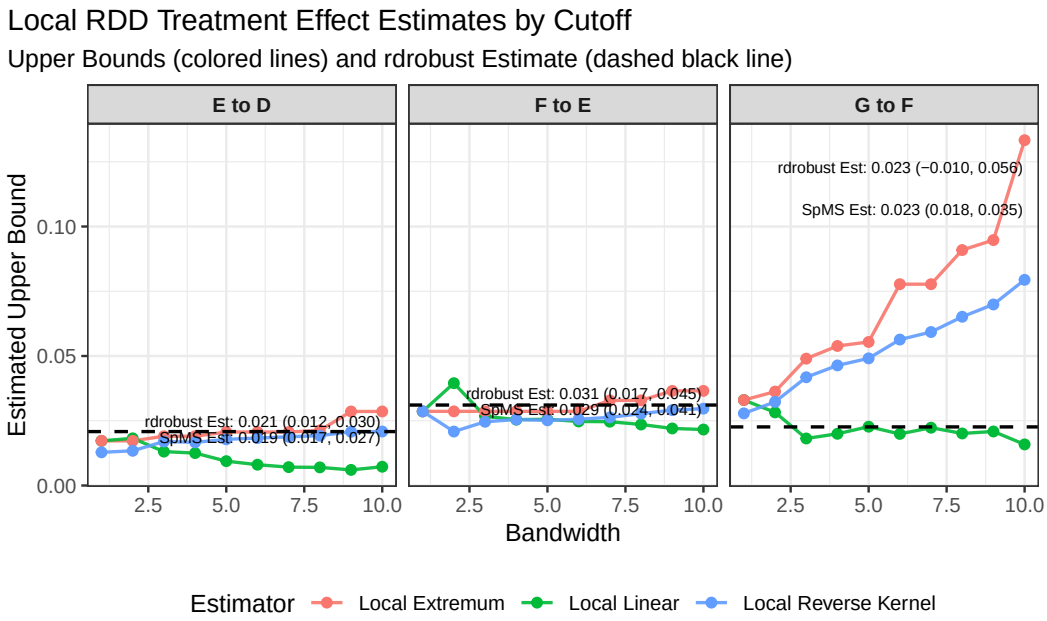

Figure 3.6: Local RDD Treatment Effect Estimates by Cutoff

Note: The estimates from rdrobust are plotted as black horizontal dashed line. Robust estimates and estimates Sejas-Portillo, Moro, and Stowasser (2025) are also shown as annotations in each panel.

Finally, I further conducted placebo tests at non-treatment values of the running variable. In continuity-based designs, the idea is to look for continuity violations at running variable values to test the underlying CEF continuity assumptions (Imbens

& Lemieux, 2008).[19] Here, placebo tests are also used (with each estimator) to see if jumps at the cutoffs are "large" in the sense of having larger estimated treatment effects than jumps at nearby placebo running variable values. Placebo jump distributions are plotted in Appendix C.8 for running variable values near the three cutoffs (E–D, F–E, and G–F).[20] Under each estimator and in each of the three cutoffs, the estimated upper bounds (and in many cases the Sejas-Portillo, Moro, and Stowasser (2025) estimates themselves) are larger than similarly estimated effect sizes at placebo running variable values. This lends further credibility to the Sejas-Portillo, Moro, and Stowasser (2025) findings that effects at cutoff values are both substantively large even under the alternative identifying assumptions here.

3.6 Conclusion

This paper develops new methodology that applies generally to regression discontinuity designs (RDD). The usual continuity assumption in the discontinuity design is modified to allow for discontinuities in the potential outcome conditional expectation functions (CEFs) at the cutoff determining treatment. Instead, monotonicity in the CEFs and boundedness in the outcome-of-interest is assumed. Under these assumptions, researchers can partially identify the treatment effect at the cutoff. Additional bounds in fuzzy designs or extrapolated bounds on ATE-type parameters is also permitted under certain assumptions.[21]

Monotone counterfactual assumptions may be useful for a variety of empirical settings. First, the monotonicity assumption is attractive in the RDD setting where the underlying data generating process is well-approximated by a polynomial model.[22] Second, monotonicity may hold where continuity-based identification is implausible: for instance, continuity may be untenable in discrete data generating processes, where there is manipulation of the running variable (McCrary, 2008), or where placebo tests of the continuity assumption at non-cutoff values are violated. Third, continuity may be difficult to visually infer or statistically verify (Korting

[19] Another possibility is placebo testing in time (Lee, Moretti, & Butler, 2004), where we expect to see "no systematic differences in any variable determined prior to the experiment."

[20] That is, all running variable values on either side of a given cutoff but not beyond the next cutoff.

[21] Appendix C.2 and Appendix C.9 further discuss possible methods for extrapolation to ATE and partial identification in the fuzzy RDD setup.

[22] There is guaranteed to be a region Ω where the CEFs are monotone for correctly specified polynomial models around a treatment cutoff. Moreover, in that region we can obtain bounds on $ATE(\Omega) = \mathbb{E}[Y_1 - Y_0 | X_i \in \Omega]$, as discussed above. Although continuity is also a plausible assumption in a polynomial model, continuity alone does not permit extrapolation on non-cutoff treatment effects (i.e., $ATE(\Omega)$).

et al., 2023), whereas monotonicity may be theoretically motivated and empirically tested.[23] Finally, the robustness of causal conclusions from different counterfactual assumptions yields credibility to a study and may strengthen resulting policy recommendations.

As with all counterfactual assumptions, monotonicity is no panacea and the reasonableness of certain assumptions are context-dependent. Partially identified sets of treatment effects or policy parameters under monotonicity may not be helpful where precise estimation is a first-order concern, or where more informative bounds are available under other tenable assumptions. Plausible theoretical or structural models, for instance, may incorporate monotonicity assumptions, ground counterfactual analysis, and output estimated bounds on treatment effect parameters. Policymakers may use the resulting causal bounds to gauge high-payoff reform areas and researchers may show estimated effects are plausible directionally and in magnitude under bounding assumptions.

References

Angrist, J., & Rokkanen, M. (2012). Wanna Get Away? RD Identification Away from the Cutoff. *NBER*, *Working Paper 18662*. https://www.nber.org/system/files/working_papers/w18662/w18662.pdf

Babii, A., & Kumar, R. (2023). Isotonic Regression Discontinuity Designs. *Journal of Econometrics*, *234*(2), 371–393. https://doi.org/https://doi.org/10.1016/j.jeconom.2021.01.008

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, *82*(6), 2295–2326. https://doi.org/https://doi.org/10.3982/ECTA11757

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs. *R Journal*, *7*(1), 38–51. https://rdpackages.github.io/references/Calonico-Cattaneo-Titiunik_2015_R.pdf

[23] As discussed in Korting et al. (2023), there is strong experimental evidence that detecting continuities or discontinuities is visually difficult, even for experts. Whether a particular RDD passes the "eye test" is highly dependent on data visualization choices under the control of the researcher. This means that, even for the observable portion of each potential outcome CEF, it is often difficult to detect discontinuities. Moreover, despite recommendations that applied researchers test for discontinuities away from the cutoff (Imbens & Lemieux, 2008), it is unclear how widespread the practice. By contrast, it may be possible to visually infer whether a CEF is compatible with monotonicity, and there are formal tests for whether a confidence band contains a monotone function.

Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2024). A Practical Introduction to Regression Discontinuity Designs: Extensions. *Cambridge Elements: Quantitative and Computational Methods for Social Science*.

Finkelstein, M., Tucker, H. G., & Veeh, J. A. (2000). Conservative Confidence Intervals for a Single Parameter. *Communications in Statistics-Theory and Methods*, *29*(8), 1911–1928. https://doi.org/https://doi.org/10.1080/03610920008832585

Gelman, A., & Imbens, G. (2018). Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, *37*(3), 447–456. https://doi.org/https://doi.org/10.1080/07350015.2017.1366909

Goldsmith-Pinkham, P. (2026). Tracking the Credibility Revolution Across Fields. *NBER*, *Working Paper 35051*. https://www.nber.org/papers/w35051

Imbens, G. W., & Lemieux, T. (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics*, *142*(2), 615–635. https://doi.org/https://doi.org/10.1016/j.jeconom.2007.05.001

Kolesar, M., & Rothe, C. (2018). Inference in Regression Discontinuity Designs with a Discrete Running Variable. *American Economic Review*, *108*(8), 2277–2304. https://doi.org/DOI:10.1257/aer.20160945

Korting, C., Lieberman, C., Matsudaira, J., Pei, Z., & Shen, Y. (2023). Visual Inference and Graphical Representation in Regression Discontinuity Designs. *Quarterly Journal of Economics*, *138*(3), 1977–2019. https://doi.org/https://doi.org/10.1093/qje/qjad011

Lee, D. S., Moretti, E., & Butler, M. J. (2004). Do Voter Affect or Elect Policies? Evidence from the U.S. House. *Quarterly Journal of Economics*, *119*(3), 807–859. https://www.princeton.edu/~davidlee/wp/voterspolicies.pdf

Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.

McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, *142*(2), 698–714. https://doi.org/https://doi.org/10.1016/j.jeconom.2007.05.005

Noack, C., & Rothe, C. (2023). Donut Regression Discontinuity Designs. https://arxiv.org/pdf/2308.14464

Sejas-Portillo, R., Moro, M., & Stowasser, T. (2025). The Simpler the Better? Threshold Effects of Energy Labels on Property Prices and Energy Efficiency Investments. *American Economic Journal: Applied Economics*, *17*(1), 41–89. https://doi.org/DOI:10.1257/app.20220329

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd). The MIT Press.

Chapter 4

COURTING THE ACADEMY: THE JUDICIAL ROLE IN POPULARIZING LEGAL SCHOLARSHIP

4.1 Introduction

The Supreme Court's recent decision in *Franklin v. New York*[1] offers a glimpse into the contemporary relationship between the judiciary and the legal academy. In *Franklin*, the Court denied *certiorari* to reconsider Confrontation Clause jurisprudence under the Sixth Amendment. Yet Justices Alito and Gorsuch issued opinions signaling a willingness to overturn a leading Confrontation Clause case.[2] Why? Because, Justice Alito writes, "recent scholarship ... casts doubt on key aspects of *Crawford's* reasoning."[3]

The Court's encouragement of particular legal scholarship is not new. Indeed, a large amount of empirical legal scholarship studies how judges use legal scholarship. What is less studied is the academic reaction to judicial acknowledgment. This reverse relationship is just as important: despite criticisms that the legal academy is disengaged with the practical matters of lawyering and the judiciary, individual academics may hold judicial mentions in high esteem. Academics with judicial ambitions can be incentivized to produce work that they believe the judiciary will invoke. To provide for the judiciary, academics attend to the judiciary's preferences. Judges can signal a need for further academic work by mentioning certain legal pieces in their opinions. Other judges may produce the scholarship themselves.[4] Still others may convey signals by indirect means, e.g., through their clerks who later become academics.

The purpose of this paper is quantification of the extent to which the academy heeds the judiciary's call. More specifically, we consider the courts' role in the rise of legal realist scholarship in the early 20th century and the rise of originalist scholarship in the late 20th century. We chose these specific academic movements for their historical and contemporary relevance and also because their adherents and

[1] *Franklin v. New York*, 604 U.S. (2025).
[2] *Crawford v. Washington*, 541 U.S. 36 (2004).
[3] *Franklin v. New York*, 604 U.S. (2025) (Alito, J., concurring).
[4] Consider Judges Posner and Easterbrook, or Justice Scalia.

detractors are (relatively) readily identifiable.[5] The objects (or units) of analysis are scholarly works (books, chapters, or articles) that are most relevant to legal realism or originalism. For each work, we collected time series data on judicial mentions[6] and an annual series of academic citations. Those specific series are obtained with a view to estimating how much of scholarly attention is attributable to court mentions. Although citations are an imperfect measure of scholarly attention, they are an oft-used proxy for academic popularity.

In contrast to citation analyses which explore the divergence between the academy and judiciary,[7] our aim is to investigate the more specific question of whether academic citations respond to judicial acknowledgment. The reason to choose such a specific aim is so that we may employ the machinery of causal inference and make causal claims, albeit with caution. In general, a causal relationship between judicial mentions and academic citations is difficult to pin down. For example, the two series may be simultaneously affected by the politics of judicial appointment.[8] Academics who see the political tides changing may get ahead of it by producing scholarly work that is adherent to a particular judicial movement.[9] Despite these difficulties, there is much variation in how[10] (and how often) works are mentioned by the judiciary, with some works mentioned much more often than others.

Challenges to causal inference are encountered across the social sciences. Because of the insolubility of many such challenges, there is a preference towards designing and implementing experiments whenever possible. Unfortunately, the question at hand is not easily amenable to experimentation. Such is the case for many questions concerning the judiciary, not least because diversions from normal operations even for scientific purposes may be seen as an illegal or unethical miscarriage of justice.[11] Yet, experimentation difficulties should not be fatal to the study of these

[5] Although there are a great many attempts by academics to launch legal movements, it is difficult to round up all such attempts and/or ascertain the boundaries of particular movements.

[6] We use the term "judicial mentions" instead of the more familiar "judicial citations" to avoid confusion between citations arising from the judiciary and the academy.

[7] See Merritt and Putnam (1995) for an example of such an analysis.

[8] Both realist and originalist movements were colored by the political leanings of the Administrations that nominated adherent judges to the court. The Roosevelt Administration nominated several legal realists and the Reagan Administration nominated originalists.

[9] They may also do so in anticipation of a judicial appointment or a blockbuster court case.

[10] Judges and their clerks may also use, for instance, citation counts as a heuristic in determining whether the work should be mentioned in their opinion or not. This potentially summons the pernicious problem of reverse causality. The possibility is not that the judicial mentions increased the scholarly popularity of the work, but its own baseline level of scholarly popularity.

[11] Cf. Ho (2005) at 1999 ("The problem for legal scholars and social scientists is that laboratory experiments are often infeasible, expensive, or unethical.")

processes, provided the counterfactual assumptions underpinning the interpretation of the derived estimates as causal hold. At its core, the counterfactual method here separates works into two groups:[12] a treatment group (works that mentioned by the judiciary) and a control group (works not mentioned).[13] The specification of such groups is necessary since it is the citation outcome of the control group that is used to construct counterfactual citations for a treatment group under alternative (i.e., counterfactual) judicial mention scenarios. But what counterfactual assumptions should one use here?

Since different causal assumptions lead to different estimates, we choose a bounding method that nests several of these causal assumptions. Our method,[14] which we call the bounded deviations approach, allows for a wide range of counterfactual outcomes in citation trajectories. Despite the statistical jargon, the counterfactual intuition is simple: absent treatment, treated units would've experienced an outcome "close to" or "nearby" the control outcomes. The upshot is that, rather than reporting a result like "the causal effect is X" under strong counterfactual assumptions, we report a range ("the causal effect is between X and Y") under weaker assumptions.[15] An important advantage of this bounding method is that it renders itself to sensitivity analysis almost immediately: do the conclusions change if confidences in the

[12] Separation of units into treatment and control groups might not be straightforward, particularly in dynamic settings such as ours. The process requires coarsening of the paths since some works receive a disproportionately large number of mentions. There may also be too few works that receive no mentions from the judiciary, even though there will be a large number of works that receive a very small number of mentions dispersed over several decades.

[13] Alternatively, in some comparisons we can compare works mentioned early versus those mentioned late, if the timing is plausibly exogenous. See infra Section 4.4 for further discussion.

[14] Our approach follows the partial identification literature in econometrics, especially that of Manski. See, e.g., C. Manski (1997), C. F. Manski (2007), C. Manski and Pepper (2018).

[15] An advantage of bounding is in applications such as ours where there is uncertainty as to whether the standard assumptions required to make causal claims are satisfied. For example, the parallel trends assumption in differences-in-differences analysis, which allows the causal effect to be the post-treatment difference less the pre-treatment difference immediately before treatment. Another source of assumption uncertainty in our setting arises from the dynamic nature of the data: for example, the uncertainty arising from the process that generates academic citations to texts overtime, or uncertainty arising from unknown treatment dynamics.

Another consideration is the sequential nature of treatments in our study. The judicial mentions treatment studied here is an entire sequence of events. The impact of this sequence or treatment history is not necessarily immediate within the year. A treatment that is a sequence of events rather than one event can cascade over time too. A bounding approach is more appropriate here, because it makes better use of the pre-treatment behavior of the time series.

The approach is flexible compared to the regression implementation for the textbook differences-in-differences. Yet, limitations exist even within this more flexible framework because the more varied the treatment paths are the more numerous their possible counterfactuals become. Statisticians call this problem the curse of dimensionality. See discussion below.

counterfactual assumptions change?[16]

Using the bounded deviations approach, we find that judicial mentions increase academic citations to realist works. These effects generally appear with some delay and are relatively robust to the identification tuning parameter we use. We show, for example, that realist works first mentioned in judicial opinions from 1976–2025 received a positive boost in academic citations. Although the magnitude of the effect is not identified, it is between 5–25 added academic citations across model specifications.

The impact of judicial mentions on originalism is less certain. Because it is a nascent judicial theory, it is too early to tell whether academic citations to originalist works increase or decrease in response to judicial mentions. The lag between issued judicial opinions and the development or notice of academic legal theories suggest the data will not pin down the effect for some time. Evidence of causality (or lack thereof) may, in other words, be premature because an uncertain amount of time may need to elapse for an academic work to trickle into the judicial and academic milieu.[17]

Our work contributes to several strands of literature. The most established of these is that of the origin of jurisprudential movements. We bring citation metrics to bear on the evolution of such movements. The beginnings of jurisprudential movements, in particular those that have partisan bearings, are usually prone to revision[18] and

[16]Since the bounds are derived from multiples of the pre-treatment difference, changing the multiple can determine how many multiples of the pre-treatment difference is required before the treatment effect's direction becomes inconclusive. Statistically minded readers will wonder if confidence intervals on point estimates perform the same function as the bounds we suggest here. The answer is no, because confidence intervals are produced by assuming that the observed data is a sample from a larger population and is chosen to reflect this sampling error. We do not suppose that the chosen texts come from a larger sample, even though our choices are not exhaustive. The main source of uncertainty here is that of what assumption best reflects the pre-treatment behavior of the two series.

See generally C. Manski (1993) at 4:

"It is useful to separate the inferential problem into statistical and identification components. Studies of identification seek to characterize the conclusions that could be drawn if one could use the sampling process to obtain an unlimited number of observations. Studies of statistical inference seek to characterize the generally weaker conclusions that can be drawn from a finite number of observations. Statistical and identification problems limit in distinct ways the conclusions that may be drawn in empirical research. Statistical problems may be severe in small samples but diminish in importance as the sampling process generates more observations. Identification problems cannot be solved by gathering more of the same kind of data."

[17]Because little of the mechanism by which judges learn of legal scholarship is known (except for their own pronouncements), it is difficult to know when it is too early to assess the causal impact of judicial (non-)mention on the fate of the publication.

[18]See Tamanaha (2009) for realism and Sawyer III (2018) for originalism.

reinvention.[19] The process of identifying influential texts and then tracing their influence on the academy and the judiciary is less prone to such temptations. Although legal movements are often analyzed from an historical viewpoint, we hope to bring quantitative methods to legal origin stories by considering the judiciary's role in the rise and fall of legal movements.

Another strand of literature, beginning with Judge Harry Edwards' observation that "law and" movements are displacing practical legal training in law schools,[20] seeks to quantify the relationship between the legal academy and the judiciary with a view to studying the divergence between the two institutions. The vast majority of such empirical studies quantify the overall volume of judicial mentions of legal scholarship. While volume is an important metric and adequate for an exploration of whether the judiciary uses legal scholarly work, it does not give a complete picture of the relationship: by focusing on particular substantive movements, we are better able to measure the effect of judicial mentions on subsequent scholarly work relevant to that movement.

Lastly, we use this empirical setting to demonstrate a transparent causal method in dynamic settings. Bounding is robust to a wide range of assumptions researchers make to pin-point a specific causal quantity, many of which go unstated. By proposing a wider range, we reduce the need for numerous or implausible assumptions. The assumption that the post-treatment difference between two counterfactual series is no more than some multiple of the pre-treatment difference is concise, transparent, and nests other popular approaches.[21]

The remainder of the paper is structured as follows. In Section 4.2, we review existing empirical and non-empirical work on how courts use academic legal theories and works. We also discuss how existing scholarship largely neglects how the courts affect legal academics. In Section 4.3, we give an overview of two important U.S. legal movements: American Legal Realism and Originalism. We discuss the observational data used in our study and some high-level descriptive findings from the data. In Section 4.4, we discuss the causal inference problem in this setting.

[19] Both legal realism and originalism has "new"-er versions.

[20] Edwards (1992)

[21] For example, our approach nests the parallel trends assumption for the difference-in-differences ("DID") strategy. See, e.g., C. Manski and Pepper (2018). In the DID approach, comparisons are made between citation trajectories re-centered at their baselines. This can inoculate against a variety of challenges to causal claim-making such as: popularity of the academic work (or the author) before judicial mentions, citation trends that are affected by factors extraneous to judicial mentions, or citation trends related to the time elapsed since publication.

We describe simplifying assumptions and illustrate how our causal strategy works. Finally, in Section 4.5, we apply our methodology to estimate the impact of judicial mentions on academic citations to realist and originalist works.

4.2 Legal Theories in the Academy and in the Courts

The relationship between the legal academy and the courts is a subject of interest for many legal scholars. Still, there is much descriptive empirical work that remains undone. Although there has been much hand-wringing about how academics affect (or not) judges and the judicial process, the reverse channel is relatively underexplored. In particular, credible empirical evidence of the causal effects of judicial activities on the development, production, and reception of academic legal theories is not available. This section discusses prior empirical and qualitative works on how the courts and legal academy interact with one another.

How Courts Use Academic Legal Articles

One well-known lament is that legal academics are largely irrelevant to judges and judging.[22] There is a smattering of quantitative, qualitative, and anecdotal evidence for this claim (Edwards, 1992; Maggs, 1929; Petherbridge & Schwartz, 2012; Schwartz & Petherbridge, 2010). Consider, for example, Judge Posner's claims along these lines:

> "So there is much to criticize in the judicial profession and therefore much room for improvement. But where is the improvement to come from? A possibility that appeals to me as a former law professor is the law schools. Law professors write a great deal about the judiciary—and mainly the federal judiciary. But there is a question about how well informed about, or helpful to, the judiciary that writing is. At present, not very, I have discovered. Not that I'm inclined to apply Brendan Behan's comment about film critics to law professors by comparing the professors (relative to judges) to 'eunuchs in a harem; they know how it's done, they've seen it done every day, but they are unable to do it themselves.' Some judges might think the comparison apt, however."[23]

Empirical findings support this hypothesis. At the Supreme Court, for example,

[22] But see Shapiro (1985, 2021) for most cited scholarship and legal scholars, measured by HeinOnline citations.

[23] Posner (2016).

Newton (2012) finds that only 20.19% of signed opinions mention a law review article from 2001-2011.[24] This may, in part, stem from ideological differences between the justices and the authors of law review pieces: the three justices with the highest percentage of opinions mentioning law reviews were liberal justices (Breyer 26.14%, Stevens 24.54%, and Ginsburg 23.08%), whereas the three justices with the fewest citations to law reviews were all conservative justices (Thomas 13.31%, Rehnquist 13.79%, and Roberts 16.04%).[25]

This trend may be changing. Although the number of mentions to law review pieces is generally low,[26] there is some evidence that law review pieces are increasingly mentioned, especially elite law reviews in high-profile Supreme Court cases.[27] Indeed, Detweiler (2020) finds that, although mentions of law review articles remained low in 2018 (1.8%), opinions increasingly mention law reviews relative to the historical minimum in 2009 (1.47%).[28] Note that the Detweiler (2020) study uses the entire universe of reported opinions from U.S. state and federal courts 1945–2018.[29]

Another important factor affecting judicial mentions of academic (and non-academic) sources is technological progress. Fronk (2010) analyzes federal appellate citation patterns and finds that the average number of cited cases increased from 15.66 in 1957 to 31.14 in 2007.[30] This was a period of technological progress, which substantially lowered the cost and difficulty of locating relevant legal sources. A similar effect may be relevant for citations to law review works, which are increasingly easy-to-access and widely shared, if law review pieces are *complements* to primary legal sources.[31]

A final consideration is that legal opinions may cite law review articles because they speak to a judge's own theories. Legal theories may not constrain judicial decisions

[24] See Table 1 in Newton (2012).

[25] See Table 1 in Newton (2012).

[26] Sirico Jr. and Drew (1991)

[27] See Feldman (2018).

[28] Detweiler (2020) attributes the overall decline from the peak in early 1980s to substitution with primary legal sources.

[29] What is driving this possible shift remains unclear. One possibility is that law review articles are becoming more sophisticated and empirical, which can "bolster a doctrinal claim." See Diamond and Mueller (2010) at 595. Indeed, the amount of "minimal empirical content" in law review articles increased from 1998-2008. See Figure 3 in Diamond and Mueller (2010).

[30] See Fronk (2010) at Table 1.

[31] But see Detweiler (2020) for discussion of law review works and primary legal sources being *substitutes*. In this case, the technological progress for easily accessing primary legal materials would decrease citations to law reviews. Indeed, Detweiler (2020) argues that: "By facilitating the discovery of primary law, these services provided researchers with a viable, and in many ways superior alternative to using law reviews as a shortcut for their own research."

writ-large, but they may persuade an individual judge to the extent they align with the judge's own theories.[32] Indeed, judicial mentions of more jurisprudential or theoretical material may represent alignment with or acceptance of theoretical precepts, means, and ends. To the extent law review works better reflect the theories of practicing judges, we may expect more citations to corroborating law review works.[33]

How Courts Affect Legal Academia

Even if judges find law reviews or legal academia "irrelevant," this does not imply that judges are irrelevant to law professors. Indeed, judges may impact the development in many ways, intentional or otherwise.

The most obvious influence of the courts on legal academia is on legal educational materials. Because of the casebook method, judges impact legal education through written opinions. Empirical descriptive work suggests that certain superstar judges (e.g., Judges Richard Posner and Frank Easterbrook) dominate the casebooks.[34] Longer service (i.e., seniority) is also positively correlated with casebook inclusion.[35] Other studies suggest that additional factors may explain casebook inclusion, especially for non-senior judges with a high degree of casebook inclusion.[36]

Another way judges can impact law professors is by changing the demand for certain legal theories. Simple economic incentive arguments imply that academics respond to the non-pecuniary benefits—e.g., prestige, fame, citations, interest—from mentions in court opinions.[37] Indeed, increased court demand (support) for particular legal theories should, on balance, increase the supply of those theories unless law professors and judges operate in fundamentally different "markets." Given the influx of scholars from other disciplines into the legal academy, academics may feel more pressure to impress like-minded colleagues rather than communicate

[32] See, e.g., D'Amato (1999); see also D'Amato (1989).

[33] If this explanation is correct, should we expect further increases in citations to law reviews as machine learning and generative language models perform increasingly well at legal reasoning tasks and law professors (and lawyers) are better able to pin down a judge's operant theories? See, e.g., Thalken et al. (2023) (discussing LLMs and legal reasoning tasks).

[34] Gulati and Sanchez (2002).

[35] Gulati and Sanchez (2002).

[36] See, e.g., Fitzpatrick and Varghese (2017).

[37] See, e.g., Vladeck (2026) (arguing legal scholars respond to "both supply and demand" in writing birthright citizenship scholarship); *id.* ("the Supreme Court bears at least some responsibility for the latter [demand]"); *id.* ("But so long as this is the Supreme Court's attitude toward legal scholarship, law professors will write articles that are designed and intended to be *cited* as much as they are designed and intended to be learned from.").

to judges on the practical matters of the law. For some judges and some law professors, this is certainly a possibility, as Judge Posner notes:

> "[T]here is a wild literature that I have avoided mentioning in which law professors in immensely long articles subject legal texts to the hermeneutic techniques of postmodernist literary theory. No judge could get *anything* out of that literature, and this unbridgeable gap is not merely a generational one."[38]

Judge Posner offers an especially vivid potential explanation for this phenomenon:

> "The process is Darwinian. In nature each animal species must find a niche for itself, critically including a food that it can find and eat without encountering destructive competition from another species. In the academy each species of professor must find an academic niche in which he can avoid destructive competition from other professors. ... Their need to communicate with persons outside of their niche, such as judges, like the need of a squirrel to learn to eat dandelions as well as nuts, is minimized."[39]

In other words, there is nothing inherently "wrong" with market specialization: judges and law professors simply occupy different niches.

Politics and ideology are also likely relevant to whether legal scholars respond to courts. Because legal scholarship is often ideological,[40] conservative legal scholars may write, for example, to respond to the current Supreme Court whereas liberal legal scholars will not. In other words, scholars may operate in separate ideological silos, which can affect whether certain law professors respond (or not) to court demand.

The focus of the remainder of this paper is testing whether law professors do, in fact, respond to judicial mentions. We examine whether certain academic law pieces increase in popularity following mentions in published court opinions. Because popularity is measured in academic citations to the original piece, we are measuring whether judicial mentions have a particular role in spurring on academic movements.

[38] Posner (1992) at 1928.
[39] Posner (2016) at 8-9.
[40] See, e.g., Table 3 and Figure 3 in Chilton and Posner (2015).

To avoid both problems of over- and under-inclusion,[41] we focus our attention on two specific jurisprudential movements and obtain as many relevant pieces that are recognized as within the movement, including those of their critics. We turn now to describing the two movements and the data we collected.

4.3 Data & Empirical Descriptions of the Movements

This paper uses data on two important legal movements: American Legal Realism and Originalism. We study these two legal movements because of their historical and contemporary impact on the law. As Posner (2016) notes, realism and originalism are two of the most influential legal philosophies in the United States:[42]

> "Legal realism...was a highly influential judicial philosophy not just in the 1920s and 1930s but in the entire period that began with the publication of Holmes's *The Common Law* in 1881 and ended with the end of the Carter Presidency a century later. Reagan's election, his conservative appointments to the federal courts of appeals and the Supreme Court, and the more or less simultaneous creation of the Federalist Society, began to shift the federal judicial balance back toward formalism, giving rise to 'originalism' and 'textualism' and increased resort to dictionaries ('literalism') as sources of statutory meaning—all backward-looking sources of judicial guidance."[43]

Because of the ideological character of these movements, their endorsement was likely to bring greater rewards.[44] Judges are likely to signal their endorsement by mentioning prominent scholarship of the movement. The mention of these by the courts can have an encouraging effect on legal academics who seek to join the fray.

To assess the impact of courts on these two movements, we collected two main types of data for each movement. The first type of data is data on judicial mentions, that is,

[41] The ideal population here would be the universe of all legal pieces that forward a jurisprudential theory and are therefore relevant to judging. Given this set, one could make comparisons among the pieces mentioned by the judiciary and those that were not. But such a universe of works is difficult to obtain because there are a great many of them. Another option is to randomly sample scholarly works and determine their relevance manually. But this might be even more unsatisfactory than the indiscriminate use of the entire sample if the random sample omits highly relevant or influential scholars.

[42] As Posner (2016) notes, the origins of these philosophies were also political.

[43] Posner (2016) at 86-87.

[44] In the case of originalism, for example, this could entail promotions for individual judges endorsing originalism. Note, however, that Judge Posner is a prominent exception, given his criticism of originalism despite being a Reagan appointee.

how often judges cite specific legal works (articles) for thinkers associated with the movement. The second type of data is information on the popularity of different legal works, as measured by academic citations. The data sources, collection process, and high-level descriptive findings are summarized for each movement below.

Legal Realists

Description of the Movement

American Legal Realism began coalescing as a movement in the 1920s and came to prominence during the Roosevelt Administration, when it began to greatly impact the development of the law in the United States. For legal education and scholarship, the realists' primary contribution was in developing a non-formalist theory of adjudication. The formalism or "mechanical jurisprudence" of the day held,[45] roughly, that judges decide cases by legal rules and reasons.[46] By contrast, the realists held that legal rules and reasons were (in at least some cases) window-dressing for judicial decision-making. Judges actually decide cases based on the facts of the case. As Leiter (2010) puts it: "judges are largely 'fact-responsive' rather than 'rule-responsive' in reaching decisions."[47] Realists suggested, moreover, that the law was, at least in "hard" cases, rationally indeterminate: the applicable legal rules or reasons did not uniquely determine a particular outcome.[48]

Although all realists accepted these central tenets, they differed substantially on what determines how judges respond to facts. Leiter (2010) describes the split as two wings: the Idiosyncrasy Wing and the Sociological Wing.[49] Broadly, the Idiosyncrasy Wing (typified by Jerome Frank) held that the individual personality, ideology, psychology, and other background characteristics of judges determine how they respond to facts. By contrast, the Sociological Wing held that certain general social facts determine judicial outcomes. The difference is largely, therefore, one of degree.[50]

As Leiter (2005) observes, the realists were proponents (in theory, if not in practice) of scientific methods of discovery and empirical testing.[51] Because most realists

[45] See Pound (1908).
[46] Leiter (2005).
[47] Leiter (2010) at 249; cf. at 257.
[48] See also Leiter (2005) (discussing how realists also subscribed to a *causal indeterminism* thesis).
[49] See Leiter (2010) at 257-262.
[50] Social scientists today might say the split was in terms of the degree of heterogeneity in the social facts that affect judicial decision-making.
[51] See Leiter (2005) at 50-51.

thought certain general social facts determined judicial outcomes, many were interested in discovering which factual patterns result in certain judicial responses. Some of the realist works then are, unsurprisingly, descriptive empirical projects to specify how particular fact patterns result in certain types of judicial decisions.

Although largely a descriptive theory of adjudication,[52] the realists also impacted the development of "law in action" and the practice of lawyering.[53] For example, realists worked to reform, improve, and summarize the law of contracts. Corbin, for example, wrote his famous treatise on contract law, whereas Karl Llewellyn was instrumental in drafting and promoting the Uniform Commercial Code. Other realists served as judges (Cardozo, Frank, Holmes, Posner) and in executive agencies (Frank, Cohen).

Data Collection

Although many scholars undoubtedly contributed to legal realism, we focus here on the more foundational and influential realist works. Specifically, we examine 16 prominent legal realists and their major works.[54] The 16 writers and their major works are included in Table 4.1.[55] For each work, we then created two annual time series: (1) academic citations, and (2) judicial mentions.

To produce the academic citations series, we use data from the JSTOR Data for Research (DfR) program. This data contains information on the articles that cite the author's last name and the title of the work. The outcome citations are from academic articles only and not, e.g., popular coverage of the work. Because some of the works have very general titles (e.g., "Statutory Interpretation"), the search we carried out is for the author's last name and an exact match to the title of the realist work to avoid false positives.

The second time series is how many times the courts mention realist works, which serves as our "treatment" path data.[56] To collect judicial mentions, we use public

[52] See Leiter (2010).

[53] Realist arguments may have also influenced the development of particular areas of law. For example, Merrill and Smith (2001) suggest that the legal realist introduction of the "bundle of rights" conception of property influenced the development of American property law.

[54] These works were chosen by using the primary sources in survey pieces on American Legal Realism. See, e.g., Leiter (2005) and Leiter (2010). Also, note that although Holmes' *The Common Law*, which was mentioned above in the Posner quote was surely influential, the name of the work is too common to obtain a reliable academic citation series.

[55] We note that many of these works are often considered to be canon. See, e.g., Kennedy and Fisher (2006); see also Mootz III (1994) discussing the legal canon.

[56] See also below Section 4.2.

Table 4.1: Legal Realists and Their Works

Writer	Work
Arthur Corbin	Corbin on Contracts
Benjamin Cardozo	The Nature of the Judicial Process
Benjamin Cardozo	The Paradoxes of Legal Science
Felix Cohen	Transcendental Nonsense and the Functional Approach
Herman Oliphant	A Return to Stare Decisis
Jerome Frank	Are Judges Human?
Jerome Frank	Law and the Modern Mind
John Chipman Gray	The Nature and Sources of the Law
Karl Llewellyn	A Realistic Jurisprudence
Karl Llewellyn	Remarks on the Theory of Appellate Decision...
Karl Llewellyn	Some Realism about Realism
Karl Llewellyn	The Bramble Bush
Karl Llewellyn	The Common Law Tradition: Deciding Appeals
Leon Green	The Duty Problem in Negligence Cases
Leon Green	The Judicial Process in Torts Cases
Max Radin	In Defense of an Unsystematic Science of Law
Max Radin	Law as Logic and Experience
Max Radin	Statutory Interpretation
Max Radin	The Theory of Judicial Decision: Or How Judges Think
Morris Cohen	Property and Sovereignty
Morris Cohen	The Basis of Contract
Oliver Wendell Holmes	The Path of the Law
Richard Posner	The Problematics of Moral and Legal Theory
Robert Hale	Force and the State
Roscoe Pound	Mechanical Jurisprudence
Roscoe Pound	The Scope and Purpose of Sociological Jurisprudence
Underhill Moore	An Institutional Approach...
Underhill Moore	Law and Learning Theory: A Study in Legal Control
Underhill Moore	Legal and Institutional Methods...
Underhill Moore	Rational Basis of Legal Institutions
Wesley Hohfeld	Some Fundamental Legal Conceptions...

data from CourtListener.[57] This data is then used to encode treatment paths for each work, which represent the number of times a particular work is mentioned each year. For example, the treatment path $\vec{k} = (0, 0, 1, 1, 2)$ means a work receives zero court mentions in years 1 and 2, one court mention in years 3 and 4, and two court mentions in the final year. To avoid false positives in judicial mentions, the search we carried out was for the author's last name and an exact match to the title of the realist work.

The cumulative counts aggregated across realist authors is shown in Figure 4.1. The left panel is the cumulative JSTOR citations series over time. This is the cumulative number of works on JSTOR that cite to one of the Table 4.1 realist works. The right panel is the number of court mentions. This is the cumulative number of published judicial opinions that mention one of the realist works.[58]

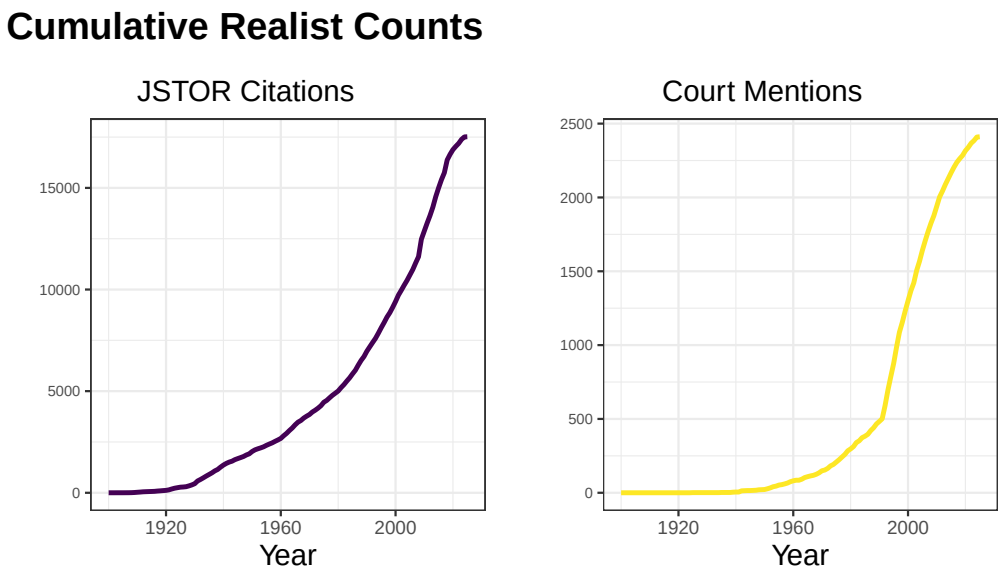

Figure 4.1: Cumulative Realist Counts

Note: The y-axis differs by panel. Academic citations are more numerous than court mentions for every year.

[57] See https://www.courtlistener.com/

[58] In the appendices, we also show plots for individual authors. In Figure D.1, for example, we plot the cumulative citation series for each work separately. In each author panel, the citation series are shown for the works in Table 4.1. We also show the citations on a yearly, non-cumulative basis in Figure D.2. In Figure D.3, we show the yearly court mentions for each author (panels) and work (lines). For example, in 1996 Arthur Corbin's *Corbin on Contracts* was mentioned in 79 court cases. By contrast, Holmes' "The Path of the Law" and Cardozo's "The Nature of the Judicial Process" were each mentioned 8 times by courts in 1996.

The works receive court mentions at different times and levels. In Figure 4.2, we show that some works are highly mentioned early on relative to their publication date (e.g., *Corbin on Contracts*). By contrast, some works only receive judicial notice many years after initial publication (e.g., Pound's "Mechanical Jurisprudence") or are never mentioned by the available CourtListener opinions (e.g., the works of Underhill Moore).[59] That so few works share similar treatment paths (i.e., treated at the same time and level) is important to the method we use to identify causal effects. Because our method can be used on individual works,[60], the fact that there is little overlap in treatment paths requires modifications to simple causal inference strategies.

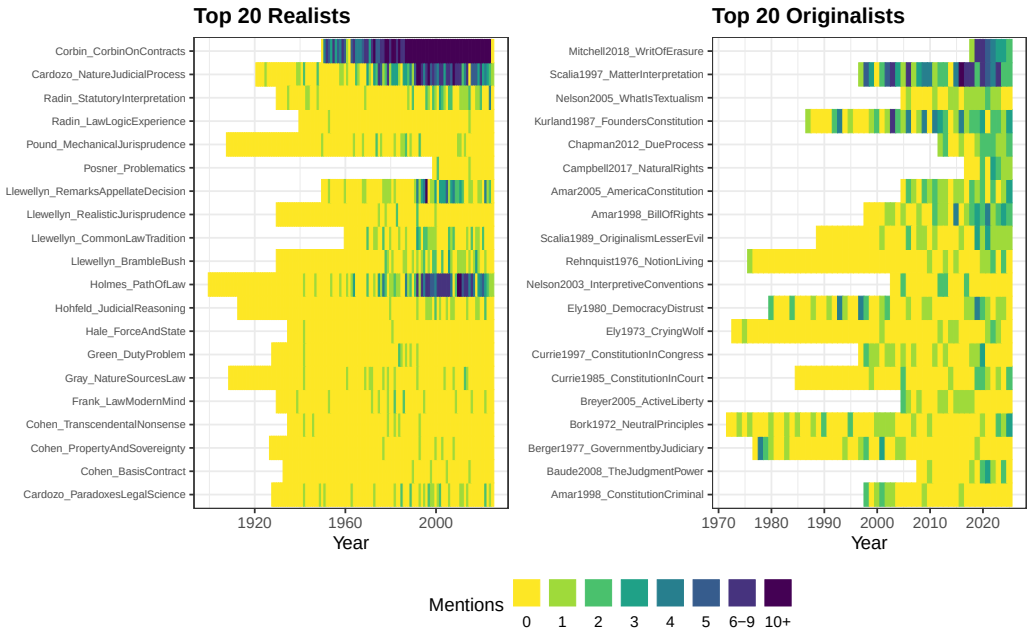

Figure 4.2: Realist & Originalist Court Mentions

Originalists

Description of the Movement

Originalism solidified in the 1980s as a movement and is one of the more influential (and controversial) legal movements of the modern era. The most basic tenet of originalism is that the authoritative meaning of constitutional text is the meaning at

[59] See Appendix D.3 for full court mention histories.

[60] As in, e.g., the synthetic control literature pioneered by Abadie, Diamond, and Hainmueller (2010).

the time of enactment.[61] This proposition, originalism proponents argue, restrains judges, restricts the permissible readings of ambiguous constitutional texts, and is "our law."[62]

Originalism is unlike legal realism in at least two important ways. First, it is concerned exclusively with constitutional law. The realists, by contrast, studied across legal subjects and were particularly interested (as academic and practitioners) in commercial law.[63] Second, whereas realism was largely a descriptive theory of adjudication, originalism advances a normative and prescriptive theory for how judges should decide constitutional law issues.[64] Accordingly, many originalist works advance particular "originalist readings" of a constitutional clause and argue that is how judges should decide cases involving that clause.

Although originalism began as a contested descriptive and normative academic theory, it's influence on the judiciary is now widely appreciated. Originalist theory has impacted how judges justify their decisions in constitutional cases.[65] And the rise in judicial acceptance of originalism has generated more demand for originalist scholarship. This is an almost uniquely American phenomenon: originalism in constitutional theory is not widely embraced outside the United States.[66]

Data Collection

Unlike legal realism, which was largely associated with two law schools (Columbia and Yale),[67] originalism claims many varied adherents across law schools with methodological and substantive differences. Because it is also an active and ongoing legal movement, the foundational originalist authors and texts are not as clear as in legal realism. Accordingly, we have included many more possible foundational originalist works than realist works.

To construct our population of originalist texts, we used a syllabus for an originalist class ("Originalism and Its Discontents") at the Harvard Law School taught by

[61] See, e.g., Whittington (2004).

[62] Baude (2015)

[63] See above for some discussion of realism and contract law.

[64] But note that some originalists also argue for the descriptive claim that originalism is the law. See Baude (2015).

[65] For example in *District of Columbia v. Heller*, both the majority and dissenting opinions devote countless pages to uncovering the 18th century meaning of the Second Amendment.

[66] See Greene (2009a) and Greene (2009b).

[67] Leiter (2005) at 51.

Professor Stephen Sachs.[68] The syllabus contains a long list of optional readings relevant to originalism and, therefore, allowed us to cast a wide net in our search for "foundational" originalist texts. We omitted all non-academic articles (e.g., court cases) that are listed on the syllabus, as well as works for which we could not find academic citation series in JSTOR, or for which the titles of the work were so common (e.g., "Commerce Clause") that they would not reflect citations to the academic article alone.

Our final list contains 191 originalism-relevant texts. As in the realist case, the list of works is neither exhaustive nor definitive. Nevertheless, it includes many popular and influential originalist texts. Importantly, note that not all of the works are by authors that adhere to originalism. Some of the works are merely *relevant* to originalism because this is a course syllabus.

As before, we generated two time series for each originalist work. The first is the number of academic citations to each work (the JSTOR citation series). The second series is how many opinions each originalist work is mentioned in each year. The cumulative time series are plotted in Figure 4.3: the cumulative JSTOR citations series aggregated across authors is displayed in the left panel; the cumulative court mentions series aggregated across authors is displayed in the right panel.

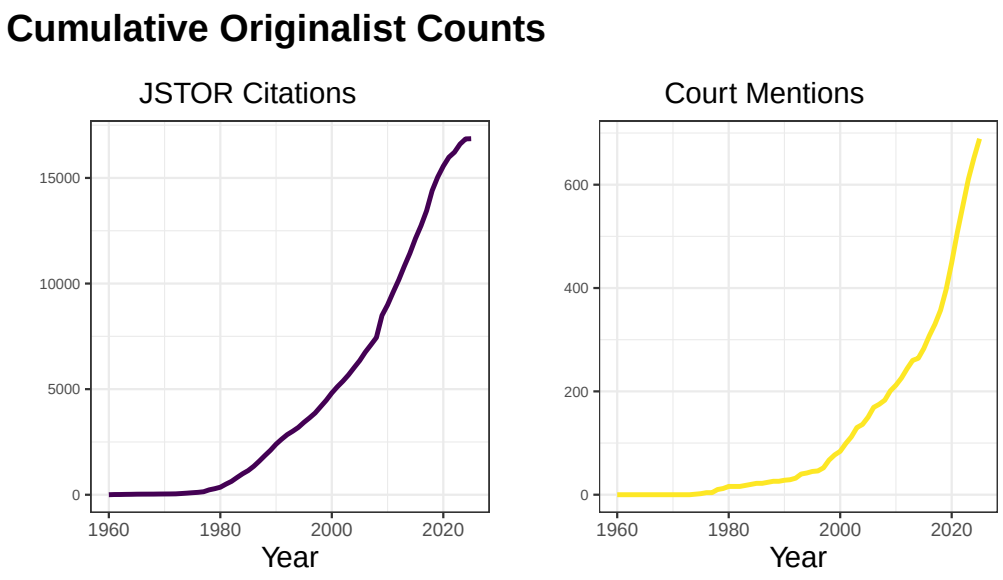

Figure 4.3: Cumulative Originalist Counts

Note: The y-axis differs by panel.

[68] We used the Spring 2025 Sachs syllabus available online at https://www.stevesachs.com/syllabi/originalism.pdf.

As in the realism case, the originalist works are mentioned by judges to varying degrees and at varying times. The right panel of Figure 4.2 also shows the levels and years the different works are mentioned in published judicial opinions for top originalist works.[69] As above, it is important to notice here that the "treated" works experience little overlap in the treatment path space.[70] Unlike the realist case, however, there are many works that are never treated (i.e., never mentioned in a judicial opinion). This is not unexpected because originalism is a newer legal theory than realism.

Descriptive Findings

Textual Data & LLM Classification

This paper is primarily concerned with the academic citation and judicial mention metrics. However, we note here briefly that the JSTOR data in particular contains additional textual data as well that may be of interest for text analysis, machine learning, and LLM classification purposes. For each realist or originalist work, the JSTOR data includes matches to articles or books that cite that work. Each match contains information on the keywords used in the article that cite the realist or originalist work. These keywords summarize that citing article. Consider, for example, that one match citing Jerome Frank's *Law and the Modern Mind* is Rumble (1965). The keywords for the Rumble article are:

> legal realists; legal realism; llewellyn; karl llewellyn; jurisprudence; judicial; american legal; legal norms; ideals; judicial decisions

Each article that cites one of the realist or originalist works has its own keywords.[71]

[69] The timelines for all works are shown in Appendix Figure D.8, Appendix Figure D.9, Appendix Figure D.10, and Appendix Figure D.11.

[70] That is, they do not experience the same number of judicial mentions at the same times. Because of this, we group works into simplified treatment cohorts. See Section 4.2 below.

[71] In Appendix Figure D.5, we show a word cloud of the 100 most common word stems among these keywords for both the realist and originalist corpora. The two word clouds contain 54 words in common. These are mostly words that relate to law and judging generally. The three most mentioned words in articles citing legal realist works that are also mentioned by originalist-citing articles are: law, legal, and judici. The three most mentioned word stems in articles citing originalist works that are also mentioned by realist-citing articles are: constitut, law, and judici. The three most mentioned words in articles citing legal realist works that are *not* mentioned by originalist-citing articles are: contract, realism, realist. The three most mentioned words in articles citing originalist works that are *not* mentioned by realist-citing articles are: origin, mean, and scalia. Unsurprisingly, the originalist-citing articles also include many keywords related to constitutional clauses that are not included in the top 100 realist keywords, such as: fourteenth, equal, protect, religion, and abort.

We use an LLM classifier to categorize citing academic articles into particular legal subjects. Using the JSTOR keywords, author name, and article title, we classify the articles which cite the set of realist and originalist works into legal subject categories. The instructions we gave to ChatGPT 4o-mini are as follows:

> **GPT prompt:** "You are a legal scholar. Given the following article metadata, classify it into exactly one of the following categories: Property, Contracts, Torts, Constitutional Law, Administrative Law, Legislation, Corporate, Tax, Criminal Law, or Jurisprudence. Respond only with the category."
>
> Metadata provided:
>
> - `Title: [title of article]`
> - `Author: [author name]`
> - `Keywords: [keywords]`
>
> **GPT output:** `Category: [category]`

Sample output for two articles from each subject category are included below in Table 4.2. The GPT-output categories seem reasonably accurate based on the titles and keywords. In Table 4.3, we perform a heuristic accuracy assessment to assess the accuracy of the labels. Specifically, we look for whether the GPT-output category stem word appears in the article title or keywords. For example, using regular expressions ("regex") we see whether the word "contract" appears in the article title or keywords columns for each JSTOR article labeled as a Contracts article. In the case of Contracts, 79.18% of articles labeled by GPT as a contracts article include the word "contract" in the title of the article or in the set of keywords.

Table 4.2: Sampled Articles by Legal Category

Category	Title	Author
Admin	Statutory Interpretation In ...	Diver
Admin	Aliens: Deportation ...	Jackson
Con Law	Constructed Constraint ...	Bradley & Siegel
Con Law	State Surveillance ...	Bhatia
Contracts	The Advent Of ...	Teeven
Contracts	The Restatement Of ...	Patterson
Corporate	Winner Of The ...	Attenborough
Corporate	Behind The Cloak ...	Stewart
Criminal	Fear Of Forgiving ...	Love
Criminal	Felony Murder ...	Binder
Juris.	Toward A Black ...	Culp
Juris.	The Problems Of ...	Cohen
Legislation	The Interaction Of ...	Hooker
Legislation	Statutory Language And ...	Miller
Property	Whose Animals? ...	Naughton-Treves, Batterbury, & Bebbington
Property	Effects Of Ownership ...	Ross
Tax	Recent Supreme Court ...	Paul
Tax	The Role Of ...	Polito
Torts	The Int'l Origins Of Torts ...	White
Torts	Autonomes Fahren ...	Schmidt & Klatt

Table 4.3: Articles with Category in Title or Keywords (Regex Check)

Category	Percentage Matched
Administrative Law	49.23
Constitutional Law	66.13
Contracts	79.18
Corporate	63.55
Criminal Law	79.91
Jurisprudence	16.08
Legislation	58.09
Property	67.03
Tax	99.38
Torts	37.70

The classification results into legal subjects across the originalist and realist corpora are shown below in Figure 4.4. The constitutional law and jurisprudence categories are most represented in each corpus. Beyond these top two categories for each corpus, the legal subject classifications of the citing articles aligns relatively with the foci of each of the legal movements. Because of its emphasis on the Supreme Court's constitutional jurisprudence, it makes sense that articles citing originalist works deal more with administrative law and legislation than the traditional common law subjects. By contrast, articles which cite one of the realist articles are more focused on the first-year common law subjects: Contracts, Property, and Torts.

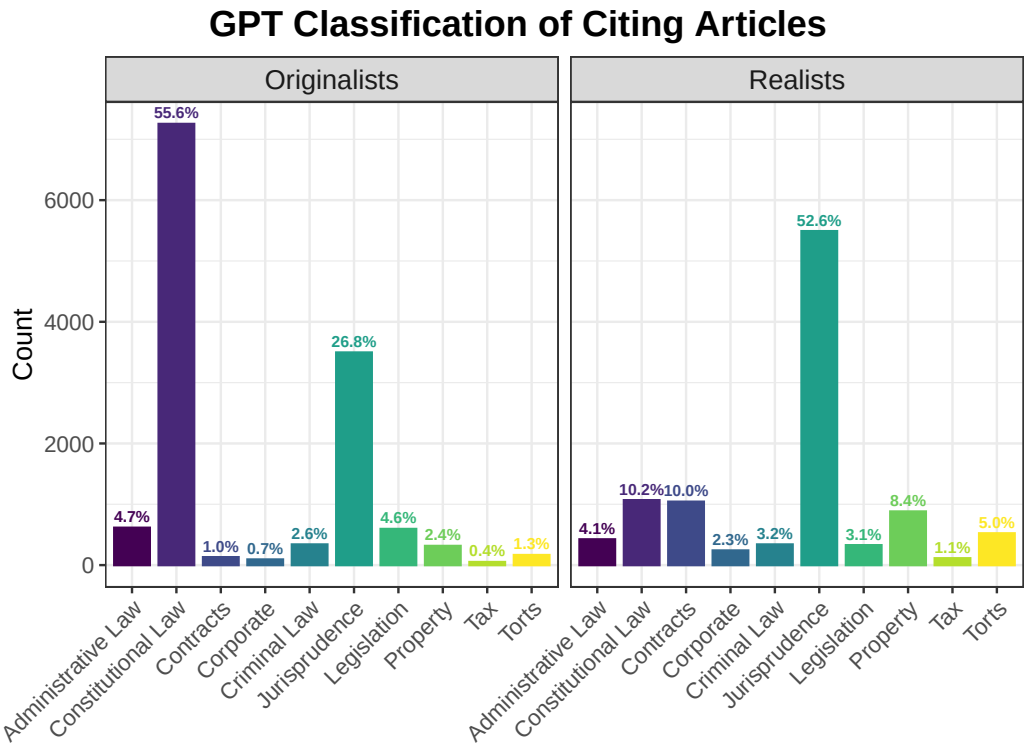

Figure 4.4: ChatGPT Legal Subject Classifications of Citing Articles

Note: The bars show the raw counts of citing articles under each legal subject. The percentage of articles in the corpus (realists or originalists) which fall into the legal subject are also listed at the top of each bar.

Citations & Mentions

The primary data of interest are the two time series for each work. Recall the two time series are: (1) the number of academic citations to an originalist or realist work (JSTOR Citations); and (2) the number of mentions for an originalist or realist work in published judicial opinions (Court Mentions).

In Figure 4.5, we show the basic linear fit for each author's total JSTOR citations on the total court mentions. The totals are plotted on the log scale for visualization purposes. We label the names for authors with four or more works in each dataset. The linear fits of both the realists (yellow line) and originalists (purple line) have positive slopes. That is, the total number of academic JSTOR citations is increasing in the total number of court mentions.

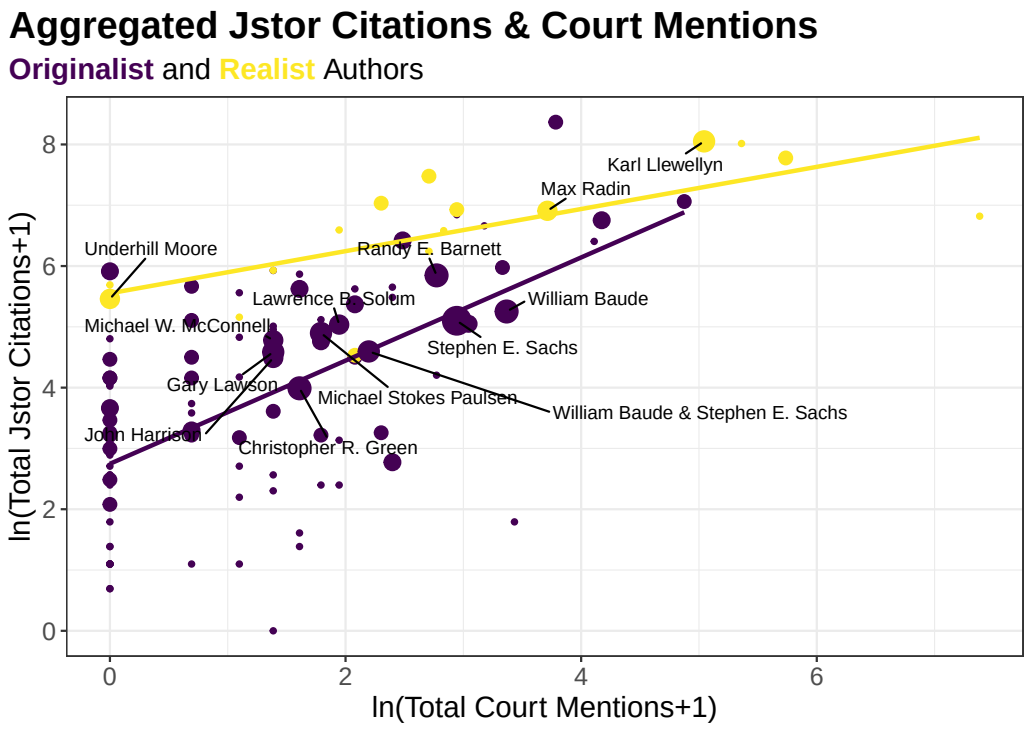

Figure 4.5: Overview of Court Mentions and Academic Citations

Note: Each dot represents a unique author grouping and has volume proportional to the number of total works. Authors with four or more total works are labeled.

Estimates calculated using regressions, an example of which is the line of best fit displayed above, do not necessarily imply a causal relationship. Since the linear fit is for the cumulative series, it does not respect the ordering of court mentions and academic citations, which is necessary for a causal implications.

The regression specification that produced the line of best fit above can be spelled out as follows:

$$\text{Total JSTOR Citations}_{ig} = \alpha_g + \beta_g \times \text{Total Court Mentions}_{ig} + \text{error}_{ig}.$$

The subscripts g relate to the author group (realists or originalists). The independent variable Total Court Mentions$_{ig}$ is the total court mentions for author i in group g.[72] The goal is to obtain an estimate for β_g (the slope in the line of best fit). Such an estimate can be interpreted as the average increase in citations that result from one more court mention for the author group.[73]

[72] Note also that the variables are on the log scale in the specifications.

[73] While the linear fit seems appropriate to the realists, it cannot account for a cluster of originalist

Since this specification obscures important variation within publications in the author group and variation in the court mention path, a different regression that is at a more granular work and year level can be also specified. Now the i subscript refers to the published work, the subscript t to the year, and the subscript g to the group of authors (realists or originalists):

$$\text{Cumulative JSTOR Citations}_{igt} = \alpha_g + \beta_g \times \text{Cumulative Court Mentions}_{igt} + \text{error}_{igt}$$

All citations and mentions are calculated as cumulative series as in Figure 4.1 and Figure 4.3. Now the interpretation of β_g is the average increase in academic citations in a year for an additional court mention in that year (for group g).

Notice, however, that this specification assumes that an effect is instantaneous. Given the possible mechanisms by which judges learn of and mention legal scholarly works in their opinions, such an assumption is overly restrictive. As Figure 4.2 illustrates, scholarly works are rarely mentioned immediately after publication. Therefore, such an instantaneity assumption can induce a downward bias on the estimate β_g for each group of authors g.[74]

Regression evidence, moreover, may be particularly problematic for causal inference where, as here, there is sparsity in the treatments. That is, we only observe limited court mention paths out of a large space of conceivable such paths. Estimates such as β that attempt to encode this process then have limited information to "chew on." To make matters worse, naive linear regression makes comparisons between units on different causal paths and does not necessarily have a causal interpretation.

The previous discussion implies that in this dynamic setting there may be several causal quantities of interest, some requiring stronger assumptions to recover than others. We now turn to detailing the challenges to causality in this setting and explaining the need for bounding.

works that receive no court mentions but have academic citations that exceed author groups that received many court mentions.

[74]There are alternatives to account for the violation of instantaneity. But they require either culling the sample or making strong assumptions about the judicial process. For example, we may consider works only after their first judicial mention. Or we may assume that works take some number of years to be recognized by the judiciary after publication, and truncate both series to account for this lag. Both are imperfect solutions and likely to overestimate the effect of judicial acknowledgments because the information removed occurs where judicial mentions are small.

4.4 Methodology: Bounding Causal Effects

The canonical machinery of causal inference is derived under simple treatment paths. The familiar potential outcomes framework is introduced in a setting where each unit's treatment path is only one time period long. There are then only two possible treatment paths: treated or untreated. It is assumed there are a sufficient number of treated and untreated units that can be used to impute the counterfactuals for each potential outcome. Under different assumptions, quantities like the average treatment effect, average treatment effect on the treated, or average treatment effect on the untreated can be estimated.

In dynamic treatment settings, treatments paths are many-periods long. The potential outcomes framework still applies but is more involved. More assumptions must be made, for example, about whether past treatments matter for later treatments, whether units dynamically select treatments (e.g., "dynamic selection on gains"), and whether units anticipate treatment. Yet, even in these dynamic settings, some treatment paths are simpler than others. If all treatment occurs at one period, the paths can essentially be reduced to the simpler setting by making simplifying assumptions.[75]

Because the number of potential outcomes grows for each individual as treatment paths increase in complexity, we now turn to certain simplifying assumptions on how full treatment history affects potential outcomes, which reduces but does not eliminate the complexity. We then propose our Bounded Deviations strategy and show how these assumptions can point-identify or bound causal effects of interest in our setting.[76]

This section first sets up the counterfactual framework for this setting and describe how it creates a dimensionality problem. Next, we briefly discuss how some simplifying assumptions can be difficult to motivate, hard to interpret, and result in biased causal estimates. Finally, we explain our counterfactual methodology, illustrating it with a simulated example and visualizations.

[75] A well-studied dynamic setting is the staggered adoption setting where the treatment may occur at different times of the path, but any unit if treated at all is only treated once. In such a setting, treated units are usually assumed to retain memory of treatment forever. See, e.g., Callaway and Sant'Anna (2021). See also discussion in Appendix D.2.

[76] Our approach is a Manski-type partial identification assumption. See, e.g., C. Manski and Pepper (2018). We compare our approach to the canonical difference-in-differences (DID) approach in Appendix D.4.

Counterfactual Framework

We now introduce notation to describe the potential outcomes as a function of the treatment histories or treatment paths for each piece i in each time period t.[77] Associated with each i are treatment histories (i.e., judicial mention paths) $\vec{K}_i = (k_{i1}, \ldots, k_{iT})$, where k_{it} is the observed treatment level for piece i at time period t.[78] So, for instance, the treatment history $(0, 1, 1)$ means a piece was mentioned zero times in the first period and once in the second and third periods. We imagine that any piece could have been subjected to any alternate judicial mention path and have shown a different trajectory of academic citations, which we refer to as a potential outcome.

We write the potential outcomes (i.e., potential academic citations) for each piece i in time period t under treatment path $\vec{k} = (k_1, \ldots, k_T)$ as:

$$Y_{it}(\vec{k}) = Y_{it}(k_1, \ldots, k_T)$$

This is the number of academic citations piece i would receive in time t if it experienced judicial mention history $(k_1, k_2, \ldots, k_T)$.[79] Finally, we define the potential outcome matrix $\boldsymbol{Y_i}$ for piece i. This matrix collects the many potential outcomes under counterfactual treatment histories.[80] For example, consider a two-time period model with a binary treatment possible in each period and an individual i that didn't receive treatment at time $t = 1$ but did receive treatment at time $t = 2$.[81] Then we write i's potential outcome matrix $\boldsymbol{Y_i}$ as follows:

$$\boldsymbol{Y_i} = \begin{bmatrix} Y_{i,t=1}(0,0) & Y_{i,t=2}(0,0) \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ Y_{i,t=1}(1,0) & Y_{i,t=2}(1,0) \\ Y_{i,t=1}(1,1) & Y_{i,t=2}(1,1) \end{bmatrix} = \begin{bmatrix} ? & ? \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ ? & ? \\ ? & ? \end{bmatrix}$$

The unobservable potential outcomes for i are denoted by the question marks. They define the **identification problem**: because unit i has does not receive treatment at

[77] All time periods can be collected and ordered in a set $\mathcal{T} = \{1, \ldots, T\}$

[78] Note that, in principle, each k_{it} can be multidimensional if there are multiple treatments, but we focus on scalar k_{it} within a discrete set $\mathcal{K} = \{0, 1, \ldots, K\}$.

[79] Implicit in this notation is that there are no treatment "spillover" effects, i.e., the potential outcome for piece i at time t depends only on i's treatment path.

[80] For a fixed ordering of the potential treatment paths $\mathcal{K} = \{\vec{k}_1, \vec{k}_2, \ldots\}$, the (i, j)-th entry of $\boldsymbol{Y_i}$ is given by $Y_{i,t=j}(\vec{k}_i)$.

[81] In other words, their observed treatment path is $\vec{K}_i = (k_{i1} = 0, k_{i2} = 1)$.

$t = 1$ and does receive treatment at $t = 2$,[82] the outcomes for i in period t at different treatment histories are unobservable.[83]

Curse of Dimensionality & Treatment Path Cohorts

The number of these unobservable potential outcomes grows as the number of time periods and number of treatments grow.[84] This creates what is sometimes called "a curse of dimensionality" problem: allowing for i's potential outcome at time t to depend on the full treatment path increases the complexity of the identification problem.[85] In terms of our potential outcome matrix, the problem is that the number of question marks inside the matrix is growing too quickly.

To take control of the explosion of paths, we define four treatment cohorts, which simplifies the treatment path space. The first cohort is the never-treated cohort: these units never receive any judicial mentions. The remaining three cohorts are those that receive at least one judicial mention. They are split according to whether they receive a court mention Early, Mid, or Late in the movement. The years were determined to ensure a roughly equal split among the three treated cohorts. This resulted in the following cohort splits for each legal movement:

Realists

Cohort	Number Works
Early (1925–1945)	9
Mid (1946–1975)	7
Late (1976–2025)	7
Never-Treated	8

Originalists

Cohort	Number Works
Early (1974-2009)	31
Mid (2009-2019)	28
Late (2020-2025)	32
Never-Treated	100

Table 4.4: Cohort Distribution for Realists and Originalists

In each table the years in parentheses indicate the times that the cohort first received a court mention. For example, if a realist work was first mentioned in a published court opinion between 1925–1945, then it is in the Early treated cohort. The cohorts are roughly equal in terms of size, except that for the originalists there are many more Never-Treated units.

[82] That is, i has treatment history $\vec{K}_i = (0, 1)$.

[83] The unobservability of some counterfactual outcomes is often called the "fundamental problem of causal inference." See Holland (1986).

[84] As the cardinality of $\mathcal{K}$ and $\mathcal{T}$ increase, that is.

[85] To address the dimensionality problem, researchers typically restrict the dependence of the potential outcomes $Y_{it}(\cdot)$ on the treatment history. We discuss further in Appendix D.2 how different assumptions that are (implicitly or explicitly) used in practice can reduce the dimensionality of the identification problem.

Bounded Deviations & Partial Identification

We now turn to explaining the framework we use to bound treatment effects in this setting. As in C. Manski and Pepper (2018), we use assumptions that bound causal effects for a pre-determined level of tolerance, technically known as a tuning parameter.[86] Our approach, which we call **Bounded Deviations**, uses information from units before their treatment paths diverge (i.e., *pre-divergence*) to impute counterfactual paths when the treatment paths no longer agree (i.e., *post-divergence*). Intuitively, the method imposes that the average treated group counterfactual must be "close" enough to the average control group observed outcome. There are many choices to encode this closeness. But we use the maximum observed pre-treatment deviation, which, to be more specific, is the largest difference in annual academic citation count between two series since they were both in publication.[87] Our method requires the treated group counterfactual to be within the maximum pre- divergence difference between observed treated and control average outcomes.[88] We use this method to estimate the average treatment effect on the treated ("ATT"), which is the average amount treatment changed outcomes for those actually treated.[89]

For demonstration and visualization purposes, we first illustrate the Bounded Deviations approach with mock simulated data. In Figure 4.6, we consider two average group outcomes for units on different treatment paths $\vec{k}_1$ and $\vec{k}_2$. For simplicity, we call the $\vec{k}_1$ path "Treated" and the $\vec{k}_2$ path "Untreated." These treatment paths are the same until 1950, after which they diverge. The observed average for Treated units is plotted as the blue line from 1900 to 2000.[90] We also plot the observed average for Untreated units in red only before 1950 for visual clarity.[91] The maximum pre-divergence deviation between the average for the two groups is 1.69 and occurs shortly after 1925.[92]

[86] Our approach is a variant of the "Bounded Variation" DID-style methodology in C. Manski and Pepper (2018), adapted to a setting with long treatment histories.

[87] This choice was motivated by its ease of interpretation, but other choices include the average deviation before treatment divergence.

[88] See Appendix D.4 for full details.

[89] For familiar readers, the Bounded Deviations method we use nests the widespread difference-in-differences method (DID) strategy. For a discussion of how the bounds from our method contain the usual DID estimates if parallel trends does hold, see Appendix D.4. We also illustrate how causal conclusions are sensitive to the tuning parameter assumption in Appendix D.5.

[90] Formally this is the observed conditional expectation function $\mathbb{E}[Y_{it}(\vec{k}_1|K_i = \vec{k}_1]$, which is plotted in blue.

[91] Formally, this is the observed conditional expectation function $\mathbb{E}[Y_{it}(\vec{k}_2|K_i = \vec{k}_2]$, which is shown in red.

[92] Note that what is plotted in Figure 4.6 are the conditional expectation functions (CEFs) for units on a particular path. For example, the solid-red line is the average observed outcome for units

The problem of causal inference here (i.e., the "identification problem") is what counterfactual outcome would units on one path experience if they were instead on another treatment path? In the left panel of Figure 4.6, the blue-striped region is the counterfactual region for the blue Treated units if they had instead been Untreated.[93] Recall our identifying assumption: if the two treatment paths had not diverged, average Treated outcomes would be "near" the observed average for Untreated units.[94] How close would they be? They'd be within the maximum pre-divergence deviation.[95] In Figure 4.6, the maximum pre-divergence deviation is 1.69 and the blue-striped region is the counterfactual region under the assumption that blue units would've continued to be within 1.69 of the observed red unit outcomes.

---

on path $\vec{k}_2$ in each time period.

[93]Recall that Treated units are those that actually experienced treatment path $\vec{k}_1$. Untreated units are those that instead experienced treatment path $\vec{k}_2$.

[94]Mathematically, we are saying the CEF for units on treatment path $\vec{k}_1$ would be close to the CEF for units with path $\vec{k}_2$.

[95]When the multiple $C = 1$.

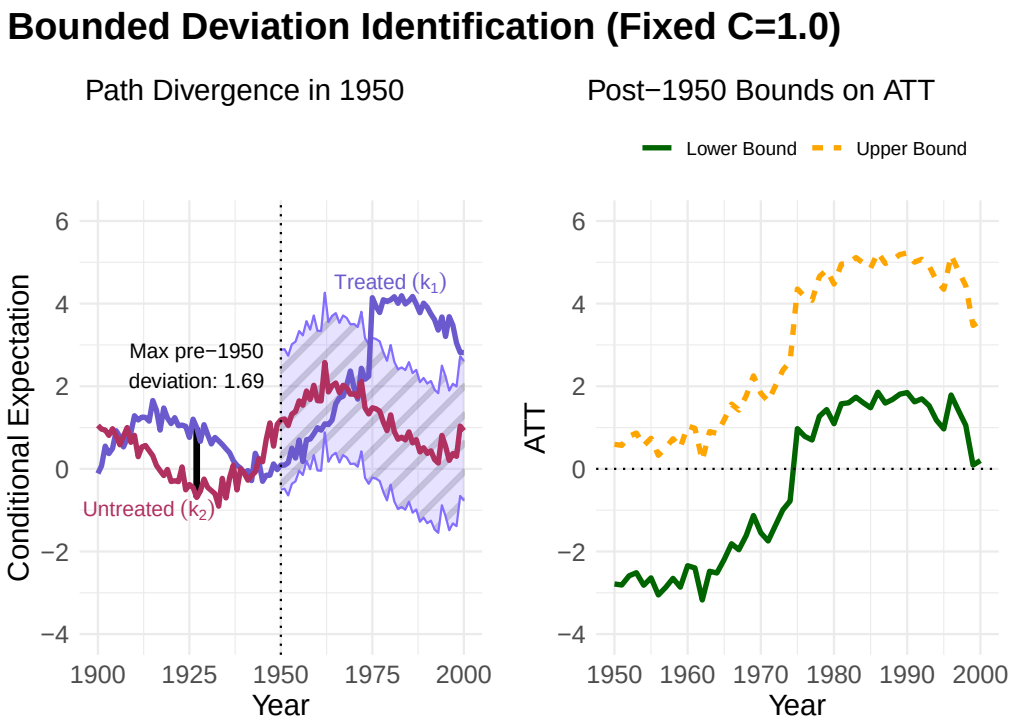

Figure 4.6: Bounded Deviations Example

Note: In the left panel, the two lines are average outcomes for Treated (blue) and Untreated (red). The black-dotted vertical line is the path divergence time period. The blue-striped region is the counterfactual region for blue Treated units: this is where counterfactual outcomes for Treated units would be if they had been Untreated. The maximum pre-divergence deviation is 1.69 and is shown by the black vertical line around 1925. In the right panel, the y-axis is the average treatment effect on the treated, with a black-dotted horizontal line is plotted for zero effect. The upper and lower bounds on the ATT for each post-1950 time are plotted as dashed-yellow and solid-green lines, resp.

This counterfactual region implies bounds on treatment effects. These bounds are shown in the right panel of Figure 4.6. In the panel, the lower bound is plotted as a solid-green line and the upper bound is plotted as the dashed-yellow line. The dotted horizontal line at zero is also shown and corresponds to no effect (i.e., $ATT = 0$).

What do these bounds tell us? Before 1975, our bounds include the possibility of a zero effect, i.e., the dotted line at zero lies between the lower bound (solid-green line) and the upper bound (dashed-yellow line). Notice that this corresponds to the pre-1975 solid-blue line lying entirely within the blue-striped counterfactual region in the left panel. Moreover, we cannot say (identify) what the direction (sign) of the effect is. That is, we cannot know whether Treated units would have experienced higher or lower outcomes relative to their counterfactual outcomes had they been

Untreated.[96]

After 1975, by contrast, the bounds do not contain the zero effect. That is, the dotted line lies below the lower bound on the *ATT* (solid-green line). Therefore, we can identify that the treatment effect of being Treated rather than Untreated increased the average outcomes for Treated units post-1975.[97] In other words, the treatment effect is sign-identified post-1975 because the upper and lower bounds are both positive, even though we cannot say precisely the magnitude of the effect.

The bounds we obtain are sensitive to a counterfactual tuning parameter C. The plots above assume $C = 1$, which intuitively means that the pre-divergence and post-divergence periods differ only in the judicial mentions but otherwise are "sufficiently alike." How alike they are is governed by how large or small C is. Large values of C will increase the size of the counterfactual region and suggest that the past is a relatively weak guide to the future. By contrast, values of $C < 1$ imply that the counterfactual outcome is even closer to the observed outcome than it was pre-treatment divergence. We show how this affects the upper and lower bounds in Appendix D.5.

4.5 Results: How Courts Impact Legal Movements

In this section, we use the methods developed above to estimate bounds on the causal effects of judges on legal realism and originalism. We quantify the impact of judicial mentions on the reception of academic legal ideas. Specifically, we bound the causal effects of judicial citation to certain legal realist or originalist works on the popularity of those works over time. We illustrate how to estimate upper and lower bounds for works if they had experienced alternative (i.e., counterfactual) treatment paths.

Individual Treatment Effects: An Illustration with Two Realist Works

First, we illustrate the methodology with a simple example of two contemporary realist works. Our causal strategy can be easily adapted to estimate individual (unit-level) treatment effects.[98] We consider, for example, comparisons between two

[96] That is, we cannot identify the sign if those Treated units on path $\vec{k}_1$ had instead been on the "Untreated" path $\vec{k}_2$ for years prior to 1975.

[97] Again, nothing here requires a binary treated-untreated setting: this methodology can be applied if some units are on treatment path $\vec{k}_1$ rather than another treatment path $\vec{k}_2$.

[98] As in the synthetic control literature. See, e.g., Abadie, Diamond, and Hainmueller (2010) and Abadie, Diamond, and Hainmueller (2015). In other words, we can relax the simplifying assumptions about the path space to bound unit-level causal effects. Recall as discussed above that few of the legal

early realist works by Holmes and Gray in Figure 4.7. In the left panel, the observed JSTOR citations for Holmes' "The Path of the Law" and Gray's *The Nature and Sources of the Law* are plotted in blue and red, respectively. For many years, neither work receives judicial mention according to the CourtListener data. After 1947, however, the treatment paths for the two works diverge, which is represented by the black dotted line. After that time, the two works receive a different number of court mentions. The levels of the court mentions are shown by the volume of the blue and red dots for each work. As shown in the left panel of Figure 4.7, Holmes' "The Path of the Law" receives many more judicial mentions (larger blue dots).

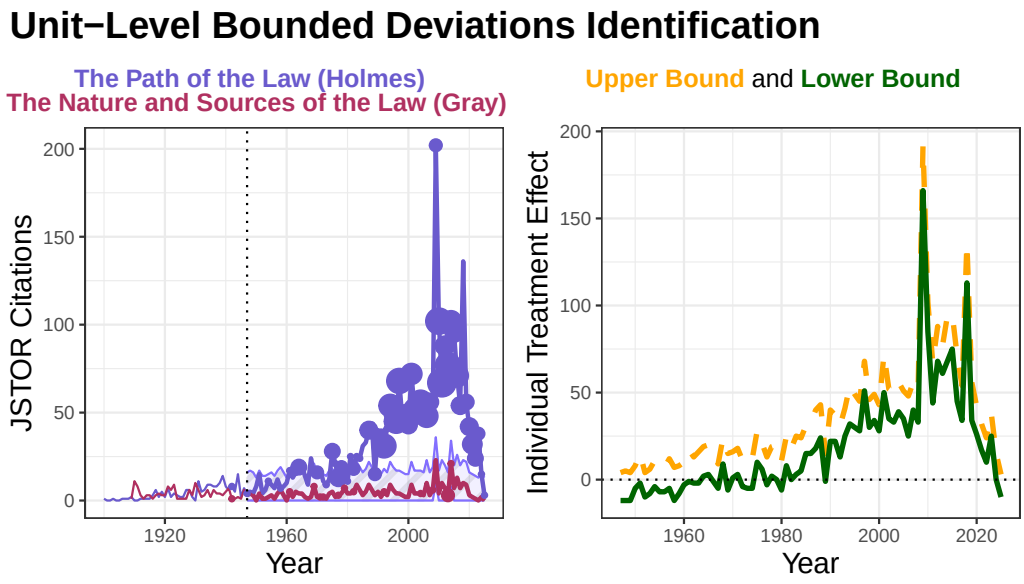

Figure 4.7: Individual (Unit-Level) Treatment Effects Illustration

Note: The left panel shows the observed academic citations for two works. The dots represent years in which the work received court mentions, with the volume of the dots corresponding to more mentions. The right panel shows the individual treatment effect upper and lower bounds for tuning parameter $C = 1$.

Using our Bounded Deviations approach for individual units, we can bound the causal effects from Holmes' work experiencing a different treatment path than

works we study experience the exact same treatment paths (judicial mentions) over the entire time. Instead, the different works are mentioned at different times and at different intensities (levels). For example, some legal works are mentioned early by many judges but then receive few (if any) judicial mentions in the late-20th or early-21st century. By contrast, some works only begin to receive a few judicial mentions many years after initial publication. And a select few works (e.g., *Corbin on Contracts*) have received high-levels of judicial mentions most years since publication. Our approach can allow for the potential popularity for each legal work to differ depending on when and how often judges mention the work in (published) court opinions. See further discussion in Appendix D.2.

Gray's work. The blue-striped region in the left panel is the counterfactual region, as discussed above: it is where we assume the citations for "The Path of the Law" would have been if it had stayed on the same judicial mention path as Gray's work. In the right panel of Figure 4.7, we use this counterfactual region to compute the upper and lower bound on the individual treatment effect. As you can see, the sign of the treatment effect in earlier years includes the zero effect possibility. However, in later years the treatment effect is positive. This is because the citations to Holmes' work are much higher than the top part of the counterfactual region.[99]

Although our method allows for bounding dynamic treatment effects for any pair, we recommend this only for those works that share a long enough pre-divergence period. Usually this means both pieces were in publication for several years before one was mentioned by the judiciary. Individual treatment bounds give us the most granular information, but the number of possible comparisons that could be made can be overwhelming. For example, even with 20 contemporary pieces there are 190 possible comparisons that we could make.[100]

Since it would be difficult to summarize such a large number of comparisons, we coalesce our treatment paths to cohorts, and continue our analysis on these cohort-level series. Cohorts are defined based on when a legal piece received their first judicial mention.

Cohort Averages

Here, we turn to identifying effects for entire cohorts of realist and originalist works. We first display the cohort average JSTOR citations for the realists and originalists. Recall that our cohorts are defined as follows: the Never-Treated cohort never receive any judicial mentions, the Early treated cohort first receives at least one judicial mention 1925–1945, the Mid treated cohort first receives at least one judicial mention 1946–1975, and the Late treated cohort first receives at least one judicial mention 1976–2025. The average number of yearly JSTOR citations for each cohort and corpus (realists and originalists) are shown below in Figure 4.8.

[99] See Appendix D.2 for further discussion.

[100] In between calculating the individual treatment effect for any pair and a cohort averages is the possibility of calculating an individual treatment effect of a select piece against the average of all other never mentioned units, such as in a synthetic control, set-up. But even with this reduced number of comparisons we are left with too many comparisons to make any informative statements about the impact of judicial mentions.

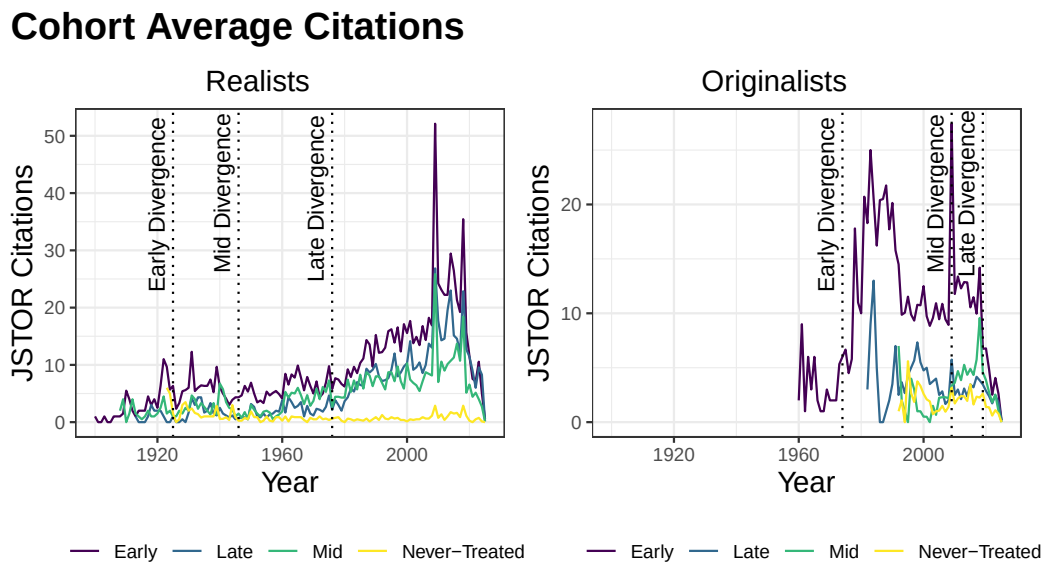

Figure 4.8: Cohort Average Citations and Divergence Periods

Note: The divergence periods that demarcate the treatment cohorts are plotted as vertical black dotted lines.

For both originalists and realists, the never-treated averages are, in most time periods, lower than the three treatment cohort averages. In general, the Early treated average (in purple) is higher than the yearly citations received for other cohorts after Early Divergence. Before Early Divergence, however, observe that the Early, Mid, and Late cohorts receive a similar number of yearly citations on average.

These cohort average citation series are the basic "building blocks" that we use to compute causal effects. The results that follow use pre-divergence deviations between the series and our counterfactual tuning parameter C to compute upper and lower bounds on the treatment effect of a judicial mention.[101]

The Effects of Court Mentions on Legal Realism

In our first application, we estimate the effect of judicial mentions on the reception for the realist thinkers mentioned above. For ease-of-reference, the cohorts for comparison purposes are as follows: the Never-Treated cohort, which received zero court mentions; works in the Early treated cohort, which first received court mention

[101] The Early treated cohort has a very short (case of legal realism) or no (case of originalism) pre-divergence time overlap with other cohorts. This makes it difficult or impossible, resp., to make appropriate comparisons between the early-mentioned and never-mentioned cohorts, so we report only the comparisons between Mid vs Never-Treated and Late vs Never-Treated within the text. See Appendix D.6 for further results.

in 1925–1945; works in the Mid treated cohort, which first received court mention in 1946–1975; and works in the Late treated cohort, which first received court mention in 1976–2025.

Our first set of results estimates the effect from being in the Mid treated cohort versus having never been treated. In Figure 4.9, we plot the upper bound (dashed-yellow) and lower bound (solid-green) on the *ATT* for being in the Mid treated cohort versus the Never-Treated cohort for units actually in the Mid treated cohort.[102] The different panels are for different values of the identification tuning parameter C.[103]

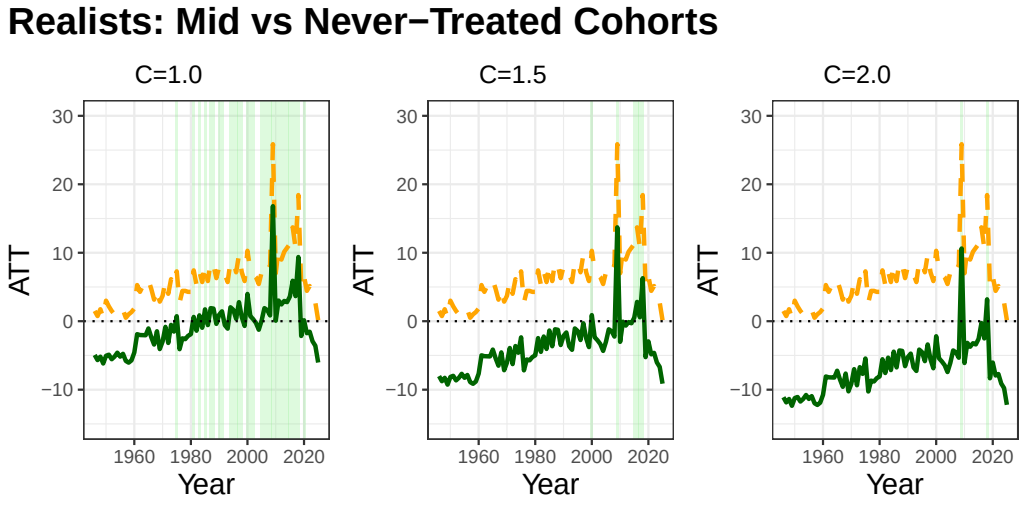

Figure 4.9: Realists: Mid vs Never-Treated Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the *ATT* upper bound is the dashed-yellow line, and the *ATT* lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

The light green-shaded region in each panel are those years for which we identify the sign of the *ATT* to be positive. As shown in the left-most panel of Figure 4.9, many of the 1980-2000 years sign-identify the *ATT* when $C = 1$. This means that the Mid units would have been much "closer" to the Never-Treated units if they had never received judicial mentions. Instead, the judicial mentions increased the

[102] In the appendices, we show results for comparisons involving the Mid treated cohort when Arthur Corbin is omitted. Although Corbin is often associated with realism, *Corbin on Contracts* is, in certain respects, distinct from the other academic realist texts insofar as its focus is largely doctrinal. In the appendix, we show that the results are largely unchanged if Corbin's work is omitted from the Mid treated cohort comparisons. See Appendix D.6.

[103] See Appendix D.5 for further details.

popularity (academic JSTOR citations) for realist works that were first treated in the 1946–1975 years (Mid cohort).

The effect is even more pronounced for realist works in the Late treated cohort. Relative to Never-Treated units, realist works that were first mentioned in judicial opinions from 1976–2025 received a positive boost in academic citations. This effect holds across values of the identification tuning parameter C, as shown below in Figure 4.10. In each panel, there are many years for which we sign-identify a positive ATT for Late treated units relative to the Never-Treated cohort. This is indicated by the large number of years which are shaded light green in each panel of Figure 4.10.

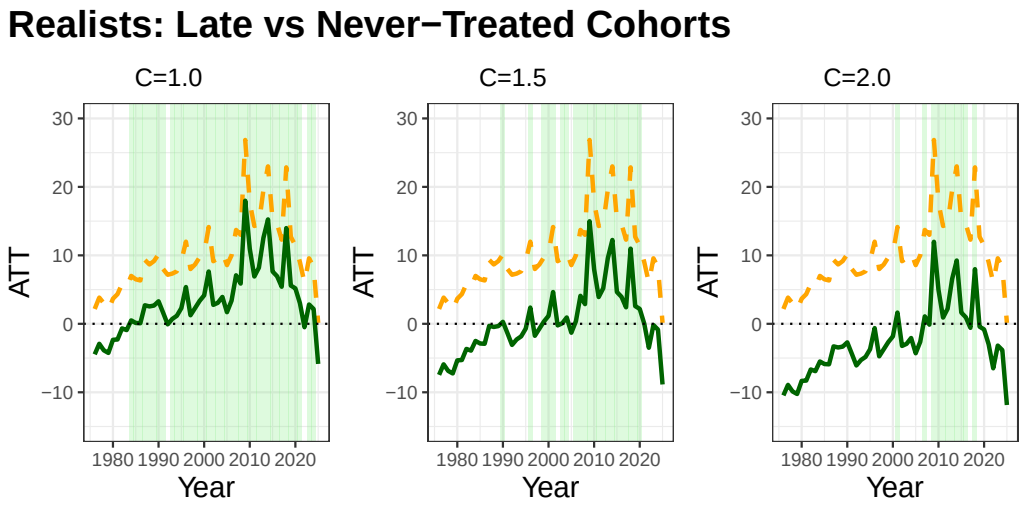

Figure 4.10: Realists: Late vs Never-Treated Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

These findings highlight an important difference between legal scholarship and judicial opinions. Existing empirical studies show that the value of court opinions depreciate rapidly. For example, Black and Spriggs II (2013) find that case precedent is less likely to be cited over time: new cases are about 30% likely to be cited, whereas 20 year old precedent is less than 10% likely to be cited.[104] By contrast, our results

[104] See Figure 1 in Black and Spriggs II (2013), which shows that the predicted probability of citation as a function of the age of precedent decreases rapidly.

suggest some legal scholarship receives more citations later in life.[105] In other words, court opinions age like milk but legal scholarship ages like wine.[106] [107]

Further comparisons between treatment cohorts are explored in Appendix D.6. To summarize the results there, Early treated units received a positive boost from court mentions relative to Mid treated units, although this effect only appears in the 2000-2025 years. It also appears that being treated in the Mid cohort decreased citations relative to being treated Late. And, finally, being treated Early versus Late resulted in a positive effect for a few years in the 21st century, although most years the sign of the effect is not identified.

The Effects of Court Mentions on Originalism

In our second application, we estimate the effect of judicial mentions on the reception for the originalist thinkers mentioned above. For ease-of-reference, the originalist cohorts are: the Never-Treated cohort, which received zero court mentions; works in the Early treated cohort, which first received court mention in 1974–2009; works in the Mid treated cohort, which first received court mention in 2009–2019; and works in the Late treated cohort, which first received court mention in 2020–2025.

Our first originalist comparison is the Mid treated cohort versus the Never-Treated cohort. As shown in Figure 4.11, the upper and lower bounds on the ATT for Mid treated units are shown in dashed-yellow and solid-green, respectively. For originalism, the sign of the ATT is identified for only one year and for only small values of the tuning parameter C. Moreover, the magnitude of possible ATT values is much smaller: the y-axis range is much smaller than in the realist results above. It is difficult to compare these results to the realist case because we have only observed the originalist movement for a much shorter span of time than the realism. Therefore, we caution against drawing a conclusion such as that judicial attention did not play a role in originalism since the effect may yet to be observed in data.

[105] See also Figure 4.8.

[106] Merritt and Putnam (1995) uses Ronald Coase's work on *Social Costs* as a specific example of scholarly writing that came to influence judicial decisions belatedly.

[107] Cf. Traynor (1962) at 10. ("Time is with the law reviews. An age that churns up problems more rapidly than we can solve them needs such fiercely independent problem-solvers preoccupied with long-range solutions.")

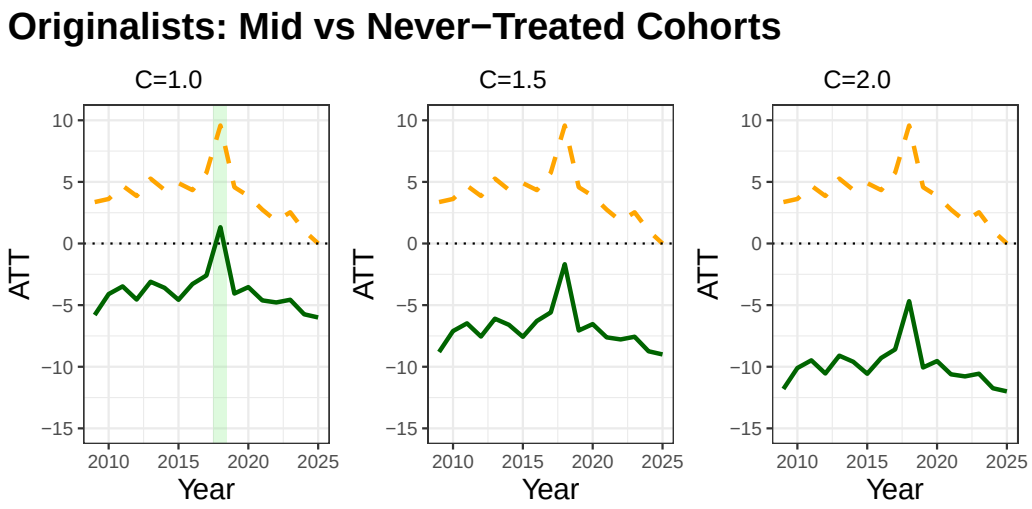

Figure 4.11: Realists: Mid vs Never-Treated Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

The comparison between Late treated originalist works and Never-Treated originalist works is similar. In Figure 4.12, we do not identify the sign of the effect from being in the Late treated cohort rather than the Never-Treated cohort. This means we cannot say whether judicial mention increased or decreased the academic citations for Late treated originalist works relative to never being mentioned in court opinions. We do show, however, that the magnitude of the effect is small either way: for $C = 1.0$ in the left-most panel, for example, the average effect of a Late judicial mention is less than 5 in absolute terms. This means, for example, that the data is consistent with judicial mention increasing (or decreasing) the academic citations to Late treated originalist works by some amount less than 5.

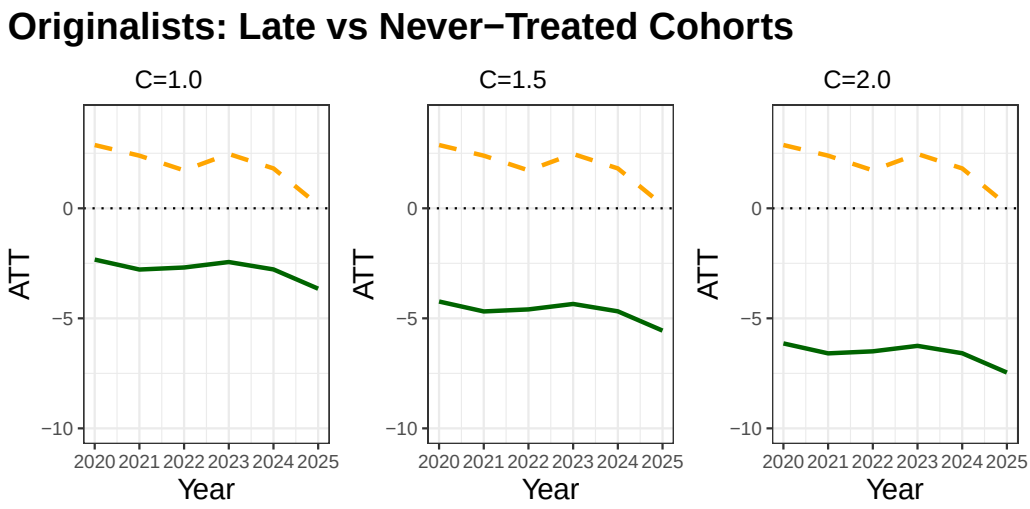

Figure 4.12: Originalists: Late vs Never-Treated Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

Again, we emphasize that we are working with data for only a small number of years.[108] Given the realist results, it appears that the impact of judicial mentions on academic citations takes many years (sometimes many decades) to noticeably affect academic citations. This makes sense given the time it takes for academics to appreciate particular works, process judicial changes to the law, and determine the canonical or important works in a particular legal tradition. In this case especially, originalist changes to constitutional law are active and ongoing. The effects of judicial originalism on legal academics may therefore take more time to appear in the data.[109]

4.6 Conclusion

Existing empirical scholarship has primarily focused on how judges *use* academic legal works but has not estimated the *effect* of judges on legal academia. The primary intent of our work is to bring tools in the causal inference revolution to legal questions. Specifically, we show how a bounding approach can estimate the impact of judicial mention on citations to realist and originalist works. Estimating counterfactuals in settings where potential outcomes may depend on the entire history or

[108] Particularly for the Late vs Never-Treated cohorts.

[109] As in the realist case, see Appendix D.6 for an additional result plot. There we compare Mid and Late treated originalist works and similarly find no sign-identification for C-values of 1.0, 1.5, 2.0.

path of treatment requires several simplifying assumptions. Bounding approaches can accommodate these simplifications by being more cautious about the causal range. As discussed above, the general strategy is to assume how unobserved counterfactual quantities evolve to determine the causal effects of treatment at different times (i.e., on cohorts). This approach can be applied to further legal settings with complicated dynamics.[110]

This paper is relevant to the alleged divergence between the legal academy and the judiciary. In contrast to empirical analyses that test for a divergence by examining the entire universe of legal scholarship, we narrow our focus to two particular jurisprudential movements. Despite this restriction, our data shows large variation in when (and to what degree) judges mention scholarly work. Some works are mentioned early and often, whereas others are mentioned (if at all) long after publication. It is this variation that we exploit for causal identification to test whether judicial mentions matter for academics scholarship or not.

In our application, we find that courts impact the popularity of legal theories. Although it is too soon to say with certainty how judges affect originalist citations, we find positive effects of judicial mentions on citations to legal realist works relative to works that are never mentioned by judges. In the case of legal realists, we also find support for the conjecture that "time is one the side of law reviews,"[111] in that delayed recognition by the judiciary does not imply that legal scholarship and judging are mutually irrelevant.

Because of our focus on the two movements, we cannot extrapolate our findings to all areas of the law. A challenge with doing the same analysis at the level of the entire scholarly corpus is to keep our comparison sets similar in their content. For we cannot compare scholarship about matters of constitutional interpretation to criminal procedure. Among other considerations, their importance to judging differs in the volume of such cases heard and the marginal impact of work across fields is non-uniform. Future empirical work must attempt to overcome these challenges to better understand the relationship between the courts and the academy.

[110] As shown above, boundedness assumptions nest the standard parallel trends assumption and are robust when counterfactual trends are not parallel. This is important because much empirical legal work relies on the difference-in-difference strategy to infer causal effects. Violations of the assumptions underpinning that strategy can result in "heroic" estimates.

[111] Traynor (1962).

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, *105*(490), 493–505. https://doi.org/https://doi.org/10.1198/jasa.2009.ap08746

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, *59*(2), 495–510. https://doi.org/https://doi.org/10.1111/ajps.12116

Baude, W. (2015). Is Originalism Our Law? *Columbia Law Review*, *115*, 2349–2408. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2672631

Black, R. C., & Spriggs II, J. F. (2013). The Citation and Depreciation of U.S. Supreme Court Precedent. *Journal of Empirical Legal Studies*, *10*(2), 325–358. https://doi.org/https://doi.org/10.1111/jels.12012

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/https://doi.org/10.1016/j.jeconom.2020.12.001

Chilton, A. S., & Posner, E. A. (2015). An Empirical Study of Political Bias in Legal Scholarship. *Journal of Legal Studies*, *44*(2), 277–314. https://doi.org/https://doi.org/10.1086/684302

D'Amato, A. (1989). Can Any Legal Theory Constrain Any Judicial Decision? *University of Miami Law Review*, *43*, 513. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1662049

D'Amato, A. (1999). The Effect of Legal Theories on Judicial Decisions. *Chicago-Kent Law Review*, *74*, 517–527. https://scholarship.kentlaw.iit.edu/cgi/viewcontent.cgi?article=3275&context=cklawreview

Detweiler, B. T. (2020). May It Please the Court: A Longitudinal Study of Judicial Citation to Academic Legal Periodicals. *Legal Reference Services Quarterly*, *39*, 87–118. https://doi.org/https://doi.org/10.1080/0270319X.2020.1818475

Diamond, S. S., & Mueller, P. (2010). Empirical Legal Scholarship in Law Reviews. *Annual Review of Law and Social Science*, *6*, 581–599. https://doi.org/https://doi.org/10.1146/annurev-lawsocsci-102209-152848

Edwards, H. T. (1992). The Growing Disjunction Between Legal Education and the Legal Profession. *Michigan Law Review*, *91*(1), 34–78. https://repository.law.umich.edu/mlr/vol91/iss1/3

Feldman, A. (2018). Empirical SCOTUS: With a Little Help from Academic Scholarship. *SCOTUS Blog*. https://www.scotusblog.com/2018/10/empirical-scotus-with-a-little-help-from-academic-scholarship/

Fitzpatrick, B. T., & Varghese, P. K. (2017). Scalia in the Casebooks. *University of Chicago Law Review*, *84*, 2231–42. https://lawreview.uchicago.edu/print-archive/scalia-casebooks

Fronk, C. R. (2010). The Cost of Judicial Citation: An Empirical Investigation of Citation Practices in the Federal Appellate Courts. *Journal of Law, Technology, and Policy*. https://illinoisjltp.com/journal-archive/volume/null/article/the-cost-of-judicial-citation-an-empirical-investigation-of-citation-practices-in-the-federal-appellate-courts

Greene, J. (2009a). On the Origins of Originalism. *Texas Law Review*, *88*(1). https://scholarship.law.columbia.edu/faculty_scholarship/673

Greene, J. (2009b). Selling Originalism. *Georgetown Law Journal*, *97*. https://scholarship.law.columbia.edu/faculty_scholarship/674

Gulati, M., & Sanchez, V. (2002). Giants in a World of Pygmies? Testing the Superstar Hypothesis with Judicial Opinions in Casebooks. *Iowa Law Review*, *87*, 1141.

Ho, D. E. (2005). Why Affirmative Action Does Not Cause Black Students To Fail the Bar. *The Yale Law Journal*, *114*, 1997. https://yalelawjournal.org/pdf/151_1sxzstar.pdf

Holland, P. W. (1986). Statistics and Causal Inference. *Journal of the American Statistical Association*, *81*(396), 945–960. https://doi.org/10.1080/01621459.1986.10478354

Kennedy, D., & Fisher, W. W. (2006). *The Canon of American Legal Thought*. Princeton University Press.

Leiter, B. (2005). *The Blackwell Guide to the Philosophy of Law and Legal Theory*. Wiley.

Leiter, B. (2010). *American Legal Realism*. Wiley.

Maggs, D. B. (1929). Concerning the Extent to Which the Law Review Contributes to the Development of the Law. *Southern California Law Review*, *3*, 181.

Manski, C. (1993). Identification Problems in the Social Sciences. *Sociological Methodology*, 1–56. https://doi.org/https://doi.org/10.22004/ag.econ.292716

Manski, C. (1997). Monotone Treatment Response. *Econometrica*, *65*(6), 1311–1334. https://www.jstor.org/stable/2171738

Manski, C., & Pepper, J. (2018). How Do Right-to-Carry Laws Affect Crime Rates? Coping with Ambiguity Using Bounded-Variation Assumptions. *The Review of Economics and Statistics*, *100*(2), 232–244. https://doi.org/https://doi.org/10.1162/REST_a_00689

Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.

Merrill, T. W., & Smith, H. E. (2001). What Happened to Property in Law and Economics? *The Yale Law Journal*, *111*, 357–398. https://scholarship.law.columbia.edu/faculty_scholarship/411

Merritt, D. J., & Putnam, M. (1995). Judges and Scholars: Do Courts and Scholarly Journals Cite the Same Law Review Articles. *Chicago-Kent Law Review*, *71*, 871. https://scholarship.kentlaw.iit.edu/cklawreview/vol71/iss3/11

Mootz III, F. J. (1994). Legal Classics: After Deconstructing the Legal Canon. *North Carolina Law Review*, *72*(4), 977–1038. https://scholarship.law.unc.edu/nclr/vol72/iss4/4

Newton, B. (2012). Law Review Scholarship in the Eyes of the Twenty-First Century Supreme Court Justices: An Empirical Analysis. *Drexel Law Review*, *4*, 399–416. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2046210

Petherbridge, L., & Schwartz, D. L. (2012). An Empirical Assessment of the Supreme Court's Use of Legal Scholarship. *Northwestern University Law Review*, *106*, 995. https://scholarlycommons.law.northwestern.edu/nulr/vol106/iss3/1

Posner, R. A. (1992). The Deprofessionalization of Legal Teaching and Scholarship. *Michigan Law Review*, *91*, 1921. https://repository.law.umich.edu/mlr/vol91/iss8/2

Posner, R. A. (2016). *Divergent Paths: The Academy and the Judiciary*. Harvard University Press.

Pound, R. (1908). Mechanical Jurisprudence. *Columbia Law Review*, *8*(8), 605–623. https://www.jstor.org/stable/pdf/1108954.pdf

Rumble, W. E. (1965). The Paradox of American Legal Realism. *Ethics*, *75*(3), 166–178. https://doi.org/https://doi.org/10.1086/291541

Sawyer III, L. E. (2018). Principle and Politics in the New History of Originalism. *American Journal of Legal History*, *57*(2), 198–222. https://digitalcommons.law.uga.edu/fac_artchop/1326

Schwartz, D. L., & Petherbridge, L. (2010). The Use of Legal Scholarship by the Federal Courts of Appeals: An Empirical Study. *Cornell Law Review*, *96*, 1345. https://scholarship.law.cornell.edu/clr/vol96/iss6/12

Shapiro, F. R. (1985). The Most-Cited Law Review Articles. *California Law Review*, *73*, 1540. https://www.jstor.org/stable/3480410

Shapiro, F. R. (2021). The Most-Cited Legal Scholars Revisited. *University of Chicago Law Review*, *88*, 1595. https://lawreview.uchicago.edu/print-archive/most-cited-legal-scholars-revisited

Sirico Jr., L. J., & Drew, B. A. (1991). The Citing of Law Reviews by the United States Courts of Appeals: An Empirical Analysis. *University of Miami Law Review*, *45*, 1051–1060. https://repository.law.miami.edu/umlr/vol45/iss5/5

Tamanaha, B. Z. (2009). Understanding Legal Realism. *Texas Law Review*, *87*, 731. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1127178

Thalken, R., Stiglitz, E. H., Mimno, D., & Wilkens, M. (2023). Modeling Legal Reasoning: LM Annotation at the Edge of Human Agreement. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9252–9265. https://arxiv.org/pdf/2310.18440

Traynor, R. J. (1962). To the Right Honorable Law Reviews. *UCLA Law Review*, *10*, 3. https://repository.uclawsf.edu/traynor_scholarship_pub/18

Vladeck, S. (2026). 219. Drunks, Lampposts, and the Birthright Citizenship Case. *One First*. https://www.stevevladeck.com/p/219-drunks-lampposts-and-the-birthright

Whittington, K. (2004). The New Originalism. *Georgetown Journal of Law and Public Policy*, *2*(2), 599–613. https://kewhitt.scholar.princeton.edu/sites/g/files/toruqf3716/files/new_originalism_0.pdf

Appendix A

SUPPLEMENTARY MATERIAL FOR CHAPTER 1

A.1 Types of Reforms

Appendix A.1 includes a table summarizing two approaches to combat evictions: procedural and structural reforms. Procedural reforms tend to focus on making the eviction process "fair" for tenants. Because tenants often face a large resource disadvantage in eviction cases, procedural reforms aim to "level the playing field" for tenants in cases against better-resourced landlords. Procedural policy reforms might include, for example, allowing longer tenant response times, permitting remote court attendance, reducing filing fees, and providing publicly-funded tenant counsel.[1] Procedural proposals prioritize helping tenants *after* they've received an eviction notice.[2]

Structural reforms, by contrast, aim to address the underlying causes of eviction. Because non-payment of rent is the primary reason for eviction, structural reforms tend to focus on reducing rent and housing prices. Because of the large impact of housing supply on prices,[3] the goal is usually (but not always) to increase the available stock of rental housing. Structural reforms may include, for example, subsidizing the building of affordable housing units, building public housing, increasing housing vouchers, or changing income thresholds to receive rental assistance. Structural proposals prioritize helping tenants *before* they've received an eviction notice.

[1] The last proposal is sometimes referred to as "civil *Gideon*," a reference to *Gideon v. Wainwright* (1963), the landmark Supreme Court decision requiring states to provide defense counsel to indigent criminal defendants.

[2] *See, e.g.*, Greiner, Pattanayak, and Hennessy (2012); Cassidy and Currie (2023).

[3] *See* Baum-Snow and Duranton (2025) (discussing extensive literature on housing supply and prices).

Characteristic	Procedural Reforms	Structural Reforms
Primary Goal	Fair process	Eliminate root causes
Policy Focus	Alter legal proceedings	Reduce rent burden and increase affordability
Intervention Timing	After eviction notice	Before eviction notice
Mechanisms	Legal aid, court procedures, fee reductions	Supply-side housing policy, subsidies, vouchers
Typical Examples	Longer response times, remote hearings, public counsel (Civil *Gideon*)	Public housing, expand vouchers, eligibility thresholds

Table A.1: Comparison of Procedural and Structural Eviction Reforms

A.2 Additional Eviction Courthouse Details

The Appendix A.2 plots below from a companion paper are included here for reference and background information.

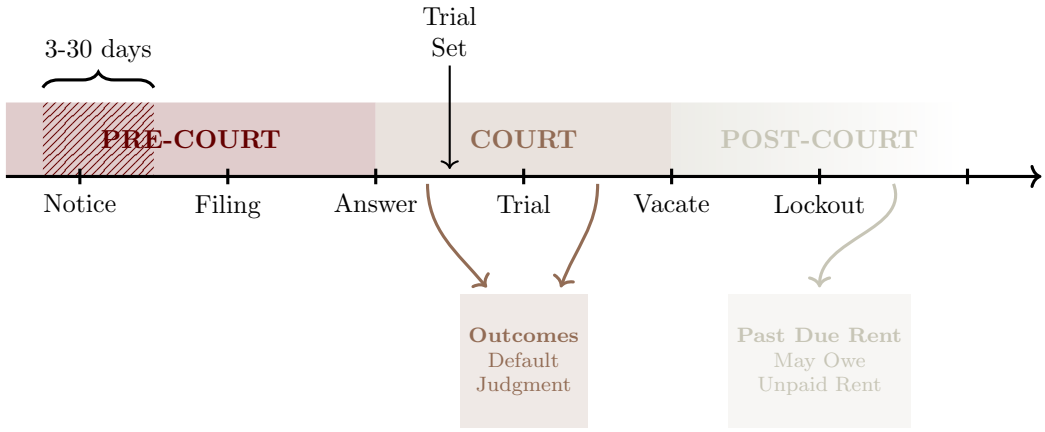

Figure A.1: Eviction Timeline

Note: This figure is reproduced from Estes and Nelson (2026).

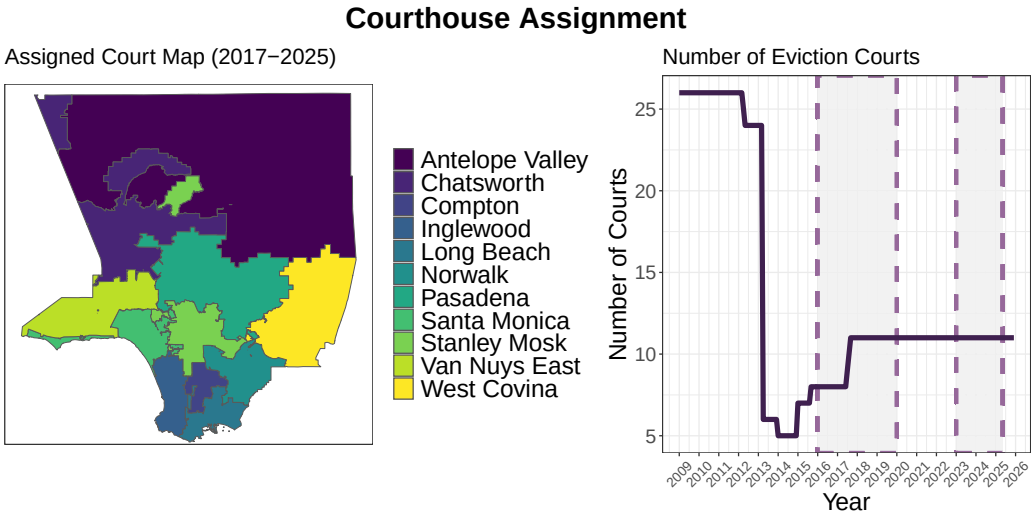

Figure A.2: LA County Eviction Court Assignment

Note: The assignment map (Late 2017–2025) is shown in the left panel. The number of courthouse districts expanded in late 2017, which is shown in the right panel. See Estes and Nelson (2026) for further discussion.

A.3 Additional Result Plots

Appendix A.3 includes additional results plots, focusing mostly on court expansion DID estimates.

Summary Statistics & Description

First, I include summary statistics, which are discussed in-text above.

Default Evictions by Filing Court

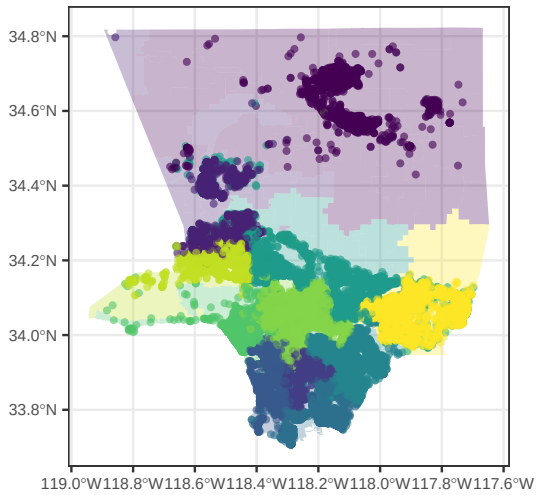

Average Distance−to−Court (miles)

Court	Pre Aug 2017	Post Aug 2017
Antelope Valley	5.24	5.14
Chatsworth	23.10	7.66
Compton	8.93	3.47
Inglewood	9.74	4.61
Long Beach	4.53	3.70
Norwalk	5.66	5.07
Pasadena	9.76	8.80
Pomona	8.32	6.18
Santa Monica	6.53	4.66
Stanley Mosk	3.65	3.70
Van Nuys East	6.66	5.48
West Covina	8.07	6.19

Summary Statistics by Court

Court	White Pr.	Black Pr.	Hispanic Pr.	Female Pr.	Income	Rent	% Represent
Antelope Valley	0.31	0.16	0.43	0.50	50193.49	1249.88	3.67
Chatsworth	0.27	0.10	0.51	0.42	83420.30	1914.55	3.41
Compton	0.20	0.14	0.56	0.49	52832.82	1320.31	4.47
Inglewood	0.30	0.17	0.40	0.41	68648.13	1598.15	3.84
Long Beach	0.27	0.13	0.49	0.44	58281.11	1412.77	1.95
Norwalk	0.14	0.06	0.72	0.44	62482.92	1441.87	2.70
Pasadena	0.26	0.09	0.48	0.39	71084.02	1760.67	4.17
Pomona	0.16	0.06	0.65	0.45	67653.91	1587.16	1.69
Santa Monica	0.35	0.15	0.35	0.39	80849.42	1938.76	3.21
Stanley Mosk	0.26	0.14	0.44	0.38	54355.54	1527.56	4.24
Van Nuys East	0.27	0.11	0.51	0.39	68333.16	1809.77	2.72
West Covina	0.17	0.07	0.62	0.45	66931.51	1623.42	3.09

Figure A.3: Observed Defaults & Descriptive Findings

Note: The top-left map shows where default evictions occur geographically in LA County (2016–2025) as colored points, with the Late 2017–2025 court assignment map plotted beneath it. The top-right table shows the average distance-to-court by observed filing courthouse under the 2016–Aug 2017 and Aug 2017–2025 courthouse assignment policies. The bottom table shows summary averages by observed filing court for all available default records using: imputed race data (first three columns), geo-linked Census data (next three columns), and representation data from the docket records (last column).

Regression Evidence: Distance-to-Court and Default Relationship

These plots are discussed briefly in text above. They show reduced form or regression relationships between distance-to-court and default evictions using assessor data and court records.

First, I use data on all rental units from the LA County Assessor dataset[4] to encode whether each unit at rental addresses has an observed default in a given year.[5] In Appendix Figure A.4, I use this data to estimate the default probability as a function of distance-to-court pooling across years. The average default probabilities are shown in 5-mile distance-to-court intervals. Next, using non-parametric local linear regressions, I also plot estimated curves in Appendix Figure A.5. In the left panel of Appendix Figure A.5, the estimated relationship shows that the default probability is low but roughly increasing as distance-to-court increases (at least up to a certain distance). Finally, I estimate the relationship between the number of observed defaults and distance-to-court for addresses observed in the court docket records only. This exercise gets at the spatial concentration of defaults across LA County and is shown in the right panel of Appendix Figure A.5: the number of observed default evictions is a complicated non-linear function of distance-to-court.[6]

Because there are some disadvantages to using either the Assessor data or court record counts only, I also repeat the reduced-form estimation in Appendix A.3 using newly available on the near universe of eviction notices in LA City. Using a new legal reporting requirement of eviction notices to the LA City Controller's Office,[7] I am able to define the population of at-risk tenant addresses as those appearing in the City Controller eviction notice data.

[4] *See also* Estes and Nelson (2026).

[5] For example, if unit A at address i in year t has a default but units B and C do not, then unit A is encoded as a 1 whereas B and C are encoded as zeroes.

[6] The number of defaults is highest at the smallest distances, largely due to a high number of evictions in downtown LA relatively close to the Stanley Mosk courthouse. The estimated curve then decreases between 0–5 miles, before increasing from around 5–13 miles, at which point it decreases and levels off at 1 default. This is because most addresses at the highest distance-to-court values have only a few observed defaults.

[7] Failure to report the eviction notice to the City Controller constitutes an affirmative defense to the eviction.

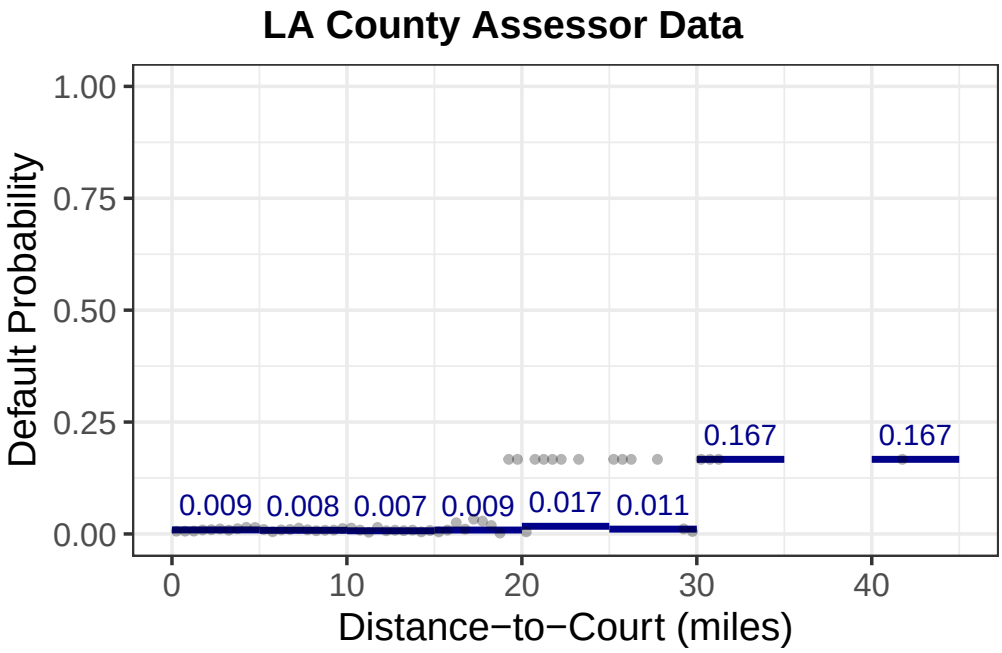

Figure A.4: Reduced Form (LA County Assessor): 5-mile Binned Averages

Note: The blue lines are binned averages (every 5 miles) and the transparent points show default probabilities every 0.5 miles.

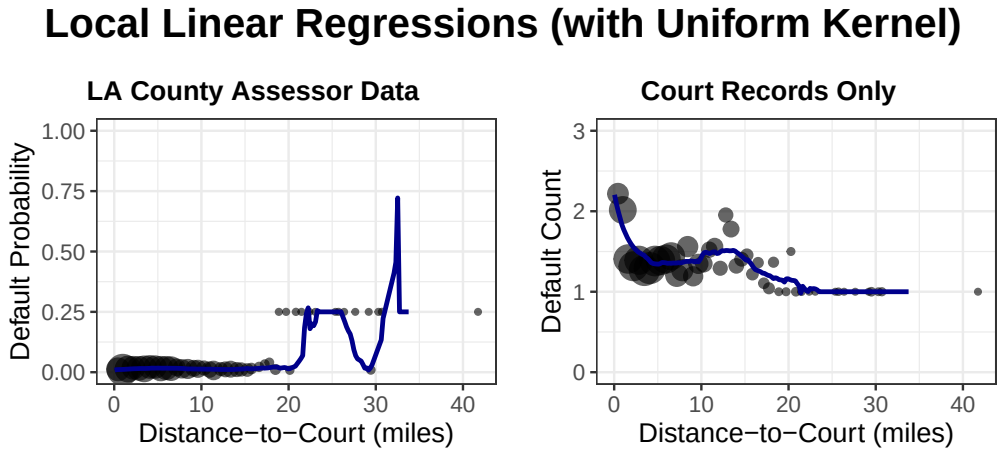

Figure A.5: Reduced Form Relationships: Default Probability & Default Counts

Note: The blue curves are local linear regression estimates (with a uniform kernel and bandwidth 5km) of the relationship between distance-to-court and default probability (left panel) or default counts (right panel).

DID Design Illustration

This subsection illustrates the DID counterfactual intuition. The plot below illustrates the average treatment effect on the treated estimation (ATT), where the ATT is the difference between the observed treatment group average outcome and the counterfactual outcome. The counterfactual outcome is assumed parallel to the control group outcomes and is shown below as a dashed blue line.

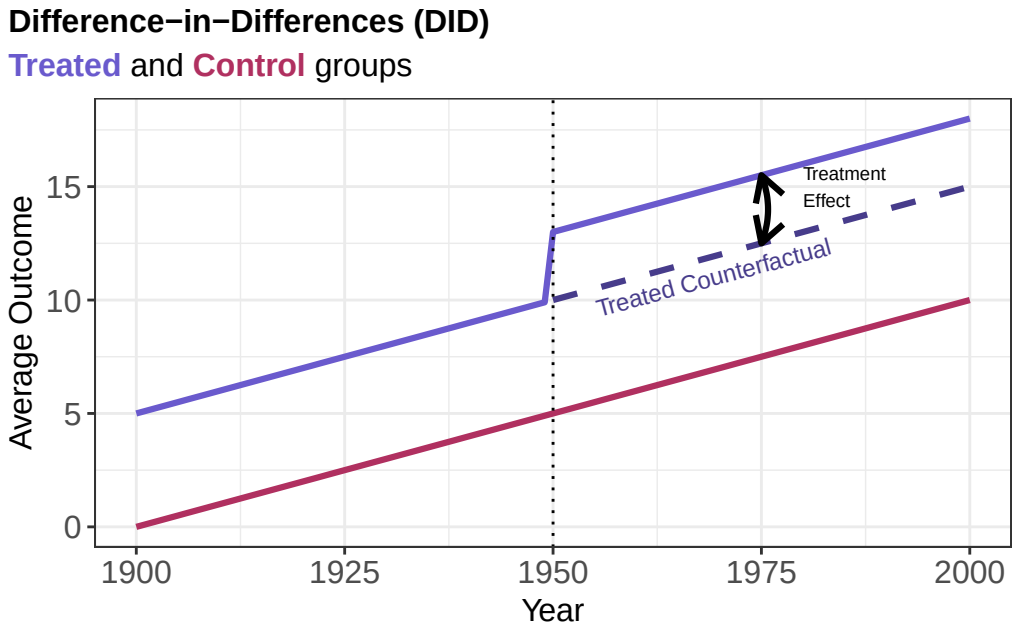

Figure A.6: Difference-in-Differences (DID) Method

Note: Average outcomes for the treated and control groups are shown in blue and red, respectively. The dotted vertical line (at 1950) is the treatment date for the Treated group. The dashed blue line is the counterfactual for Treated units: this is the average outcome Treated units would have experienced if they were untreated. The average treatment effect on the treated (ATT) is the difference between the solid and dashed blue lines.

LA City Controller Data Only Results

To address potential concerns with using either assessor-based addresses as the population of rental addresses or only addresses with at least one observed default, I consider here using new data from the LA City Controller's Office. A rule change in 2023 mandated that all landlords in LA City record eviction notices with the City Controller's Office. The data includes address information for each eviction notice, allowing me to identify the universe of apartment buildings at-risk of eviction or default eviction. Using only these addresses in LA City, I reproduce the main-text results: reduced-form relationships and dynamic DID event-study type plots.

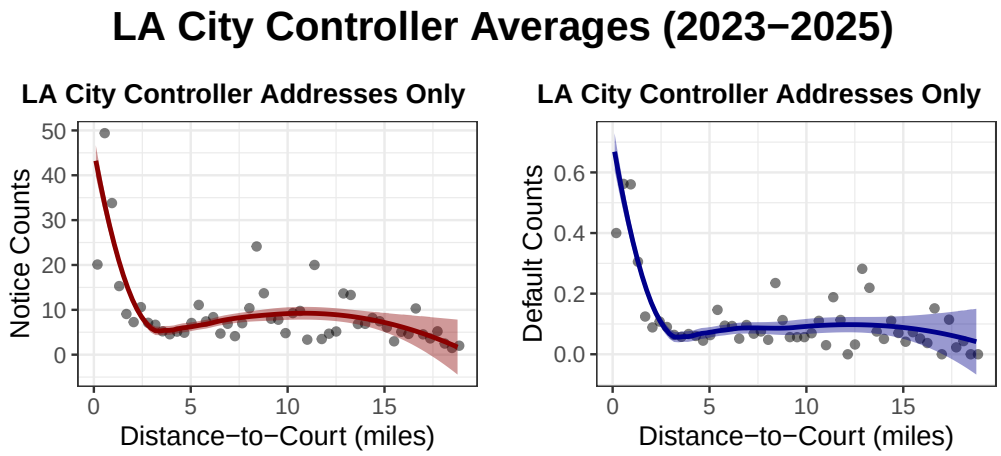

Figure A.7: Reduced Form (LA City Controller Data)

Note: The left panel shows the smoothed LOESS estimates of the average number of notice counts as a function of distance-to-court. The right panel shows the smoothed LOESS estimates of the average number of defaults as a function of distance-to-court.

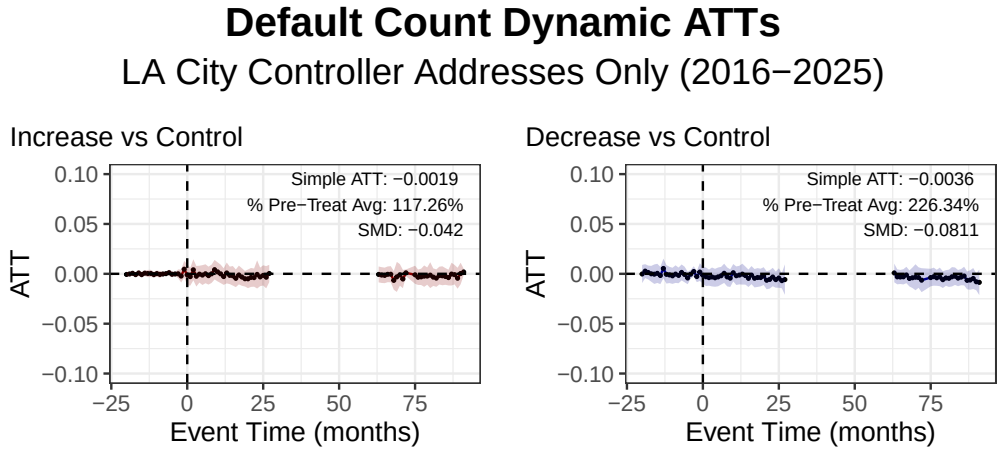

Figure A.8: Dynamic DID Estimates (LA City Controller Addresses Only)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

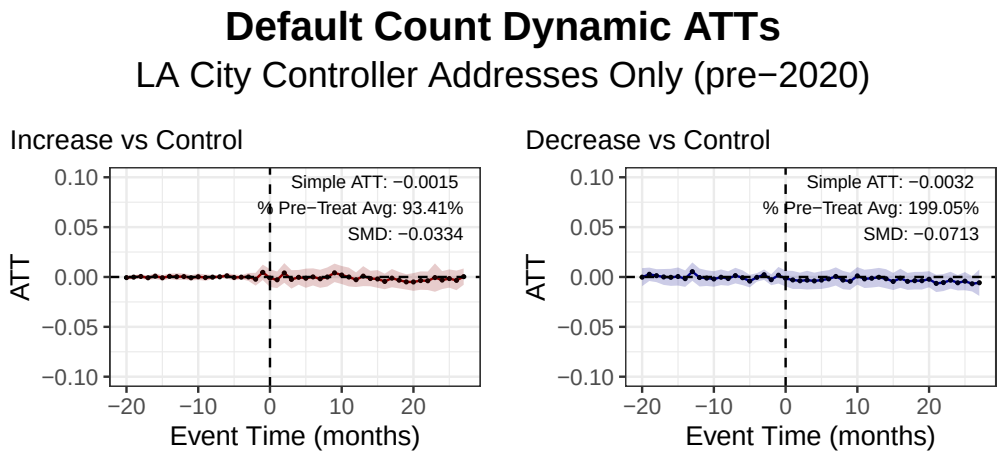

Figure A.9: Dynamic DID Estimates Pre-2020 (LA City Controller Addresses Only)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

Assessor Unit-Weighted Results

I also include results with outcomes weighted by the number of apartment units. Using the LA Assessor data, I match addresses in the eviction dataset to assessment data on how many units there are in each building. The exercise is complicated from a data perspective, as some rental units may be empty or matches to assessment data may be subject to assessor recording or matching error. But the DID-style estimates allow us to estimate how the court expansion affected tenants in all units living at the addresses in the eviction dataset. The weighted estimates are produced below for all years and for the pre-2020 years separately. The results are subject to very large uncertainty (much larger SEs) and insignificant in both the point estimates and the aggregate. Moreover, the direction of the effect does not align with expectations for the Increase versus Control comparisons.

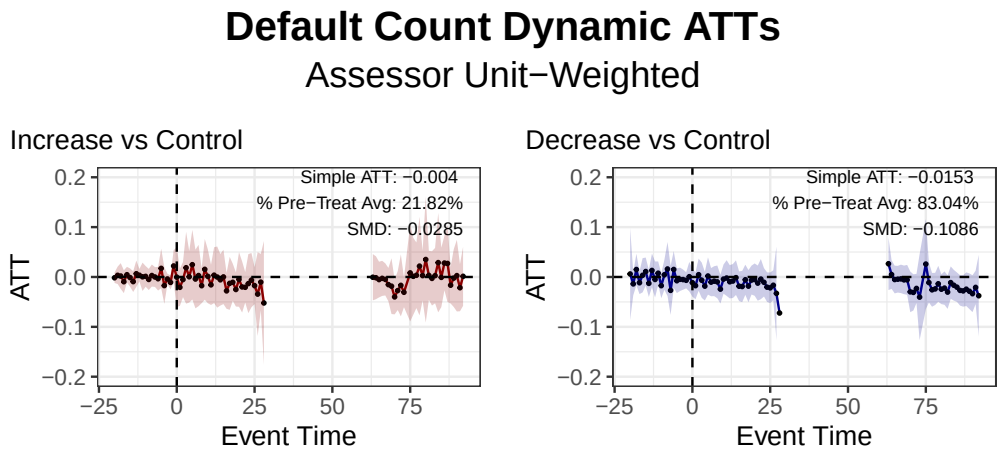

Figure A.10: Dynamic DID Estimates (Assessor Unit-Weighted)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

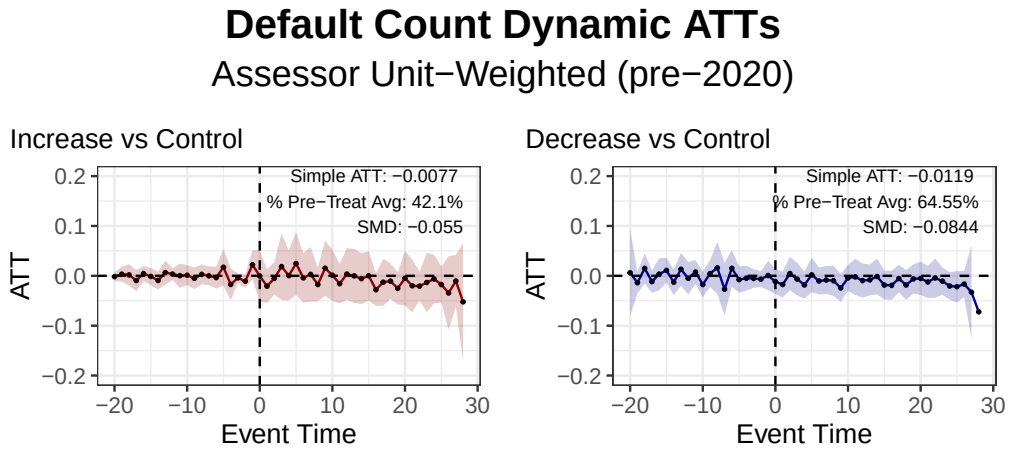

Figure A.11: Dynamic DID Estimates Pre-2020 (Assessor Unit-Weighted)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

Pre-Aug 2017 Buildings Only Results

Finally, to address potential bias from including all buildings with observed evictions in the main analysis, I include results for only those buildings with observed

default evictions pre-Aug 2017. The results for the Increase-Treated and Large Increase-Treated versus Control group are both insignificant; moreover, the Large Increase ATTs are mostly negative, which does not align with the predicted effect: those buildings with an increase in distance-to-court are expected to have increased defaults. Conversely, the Decrease-Treated and Large Decrease-Treated comparisons with the Control group are significant, but neither effect aligns with expectations: the point estimates and aggregate simple ATTs are all positive. The results reinforce the conclusions in the main body of the paper: either there is large uncertainty (Increase comparisons) or the results do not align with expectations (Decrease comparisons).

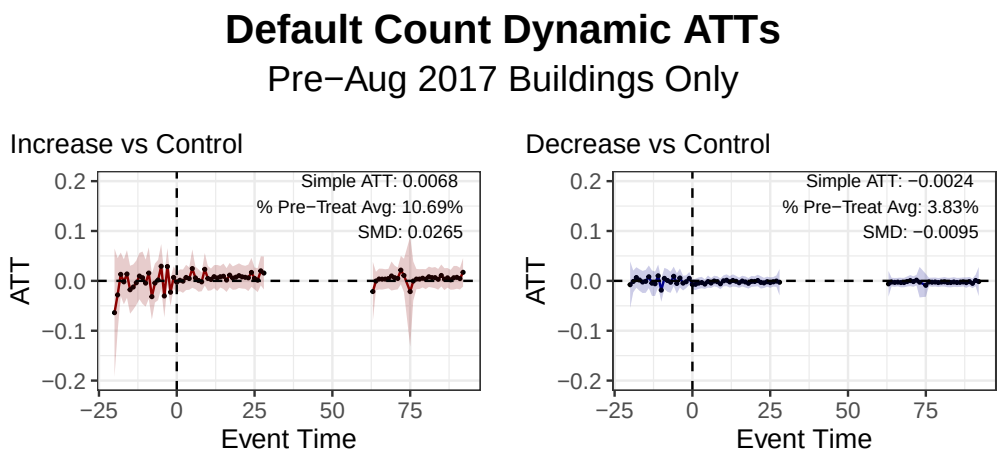

Figure A.12: Dynamic DID Estimates (pre-Aug 2017 Buildings Only)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

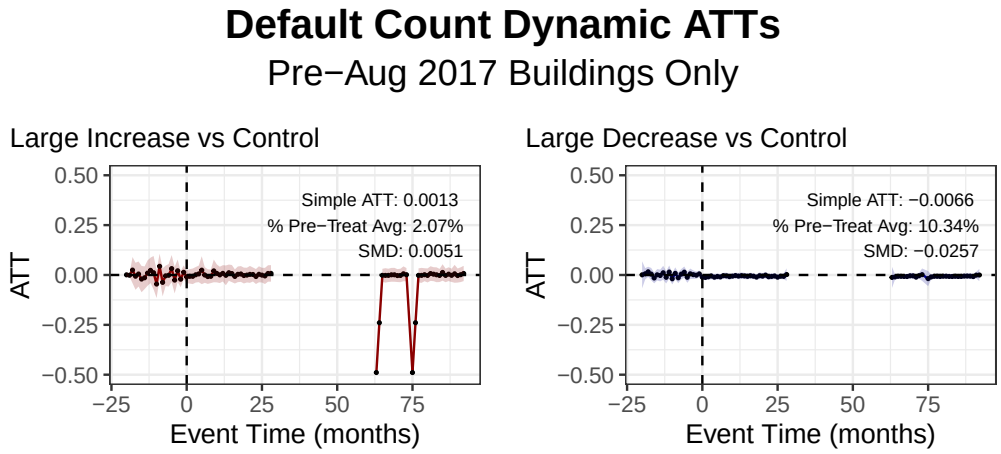

Figure A.13: Dynamic DID Estimates: Above-Median Comparisons (pre-Aug 2017 Buildings Only)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

A.4 Commute Times

Appendix A.4 includes additional findings using commute times rather than distance-to-court as the treatment variable of interest. Commute times were calculated for each address using the Mapbox API.

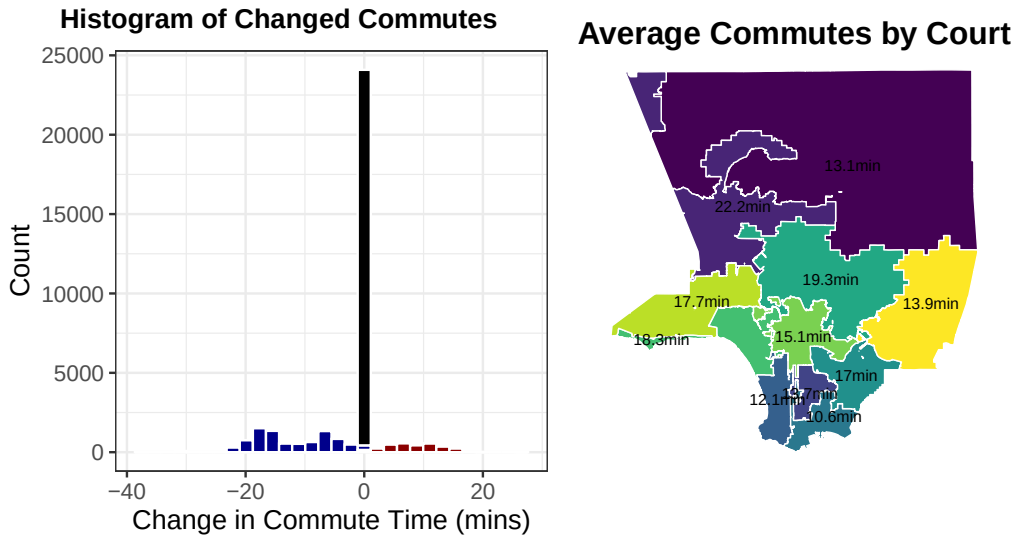

Correlation: Distance−to−Court and Commute Time

Overall Correlation = 0.843

Filing Courthouse	Correlation	N	Filing Courthouse	Correlation	N
Antelope Valley	0.968	2964	Pasadena	0.916	6417
Chatsworth	0.949	1639	Pomona	0.948	1777
Compton	0.877	1632	Santa Monica	0.924	3890
Inglewood	0.953	2004	Stanley Mosk	0.932	12406
Long Beach	0.912	6505	Van Nuys East	0.940	2717
Norwalk	0.963	4890	West Covina	0.760	2713

Figure A.14: Commute Time Histogram, Averages, and Correlations

Note: The top-left panel shows the histogram of changes in commute times following the policy shock. The top-right panel maps the average commutes by court district under the post-reform assignment map. The bottom table shows the correlation between distance-to-court and commute times for each observed filing courthouse (overall corr. ≈ 0.843).

DID results are shown below. As in the other cases, the results are not significant, do not align with expectations as to the sign of the effect (e.g., Decrease-Treated and

Large Decrease-Treated units have positive aggregate simple ATTs), or both. Note as well the small standardized mean differences (SMDs) in each panel.

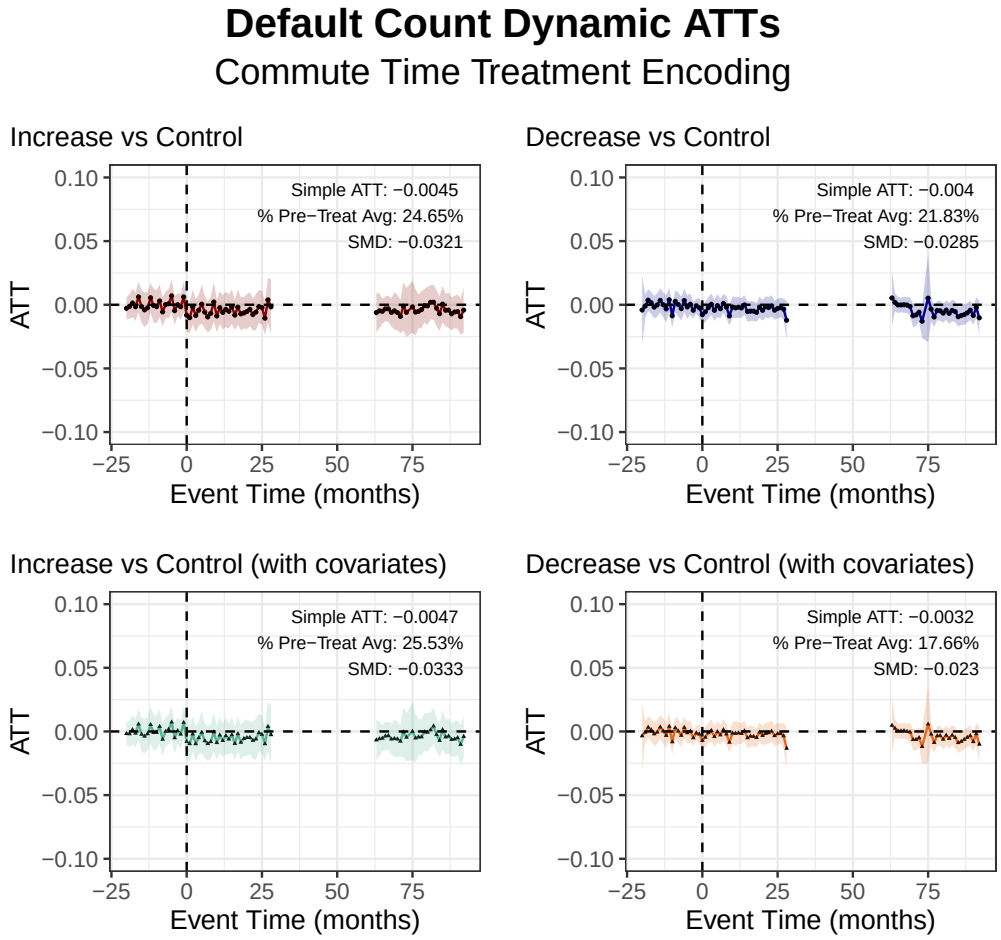

Figure A.15: Dynamic DID Estimates: Comparisons with Controls (Commute Times)

Note: The panels show estimates of the average treatment effect on the treated for Increase-Treated (left panel) and Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations. The covariate models include gender, race, and income variables.

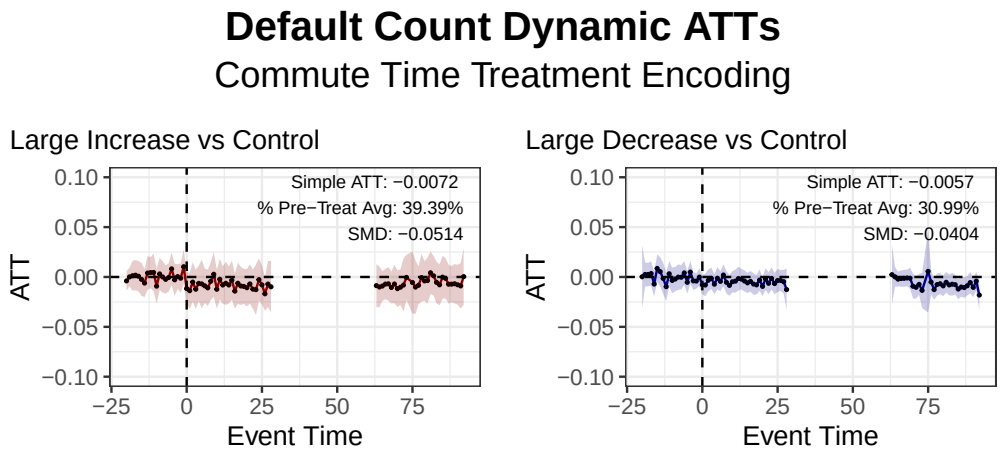

Figure A.16: Dynamic DID Estimates: Above-Median Comparisons with Controls (Commute Times)

Note: The panels show estimates of the average treatment effect on the treated for Large Increase-Treated (left panel) and Large Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations.

A.5 Relaxing Parallel Trends

Appendix A.5 includes additional results which weaken the key identification assumption. Meta-analyses of existing DID research suggests violations of parallel trends is common and requires validation, sensitivity analysis, or relaxation.[8] Although there are a few ways to relax the parallel trends assumption, I examine here a Manski-style partial identification approach following C. Manski and Pepper (2018) and Estes and Clark (forthcoming 2026). This approach accommodates a form of sensitivity analysis to a tuning parameter, which I discuss below and follows the exposition in Estes and Clark (forthcoming 2026).

To begin, I specify the treatment effect estimand of interest: the *average treatment effect on the treated* ("ATT") for units in two treatment cohorts $\vec{k} = (k_1, k_2, \dots)$ and $\vec{k}' = (k'_1, k'_2, \dots)$. This notation for treatment $\vec{k}$ says that k_1 is the treatment level received at time 1, k_2 the treatment level at time 2, etc. The average treatment effect in time t on units treated with $\vec{k}$:

$$ATT(t, \vec{k}, \vec{k}') = \mathbb{E}[Y_{it}(\vec{k}) - Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}]$$

This is the average treatment effect from switching to treatment $\vec{k}'$ for units on treatment $\vec{k}$ at time t. The identification problem is that $Y_{it}(\vec{k}')$ is unobservable for units with observed treatment $\vec{K}_i = \vec{k}$.

For two distinct treatment histories $\vec{k}$ and $\vec{k}'$, let the first time period the treatments diverge (i.e., do not agree) be denoted $T_{div} = \text{argmin}_t(k_t \neq k'_t)$. I use information from the pre-divergence periods (i.e., time periods before T_{div}) and the following assumption to impute counterfactual outcomes.

Assumption 5. *[Average Bounded Deviation] Let $T_{div} =_t (k_t \neq k'_t)$ for distinct treatments $\vec{k}$ and $\vec{k}'$. Then the following holds for each i, $t \geq T_{div}$, any $(\vec{k}, \vec{k}')$ pair, and some $C \in \mathbb{R}_+$:*

$$\left|\mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}']\right| \leq C \cdot \max_{t<T_{div}} \left|\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}']\right|$$

where $C = C(t, \vec{k}, \vec{k}')$ can depend on the time period and both treatments.

This assumption says the counterfactual average outcome from treating with $\vec{k}'$ instead of $\vec{k}$ for units with observed treatment $\vec{K}_i = \vec{k}$ is "near" the observed outcome for units with treatment $\vec{k}'$. Specifically, the counterfactual average for the

[8] See Chiu et al. (2026).

$\vec{k}$-treated group is within C times the maximum absolute difference in means for units on the two treatments before they diverge in time.

The method is illustrated graphically below in Appendix Figure A.17. In the left panel, the two lines are average outcomes for Treated (blue) and Untreated (red) units. The black-dotted vertical line is the path divergence time period, which in this example is 1950. The blue-striped region is the counterfactual region for blue Treated units: this is where counterfactual outcomes for Treated units would be if they had been Untreated. The maximum pre-divergence deviation is 1.69 and is shown by the black vertical line around 1925. In the right panel, the y-axis is the average treatment effect on the treated (ATT), with a black-dotted horizontal line plotted for zero effect. The upper and lower bounds on the ATT for each post-1950 time are plotted as dashed-yellow and solid-green lines, resp.

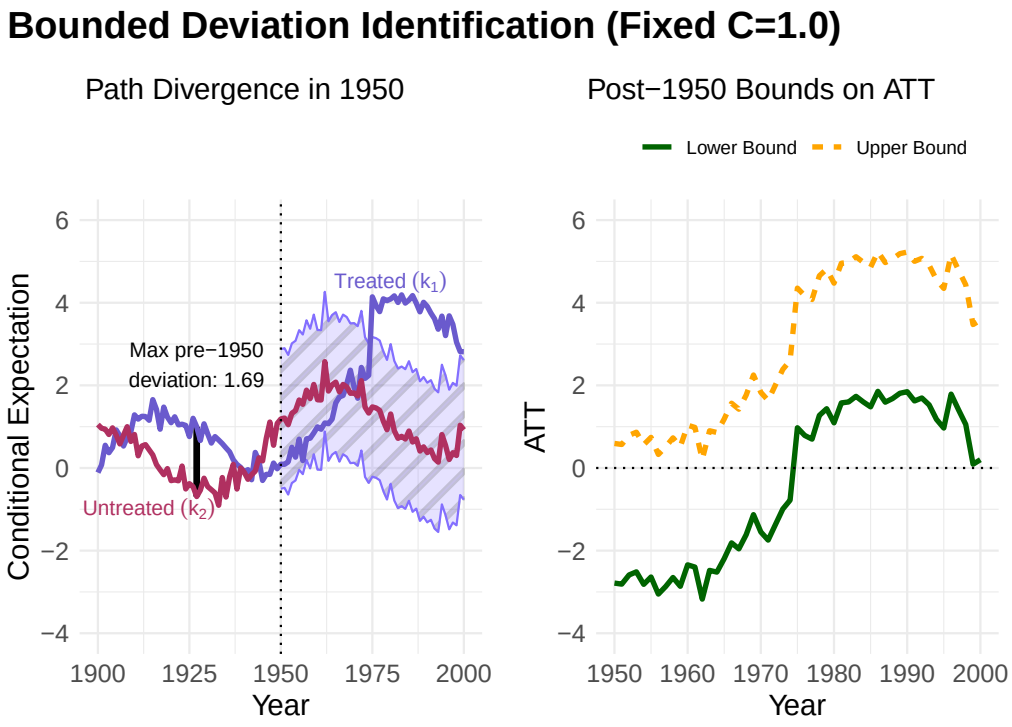

Figure A.17: Partial Identification in the Bounded Deviations Approach

Note: The figure is reproduced from Estes and Clark (forthcoming 2026).

The upper and lower bounds implied by this exercise for three different C values are shown in Appendix Figure A.18. The top row shows estimated ATT bounds comparing the Increase-Treated and Control units, whereas the bottom row shows

estimated ATT bounds comparing the Decrease-Treated and Control units. Note that the ATTs have been standardized by the control group standard deviations in the pre-divergence period for interpretation in terms of standard effect sizes. Because the standardized ATT upper and lower bounds are generally less than 0.20, this suggests that effect of the Aug 2017 court expansion are small even when weakening the parallel trend assumption.

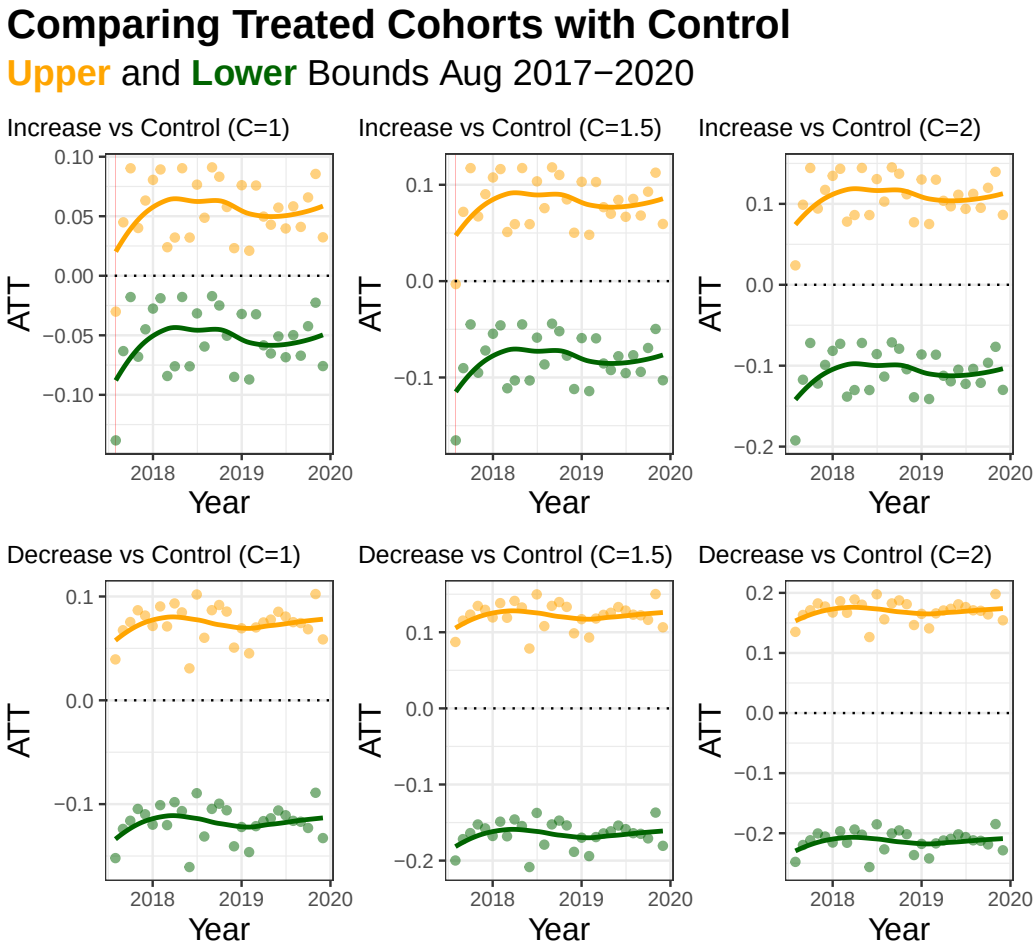

Figure A.18: Relaxing Parallel Trends: Bounded Deviations

Note: The upper (yellow) and lower (green) bounds on the ATT in each post-divergence time are shown from Aug. 2017–Dec. 2020. The points show standardized ATT estimates in each time period, and the solid lines are smoothed fits of those point estimates.

A.6 All Observed Evictions

Appendix A.6 includes DID results using any observed eviction as the outcome-of-interest. Results from the baseline specification with and without the same covariates are shown below.

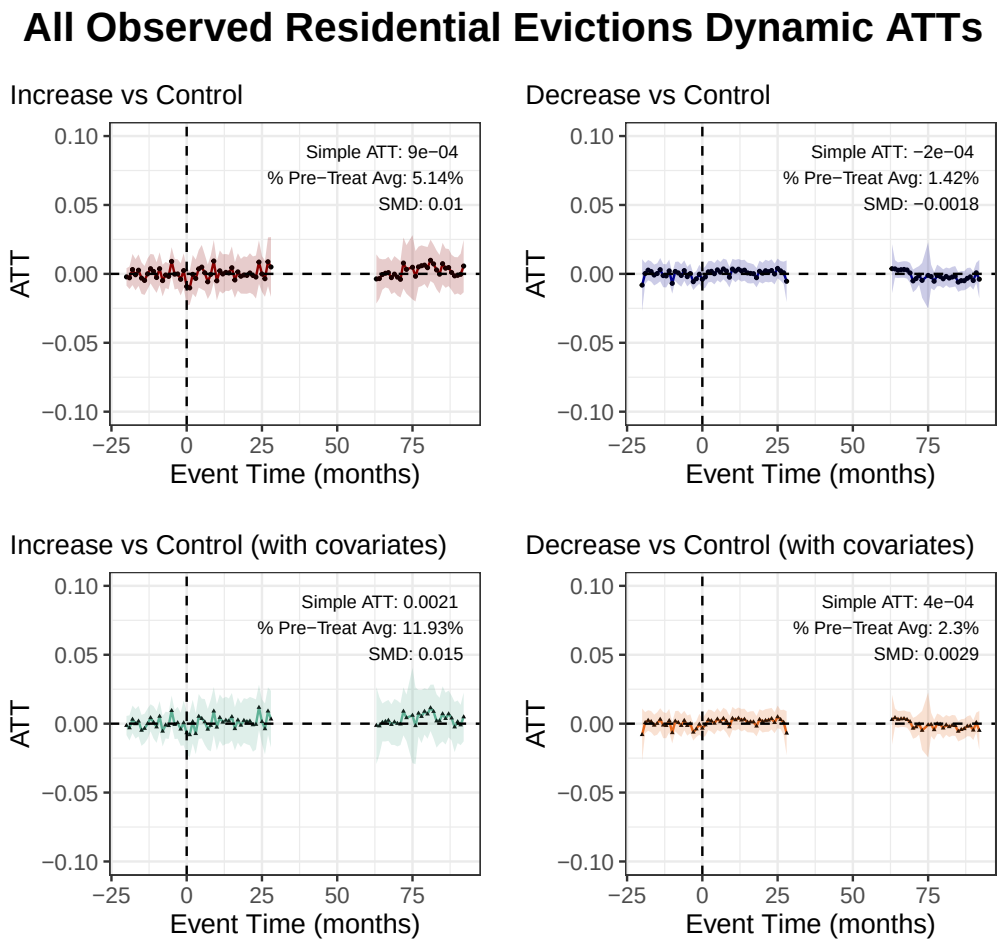

Figure A.19: Dynamic DID Estimates: All Observed Outcomes

Note: The panels show estimates of the average treatment effect on the treated for Large Increase-Treated (left panel) and Large Decrease-Treated (right panel) units. The Simple ATT aggregates post-treatment ATT estimates as in Callaway and Sant'Anna (2021a). The percent pre-treatment average is the ratio of the ATT and the control cohort average outcome in the pre-treatment period. The standardized mean differences (SMDs) suggest small effects under standard rule-of-thumb interpretations. The covariate models include gender, race, and income variables.

A.7 Covariate Specification Curves

Appendix A.7 includes specification curves of ATT estimates for different covariate combinations with the baseline specification of addresses and two outcomes: defaults only (top panels) and all observed eviction outcomes (bottom panels).

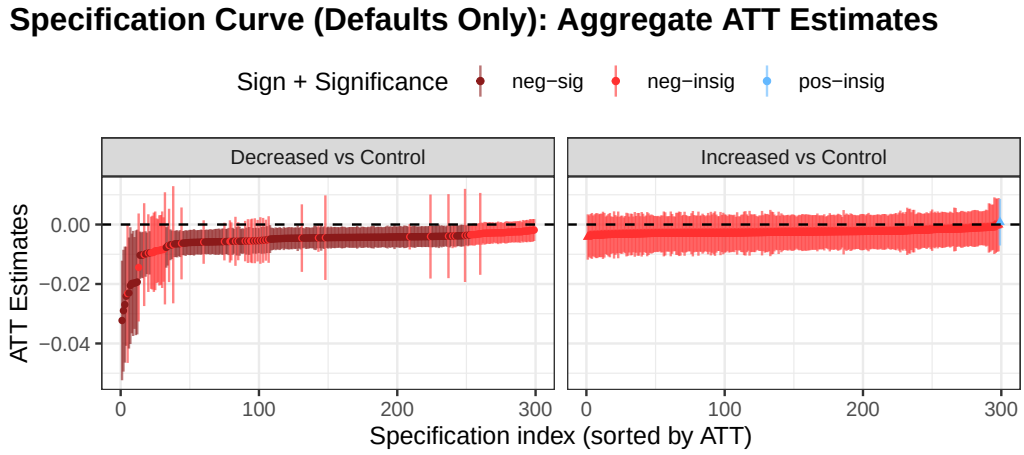

Figure A.20: Covariate Specification Curve (Defaults Only)

Note: Positive and negative aggregate ATT estimates are shown in blue and red, resp. Dark colors (resp. light colors) indicate significant (resp. insignificant) aggregate estimates.

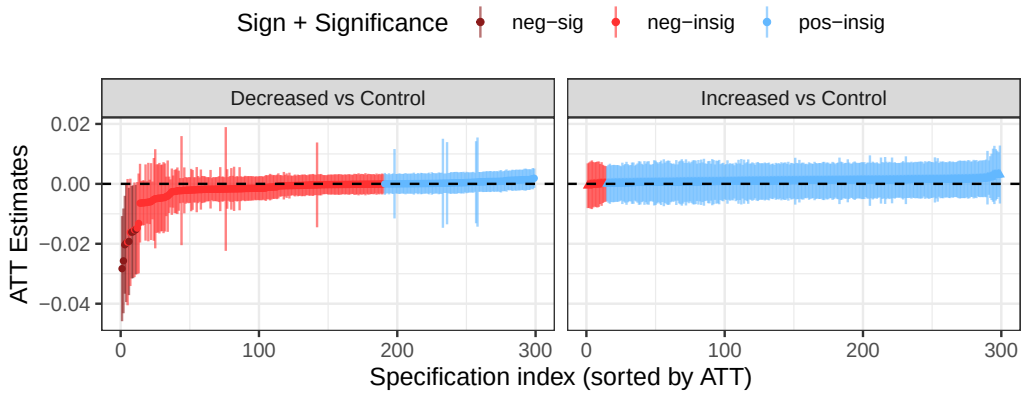

Figure A.21: Covariate Specification Curve (All Outcomes)

Note: Positive and negative aggregate ATT estimates are shown in blue and red, resp. Dark colors (resp. light colors) indicate significant (resp. insignificant) aggregate estimates.

A.8 Comparative California Eviction Rules

Appendix A.8 includes additional information on California eviction rules to facilitate comparison and further additional study. In particular, the results here may help shed light on where the LA County specific estimates are likely or unlikely to generalize, as well as suggest potential counties where distance-to-court may be a high tenant cost.

County Filing Rules Summarized

The table below summarizes the county filing rules. Note that Contra Costa, Riverside, and San Bernardino Counties use multiple courthouses to hear eviction cases. Contra Costa uses a zip-code filing rule similar to LA County, and shown below. Riverside assigns cases to the courthouse nearest the tenant's rental unit. San Bernardino uses a complicated neighborhood-case type filing rule, which is discussed further below. Finally, please note that for Kern County two buildings hear civil cases: the Metropolitan Division (1415 Truxtun Ave) and Metropolitan Division Justice Building (1215 Truxtun Ave). Although it is somewhat unclear whether unlawful detainer cases are always assigned to one building or another, they are located across the street from one another. Accordingly, the 1215 Truxtun Ave lat-long coordinates were used for approximating the distance-to-court Kern density.

Table A.2: Courthouses and Locations for Estimating Distance-to-Court Densities

County	Court Code	Courthouse Name	Latitude	Longitude
Alameda County	A1	Hayward Hall of Justice	37.65760	-122.0941
Contra Costa County	CC1	George D. Carroll Courthouse (Richmond)	37.93233	-122.3338
Contra Costa County	CC2	Richard E. Arnason Justice Center (Pittsburg)	38.01950	-121.8899
Contra Costa County	CC3	Wakefield Taylor Courthouse (Martinez)	38.01850	-122.1337
Fresno County	F1	B.F. Sisk Courthouse	36.73917	-119.7863
Kern County	K1	Metropolitan Division	35.37289	-119.0158
Orange County	O1	Central Justice Center	33.75136	-117.8754
Riverside County	R1	311 E. Ramsey St. (Banning)	33.92652	-116.8727
Riverside County	R2	265 N. Broadway (Blythe)	33.61335	-114.5971
Riverside County	R3	505 S. Buena Vista (Corona)	33.87802	-117.5770
Riverside County	R4	27401 Menifee Center Dr (Menifee)	33.68244	-117.1812
Riverside County	R5	13800 Heacock St. (Moreno Valley)	33.92033	-117.2429
Riverside County	R6	3255 E. Tahquitz Canyon Way (Palm Springs)	33.82234	-116.5113
Sacramento County	S1	Carol Miller Justice Center	38.55072	-121.4067
San Bernardino County	SB1	Barstow District	34.89396	-117.0254
San Bernardino County	SB2	Fontana District	34.10043	-117.4167
San Bernardino County	SB3	Joshua Tree District	34.13448	-116.2757
San Bernardino County	SB4	Needles District	34.83527	-114.6174
San Diego County	SD1	Hall of Justice	32.71628	-117.1660
San Francisco County	SF1	Civic Center Courthouse	37.78045	-122.4188
San Joaquin County	SJ1	180 E Weber Ave.	37.95332	-121.2890
San Mateo County	SM1	800 North Humboldt St. (Central Branch)	37.58179	-122.3285
Santa Clara County	SC1	Downtown Superior Court	37.33838	-121.8933
Ventura County	V1	Hall of Justice (800 South Victoria Ave.)	34.26847	-119.2101

The locations of the different counties and eviction court locations studied here are plotted below for additional reference.

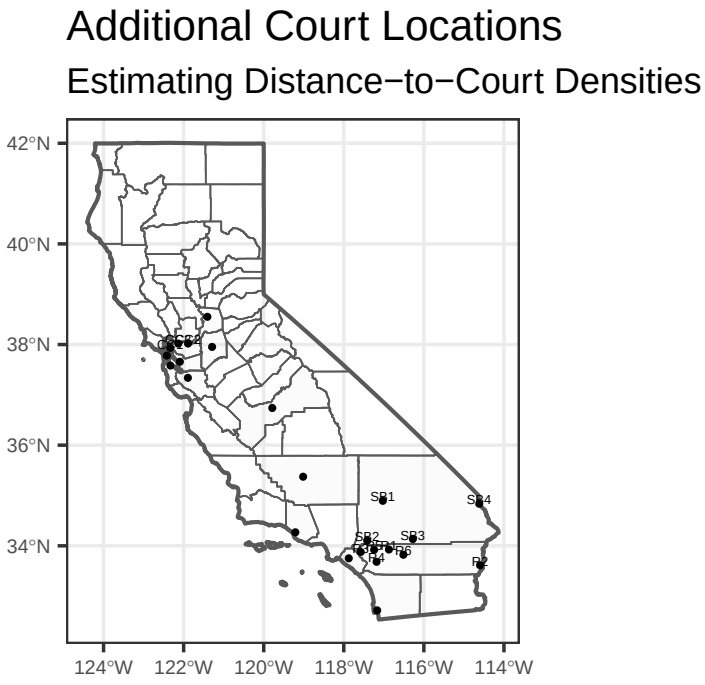

Figure A.22: Additional Eviction Court Locations

Contra Costa County Eviction Filing Rule

The Contra Costa County filing rule is shown below.[9] Like LA County, Contra Costa County uses zip codes to sort eviction cases to courthouses. Cities are also part of the filing rule but, where all cities with a given zip code are assigned to same courthouse, the rule below is simplified. There are three eviction courthouses in total.

[9] https://contracosta.courts.ca.gov/forms-filing/where-file

Table A.3: Contra Costa County Filing Rule

Court Code	Courthouse Name	Zip Codes
CC1	George D. Carroll Courthouse (Richmond)	94525
		94707
		94564
		94530, 94801 - 94808, 94820, 94850
		94547, 94572
CC2	Richard E. Arnason Justice Center (Pittsburg)	94509, 94531
		94565
		94511
		94513
		94505, 94514
		94548
		94561
CC3	Wakefield Taylor Courthouse (Martinez)	94507
		94516
		94517
		94518 - 94522, 94524, 94527, 94529
		94506, 94526
		94528
		94549
		94553
		94556, 94570, 94575
		94563
		94523
		94569
		94582, 94583
		94595 - 94598

San Bernardino County Eviction Filing Rule

The San Bernardino County eviction filing rule is shown below for incorporated cities. Because of the complexity of the filing rule and the difficulty in ascertaining the boundaries of unincorporated areas, a small number of Eviction Lab census tracts were excluded. Essentially the difficulty is in determining to which neighborhood an unincorporated census tract belongs and then mapping that into one of the "areas" and case type served by a particular San Bernardino County courthouse.

Table A.4: San Bernardino County Filing Rule

City Name	Court Code	Courthouse Name
ADELANTO CITY	SB1	Barstow District
BARSTOW CITY	SB1	Barstow District
BIG BEAR LAKE CITY	SB1	Barstow District
HESPERIA CITY	SB1	Barstow District
TOWN OF APPLE VALLEY	SB1	Barstow District
VICTORVILLE CITY	SB1	Barstow District
CHINO CITY	SB2	Fontana District
CHINO HILLS CITY	SB2	Fontana District
COLTON CITY	SB2	Fontana District
FONTANA CITY	SB2	Fontana District
GRAND TERRACE CITY	SB2	Fontana District
HIGHLAND CITY	SB2	Fontana District
LOMA LINDA CITY	SB2	Fontana District
MONTCLAIR CITY	SB2	Fontana District
ONTARIO CITY	SB2	Fontana District
RANCHO CUCAMONGA CITY	SB2	Fontana District
REDLANDS CITY	SB2	Fontana District
RIALTO CITY	SB2	Fontana District
SAN BERNARDINO CITY	SB2	Fontana District
UPLAND CITY	SB2	Fontana District
YUCAIPA CITY	SB2	Fontana District
TWENTYNINE PALMS CITY	SB3	Joshua Tree District
YUCCA VALLEY CITY	SB3	Joshua Tree District
NEEDLES CITY	SB4	Needles District

Commute Time Densities by CA County

I also include density estimates using commute-to-court times rather than distance-to-court times across CA counties. Using Eviction Lab data and the same algorithm as discussed for estimating distance-to-court, I estimate commute-to-court time densities in Appendix Figure A.23. By this metric, the LA reform in late 2017 reduced average commute times from 18.2 min to 15.6 min. The relative ranks of the counties is mostly unchanged for the highest and lowest tenant cost counties: for example, Kern, San Bernardino, San Diego, and Ventura Counties are the highest tenant cost counties whether we use distance-to-court or commute-to-court times as the metric.

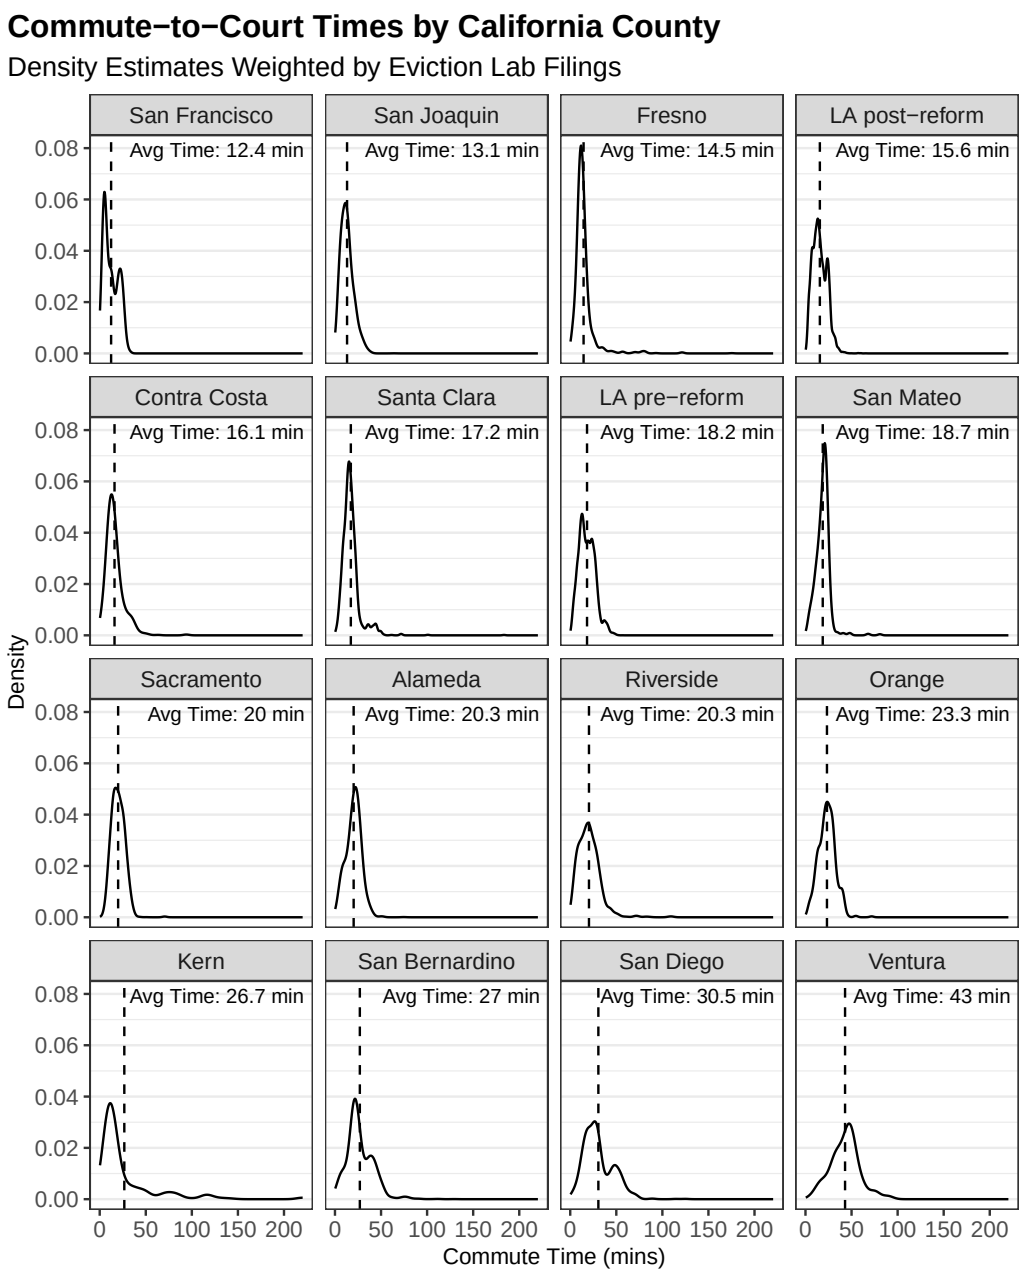

Figure A.23: Commute-to-Court Times by CA County

Note: Except for LA County, the densities are approximated using Eviction Lab data on census tract filings from 2010-2018. The centroid of each census tract is used to compute commute-to-court time densities weighted by the number of Eviction Lab filings reported. The average commute-to-court time is shown as a vertical dashed line in each county panel. The LA County pre- and post-reform density estimates are constructed using the collected micro-data discussed in-text.

Appendix B

SUPPLEMENTARY MATERIAL FOR CHAPTER 2

B.1 Creating the Courthouse Assignment Map

Appendix B.1 details the creation of the courthouse assignment map and shapefiles. As discussed in the main body of the paper, the LA County filing rule for evictions is unique to eviction cases. Assignment to a court is determined by a neighborhood–zip code pair. To create the assignment shpaefile, we started with zip-codes where the filing rule uniquely specifies a courthouse. Using zip code shapefiles from LA County geohub, this determines the courthouse assignment map for most of LA County. In the remaining zip codes where there are multiple eviction courthouses, we had to determine the neighborhood boundaries within the zip code. We used other LA County Geohub shapefiles (i.e., the census tract shapefiles) to determine neighborhoods within the remaining unassigned zip codes by matching LASC assignment table neighborhood names to census tract neighborhood names in the LA County Geohub shapefiles.

Next, if a particular neighborhood/city is always assigned to a particular courthouse (across all zip-codes), then we assigned the census tracts with those neighborhood/city names to that courthouse. We then manually created the courthouse assignment shapefiles for zip codes where assignment is done by a street divider. This involved using street and highway shapefiles from LA County and the Census shapefiles. Finally, for any remaining unassigned census tracts, we used the scraped eviction data from 2018 to determine where the vast majority of cases for that census tract are assigned. We assumed that the majority assigned courthouse was the correct one, i.e., we minimized classifying observed evictions as "misfiled."

Due to the complex governing structure of LA County, some of the city/neighborhoods listed in the official Zip Code Table are not official, incorporated governing regions. Therefore, the precise boundaries of some neighborhoods (and, consequently, some city-zip pairs) is difficult to ascertain. Fortunately, 89% (480 of 539) of the zip codes in the Zip Code Table are uniquely assigned to one courthouse. For example, in Appendix Table B.1 the highlighted rows for 90006 show that both Koreatown and Los Angeles evictions are assigned to Stanley Mosk Courthouse.

Table B.1: Simplifying Assignment: Unique Courthouses

Zip Code	City/Neighborhood	Courthouse
90006	KOREATOWN	Stanley Mosk Courthouse
90006	LOS ANGELES	Stanley Mosk Courthouse
90007	LOS ANGELES	Stanley Mosk Courthouse
90008	BALDWIN HILLS	Stanley Mosk Courthouse
90008	CRENSHAW	Stanley Mosk Courthouse
90008	LIEMERT PARK	Stanley Mosk Courthouse
90008	LOS ANGELES	Santa Monica Courthouse

Accordingly, we work with a Simplified Zip Code Table that aggregates neighborhoods in the event that all neighborhoods within a zip code are assigned to one unique courthouse. For instance, Appendix Table B.2 below highlights the Simplified Zip Code Table for the 90006 zip code.

Table B.2: Simplified Zip Code Table

Zip Code	City/Neighborhood	Courthouse
90006	KOREATOWN-LOS ANGELES	Stanley Mosk Courthouse
90007	LOS ANGELES	Stanley Mosk Courthouse
90008	BALDWIN HILLS	Stanley Mosk Courthouse
90008	CRENSHAW	Stanley Mosk Courthouse
90008	LIEMERT PARK	Stanley Mosk Courthouse
90008	LOS ANGELES	Santa Monica Courthouse

Using the Simplified Zip Code Table and the official LA County Zip Codes shapefiles, we are able to unambiguously determine where the vast majority of eviction cases in Los Angeles County must be filed. For the remaining zip codes where eviction cases are assigned by neighborhood or location relative to a street to different courthouses, we determine the contours of particular neighborhood-zip code pairs using official shapefiles from the LA County eGIS Program.

We then checked the map against eviction records. If a neighborhood-zip code pair had a majority of filed cases with observed courthouses that did not match the courthouse we specified originally for the region, then we corrected that region's filing courthouse. Even after incorporating this correction into the courthouse filing map, approximately 2% of cases have observed filing courthouses that do not match

the "true" or "correct" filing courthouse for that region. See Appendix B.8 for misfiling discussion and analysis.

B.2 Docket Records (Scraping the LASC Portal, Sampling Rates, Geocoding Addresses)

Appendix B.2 describes how the docket record data was obtained. We obtained docket records from scraping the LA Superior Court public case portal. Court filing number is determined by four pieces of information in a pre-specified format: year, court location, case type, and case number. Therefore, we are able to scrape every number up to the total number of records, which we know from FOIA requests. In general, we added a search buffer in case the FOIA volume data was slightly off. Because the records themselves have dates, we check that the final docket we find is close to "12-31-YEAR". In unlawful detainer cases, this means that case numbers look like:

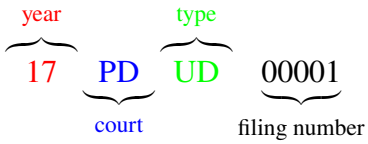

which is a 2017 Pasadena courthouse unlawful detainer (UD) case with filing number 00001.

Over time, the filing formats changed, which are recorded by local court rules. Changes to sealing records rules changed too, effective Jan. 1, 2017. The local rules of court contain information on the filing numbers and sealing laws, which were used to obtain docket records.

Default Evictions Sampling Rates

This section describes sampling rates from the FOIA total filing numbers and the number of defaults we observe. Note that FOIA data collection includes only years 2016–2023 in full, whereas scraped data includes additional cases from later years (2024–2025). Here are sampling rates for years with available FOIA data:

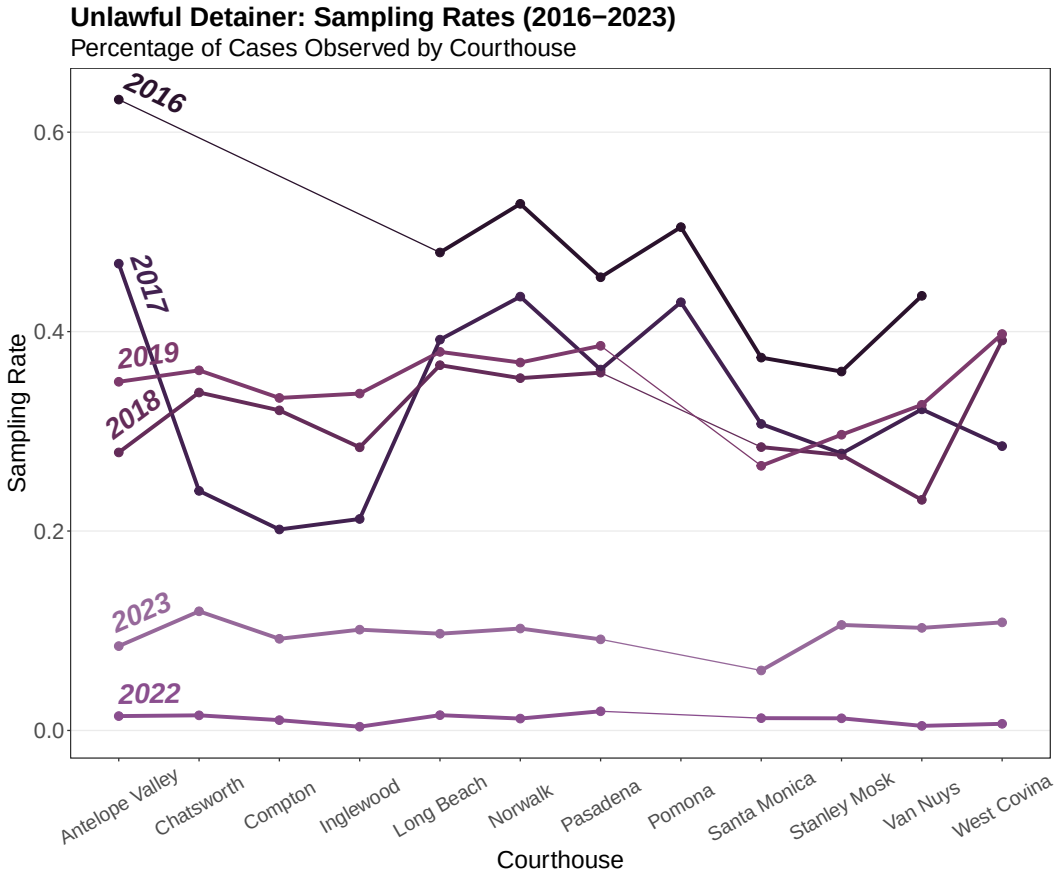

Figure B.1: Sampling Rates using FOIA Data: 2016-2023

Note: The sampling rates are the number of docket records collected divided by the case number reports from FOIA data. The FOIA data are only complete from 2009-2023, with totals also available for Jan 2024–Mar 2024. Note also that the court assignment rule changed in Late 2017, but is unchanged from 2018–2025.

Most of the observed cases are residential default judgments, so the sampling rate is close to the default judgment rate at each courthouse for each year. Note also that the denominators are the higher of either the FOIA number reported to us by LASC or the last scraped case number, in cases of disagreement. Because 2022 case records are largely unavailable (near zero in Appendix Figure B.1), we exclude 2022 from the analyses. Records from the 2020–2021 years were also unavailable to scrape due to different pandemic era rules (e.g., moratoria). However, note that FOIA data from LASC on filing numbers is shown in the top-right panel of Appendix Figure B.13. Filings were much lower in 2020 (13,796 filings) and 2021 (12,646 filings), with an uptick in 2022 (34,398 filings).

The main eviction rules and regulations affecting the 2020–2022 years are briefly

described here. The discussion here is not comprehensive but is meant to illustrate why the 2020–2022 years do not have many (if any) observable docket records for limited residential unlawful detainer actions in LA County.

Local. At the local level, the LA County COVID-19 Tenant Protections Resolution modified eviction rules, which was in effect from March 4, 2020 until March 31, 2023.[1] The Resolution protected tenants from eviction who were unable to pay rent due to the pandemic. Expiration of this LA County Resolution meant that "[n]ormal rent payments must resume for rent due on or after April 1, 2023 to avoid being evicted."[2]

State. At the state level, California passed several acts modifying eviction rules and providing tenants protection who were unable to pay rent due to COVID-19 circumstances, including: California Assembly Bill 3088 (COVID-19 Tenant Relief Act of 2020); California Senate Bill 91, which extended Assembly Bill 3088 protections; and California Assembly Bill 832 (COVID-19 Rental Housing Recovery Act). These acts provided tenant protections from September 2020 to September 2021, with some rental assistance and modified eviction rules lasting until March 2022. The thrust of the acts were to protect those "who have experienced COVID-19-related financial distress from being evicted for failing to make rental payments."[3]

Federal. At the federal level, the CDC's eviction moratorium was in effect from September 2020 until August 2021, when it was struck down by the Supreme Court in *Ala. Ass'n of Realtors v. HHS*, 594 U.S. 758 (2021).

Geocoding the Addresses

The number of records observed at each step of the geocoding process are displayed in Appendix Figure B.2. The percentages at the top of each year grouping give the proportion of eviction records matched in the final step to assessor data from the set of eviction records that contain a nonempty address. In the process of cleaning, subsetting to residential defaults, geocoding, and matching to assessor information

[1] https://dcba.lacounty.gov/noevictions/

[2] https://dcba.lacounty.gov/wp-content/uploads/2024/01/LA-COUNTY-TENANT-PROTECTIONS-RESOLUTION-FAQ-1.24.2024-1.pdf

[3] Per AB 3088, landlords were required to notify "tenants who, as of September 1, 2020" did not pay rent of the following "Notice from the State of California":

> "The California Legislature has enacted the COVID-19 Tenant Relief Act of 2020 which protects renters who have experienced COVID-19-related financial distress from being evicted for failing to make rental payments due between March 1, 2020, and January 31, 2021. 'COVID-19-related financial distress' means..."

we keep more than 70% of the total records.

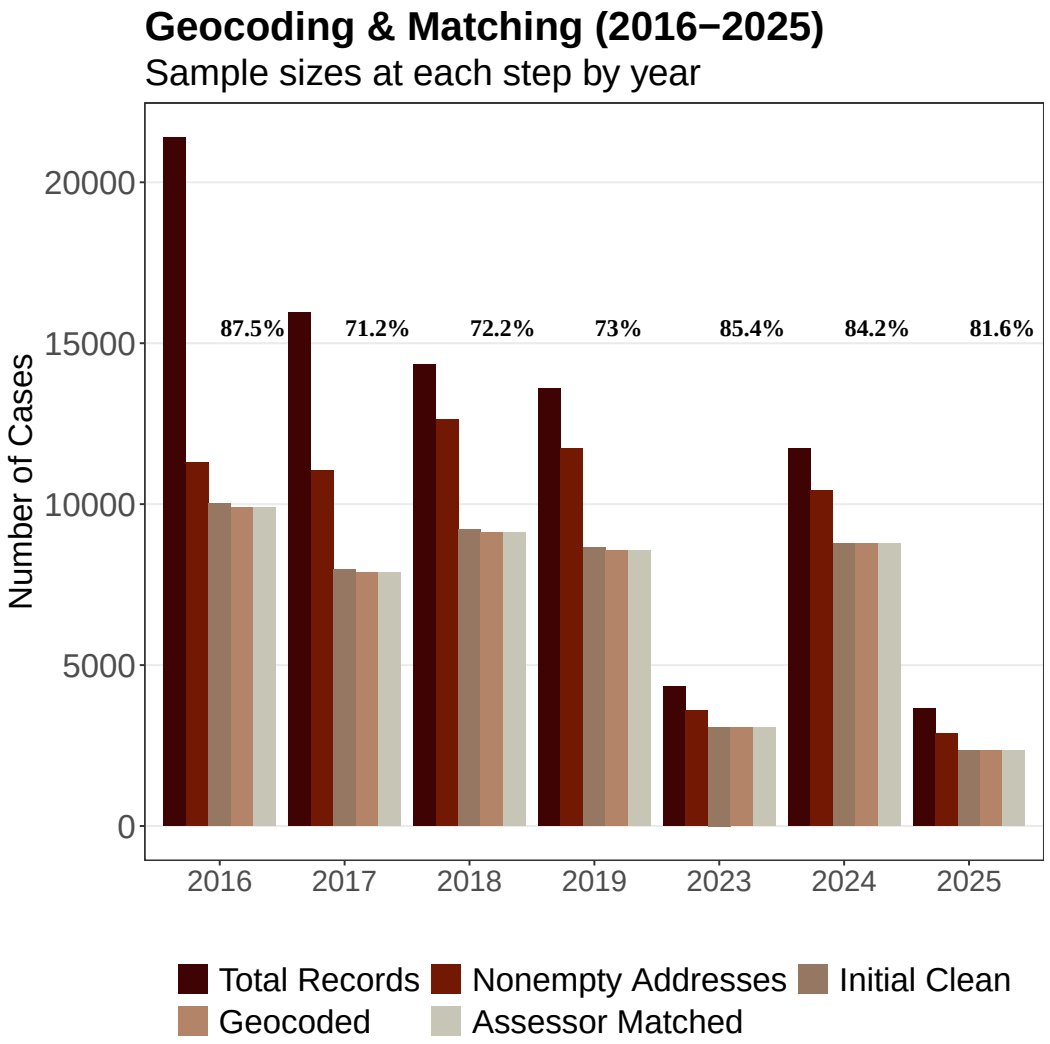

Figure B.2: Defaults Geocoding & Matching Rates

Note: The Total Records are the number of total case docket records collected from LASC. The Nonempty Addresses are the number of records with an ascertainable address from the docket records. The Initial Clean is the number of case records flagged as a "residential" eviction with a "default" outcome. Geocoded numbers are the number of cases with a "High" or "Exact" geocoding match from the Mapbox API. Finally, the Assessor Matched numbers are those records with a match from the LA County Assessor records with non-empty "Number of Units" in the apartment building. The percentages at the top of each year are the number of eviction records matched in the final step over the number of eviction records that contain a nonempty address.

B.3 Time-to-Default Plots

Appendix B.3 contains additional plots documenting how long it takes for the default cases to reach default status. We calculate the number of days to default status using regular expressions on the raw docket data.

First, we extract three dates from the raw case data: the case start-date (measured as the date on the docket the case is assigned to a judge), the earliest possible default entry date (measured as the first appearance of "C.C.P. 415.46" or "C.C.P. 1169" on the docket), and the latest possible default entry date (measured as the last appearance of "C.C.P. 415.46" or "C.C.P. 1169" on the docket).[4] In some (but not many) cases, the last appearance of a default entry differs from the first default entry date. This makes only a small difference in the mean across all years and courthouses (see plots below).

Next, we compute different timing measures. The first set of measures is the time-to-default: the number of days from the assignment of the case to a judge and the first entry of the default judgment on the docket. For each default case j, this is computed as follows:

$$\text{time-to-default}_j = \text{default entered on docket}_j - \text{judge assigned}_j.$$

In some cases, the docket record does not contain the date default judgment is entered. We use early-estimated (E) and late-estimated (L) time-to-default measures, which use the different default entry dates (earliest possible and latest possible).

The second measure calculates the total case time in case j. Using regular expressions, we extract all dates of the format "MM/DD/YYYY." We then compute the difference between the earliest and latest date in case j:

$$\text{total time}_j = \text{latest date on docket}_j - \text{earliest date on docket}_j.$$

This measure is computed for every case in our dataset.

Using these measures, we observe that the time-to-default and total case times are highest in 2023. For example, Appendix Figure B.3 shows the 2023 time-to-default density is shifted further rightward than pre-pandemic years.

[4] To cast a wide net, the regex search is for the dates where either "415.46" or "1169" appear.

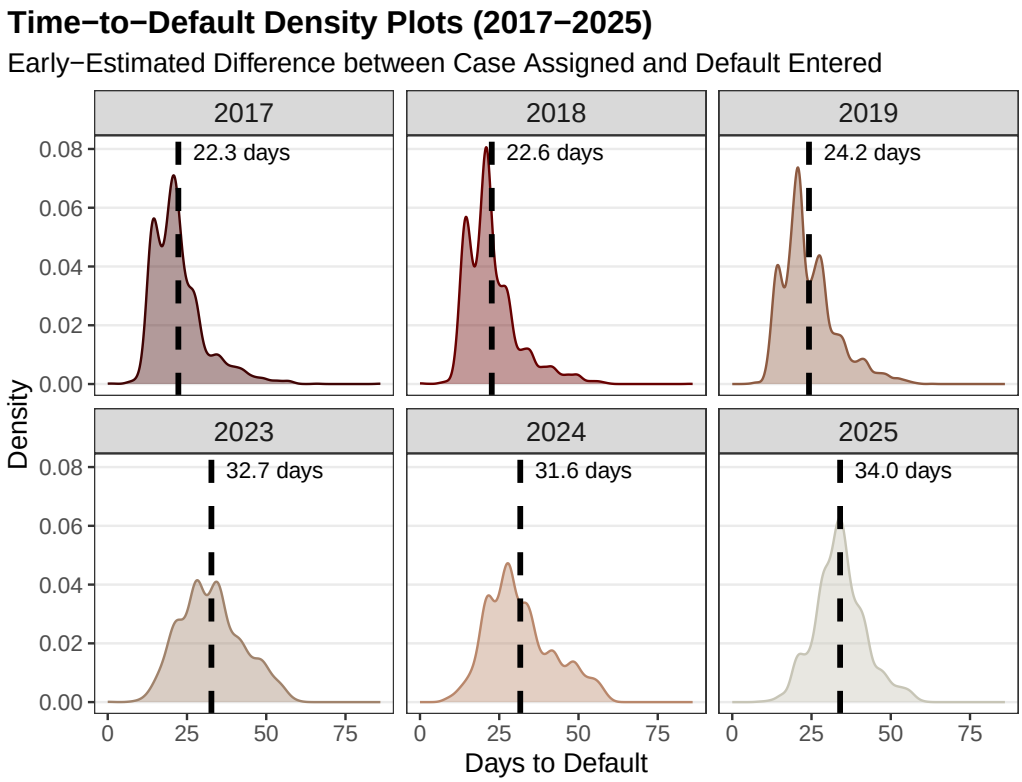

Figure B.3: Time-to-Default: Early-Estimated Default Entered Date

Note: Using only residential default evictions from after Sep. 1, 2017, the mean early-estimated time-to-default is plotted as the dashed black vertical line in each year panel. Densities are plotted for the subset of cases with time-to-default less than or equal to 90 days.

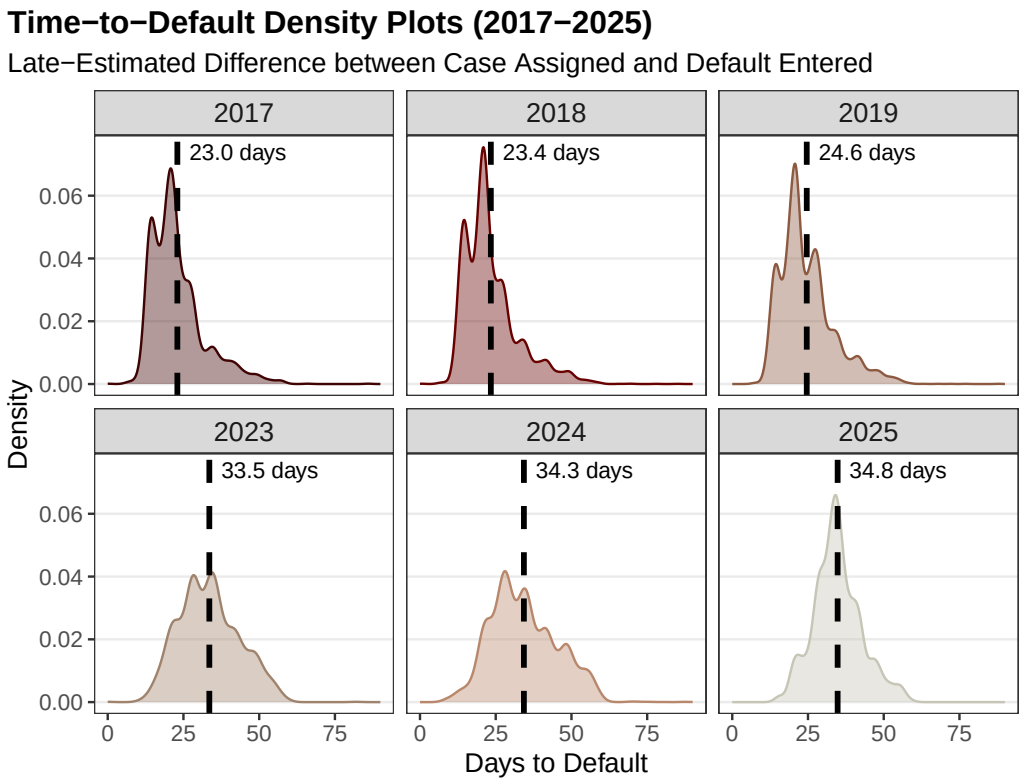

Figure B.4: Time-to-Default: Late-Estimated Default Entered Date

Note: Using only residential default evictions from after Sep. 1, 2017, the mean late-estimated time-to-default is plotted as the dashed black vertical line in each year panel. Densities are plotted for the subset of cases with time-to-default less than or equal to 90 days.

We also record the box plots of the early-estimated time-to-default below for each courthouse pooled across all years: although there is variation across courthouses, the median time-to-default is similar for the 2017-2023 years.

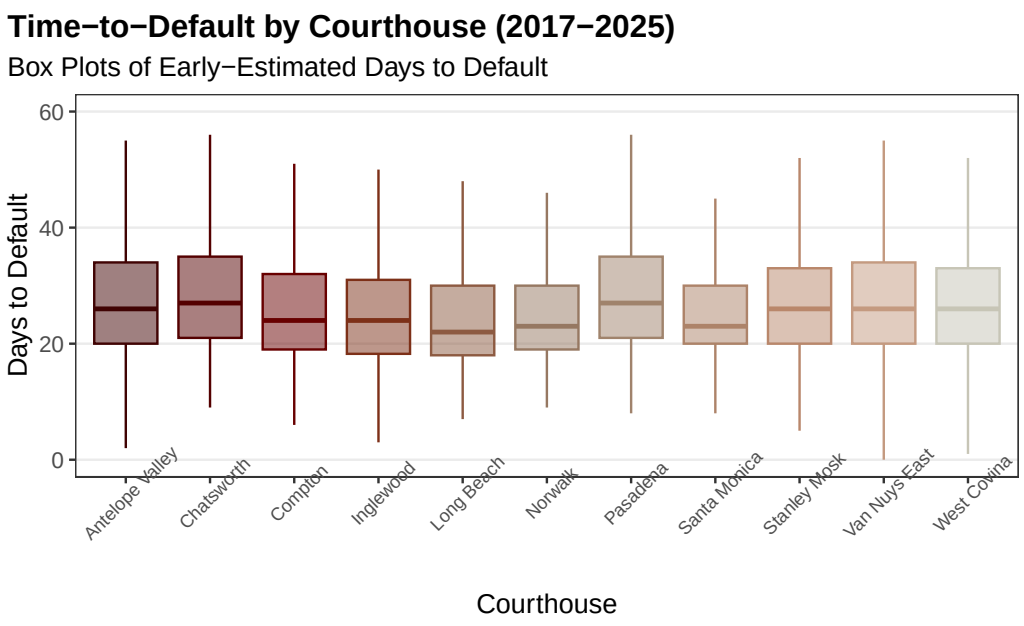

Figure B.5: Time-to-Default: By Courthouses

Note: Analysis is restricted to residential default evictions from after Sep. 1, 2017 here. The data here is also restricted to cases that default in fewer than 60 days and the eleven main eviction courthouses.

Alternatively, we can estimate the start date by the earliest date that appears within the docket. Then we can compute the difference in days between this start date and the estimated default date (early or late). The plots below in Appendix Figure B.6 give densities for estimated time-to-default under this different way to measure the start date: the new mean estimated time-to-default is approximately 28 days (early-estimated default date) or 29 days (late-estimated default date).

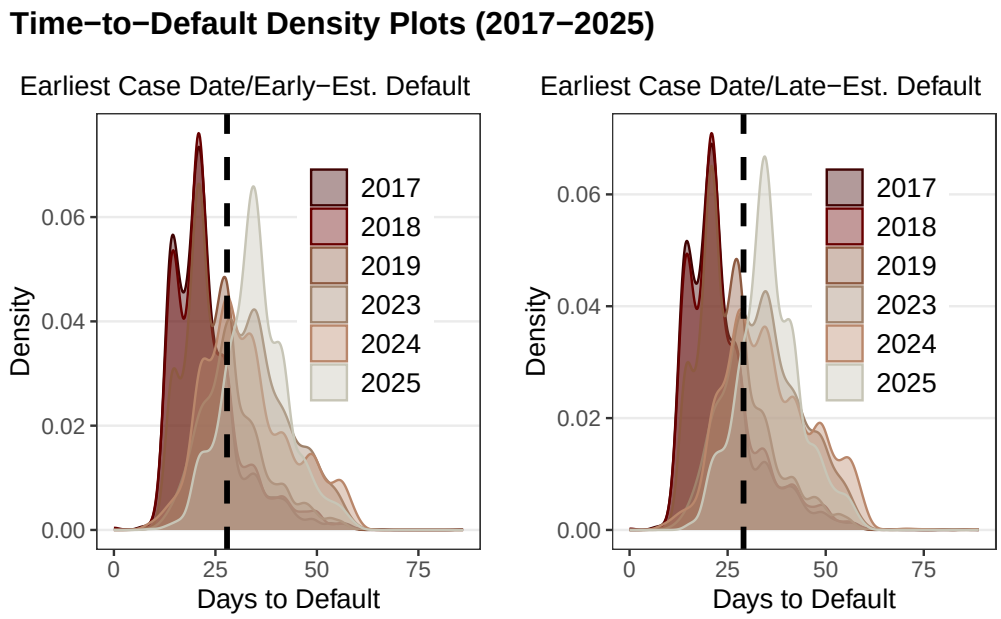

Figure B.6: Time-to-Default: Using Earliest Case Dates

Note: The left panel plots yearly densities for the number of days to default, using the earliest case date and early-estimated default date. The mean time-to-default (27.9 days) is plotted as the dashed black vertical line. The right panel plots yearly densities for the number of days to default, using the earliest case date and late-estimated default date. The mean time-to-default (29.1 days) is plotted as the dashed black vertical line. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

B.4 Total Case Time Plots

Appendix B.4 contains plots on total case times. We extract all dates in the format Month/Day/Year (e.g., "01/01/2018") from the raw docket data and take the earliest and latest dates that appear in the record. The total time is the difference between the latest and earliest dates in the docket data for a given eviction case. The two plots below contain the density plots for the subset of cases that are finished within 1 year or 3 years, respectively.

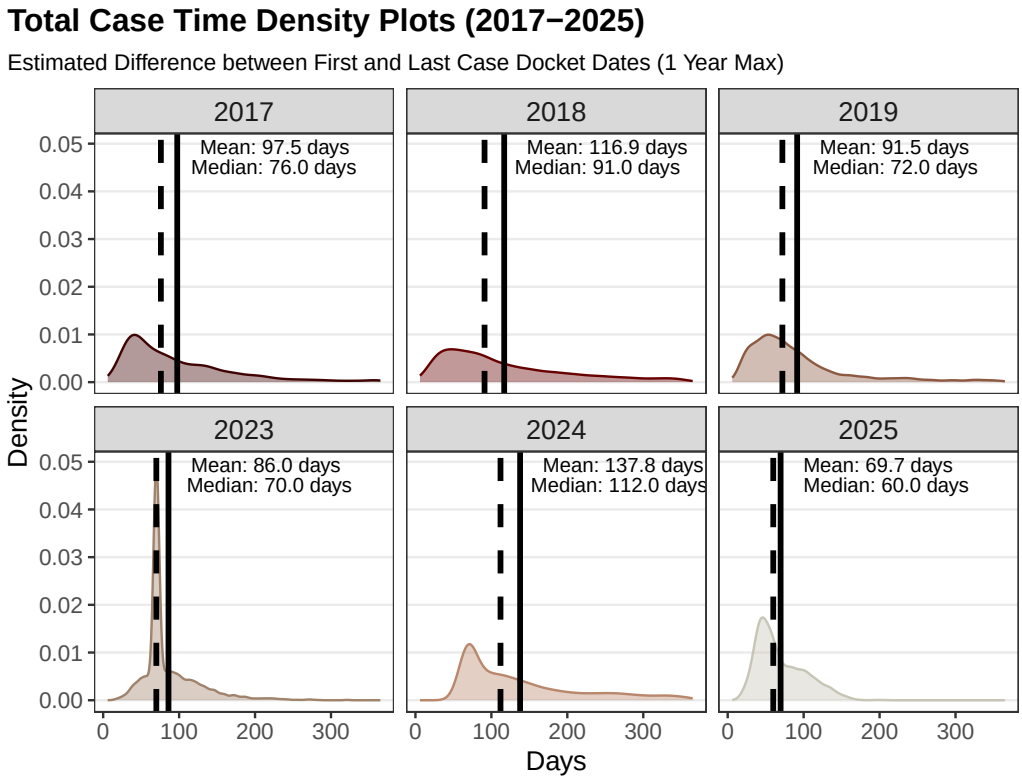

Figure B.7: Total Case Time (< 1 year cases only)

Note: The yearly mean case time and the median case time are plotted as the solid black vertical and dashed lines, respectively. Densities are plotted for the subset of cases with case times less than or equal to one year. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

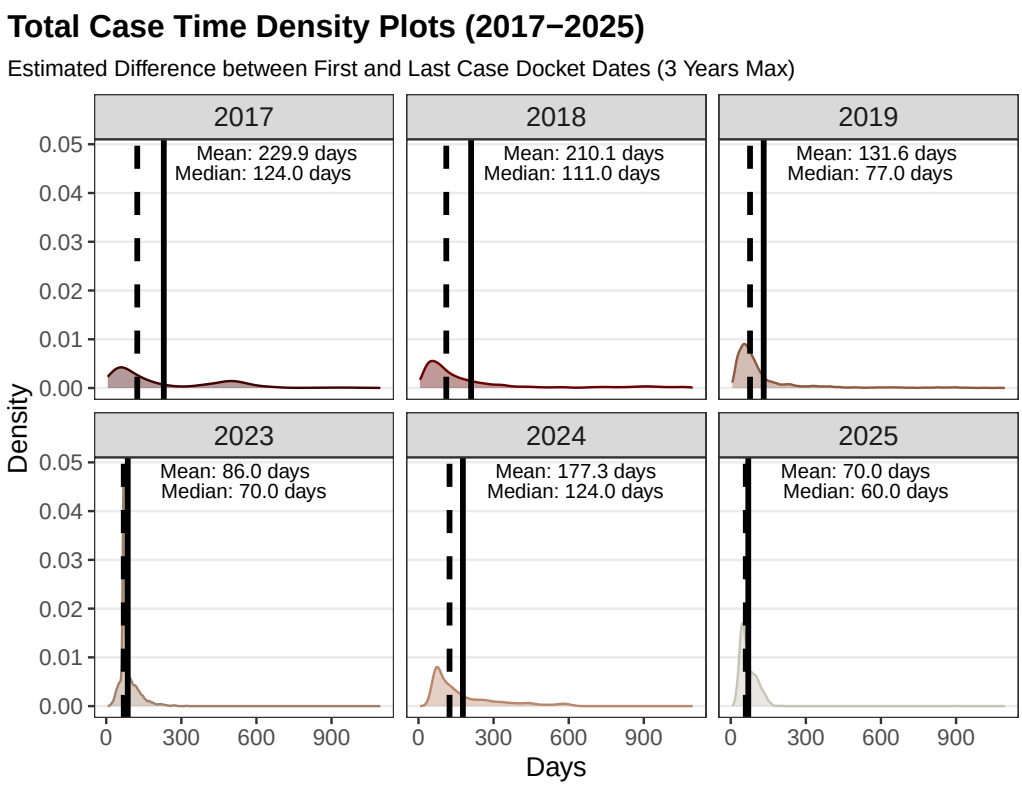

Figure B.8: Total Case Time (< 3 years cases only)

Note: The yearly mean case time and the median case time are plotted as the solid black vertical and dashed lines, respectively. Densities are plotted for the subset of cases with case times less than or equal to one year. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

B.5 Money Judgment Plots

Appendix B.5 contains plots concerning money judgments in the observed eviction cases. We first plot the densities for money judgments across years (pooled courthouses) and density plots for courthouses (years pooled).

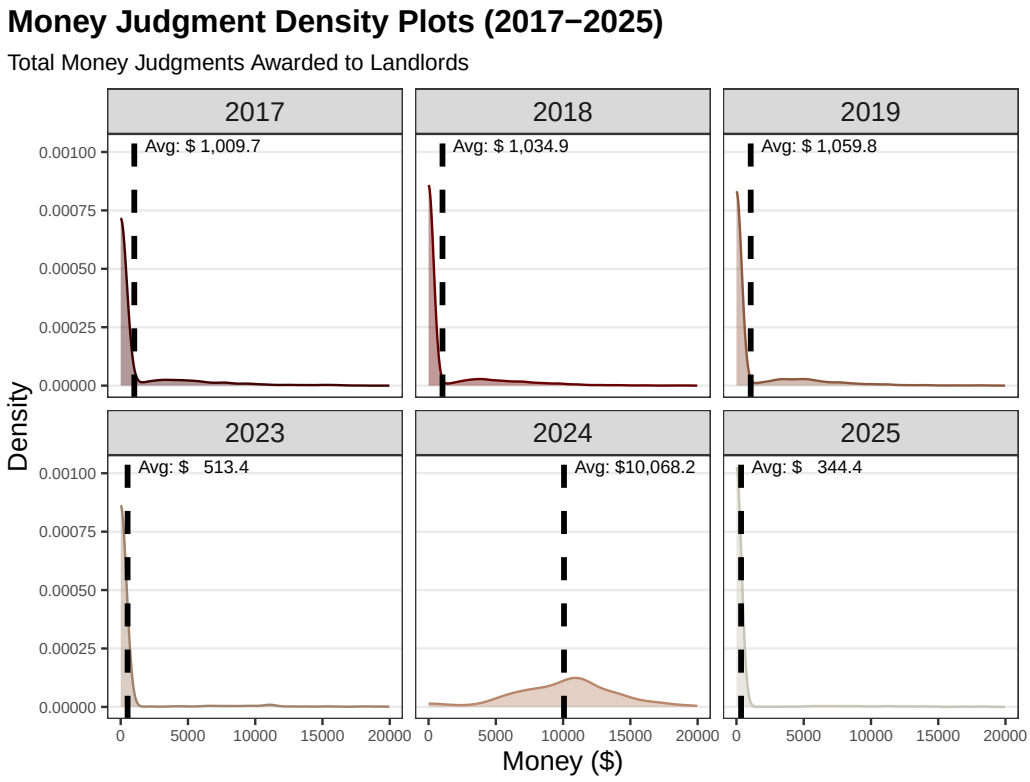

Figure B.9: Yearly Money Judgments

Note: The yearly mean total money judgment is plotted as the dashed black vertical line. Densities are plotted for the subset of cases with money judgments less than $\$20,000$. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

We also plot the densities for each courthouse across all years in the sample for different subsets: cases with total money judgments less than \$5,000, \$10,000, and \$20,000.

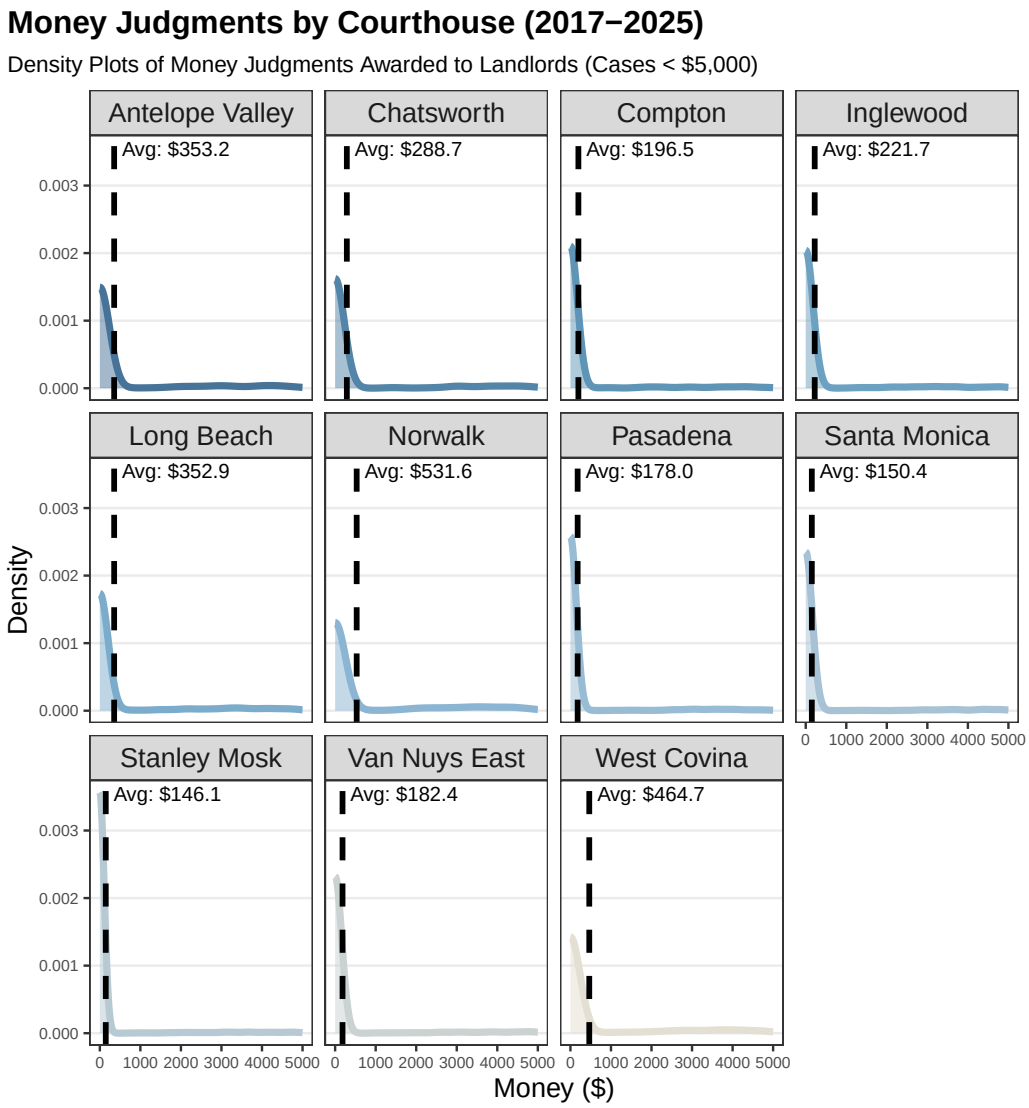

Figure B.10: Money Judgments by Courthouse ($< \$5,000$ cases only)

Note: The black vertical lines show the means for each courthouse. Only the main eleven courthouses are shown and cases with no money judgment information are excluded. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

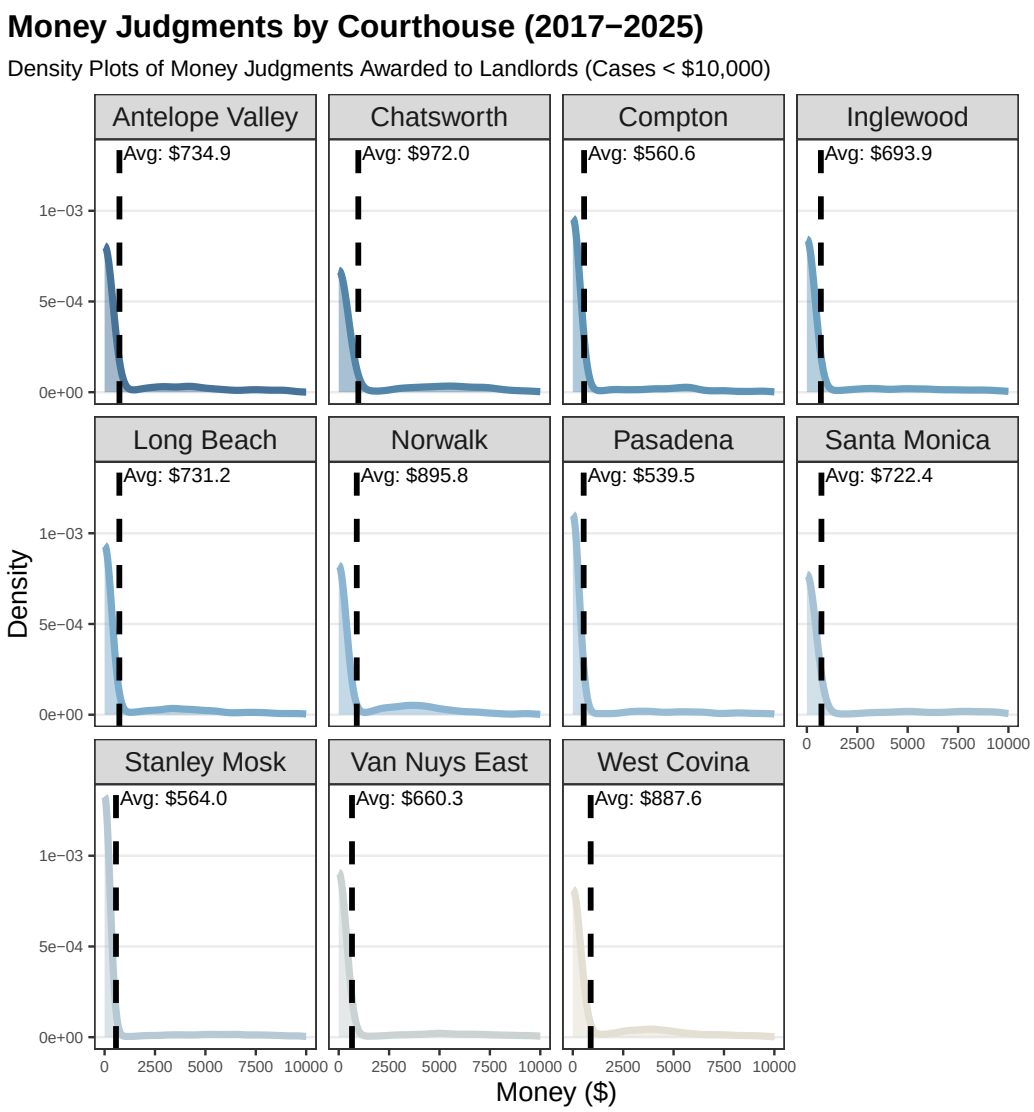

Figure B.11: Money Judgments by Courthouse (< $10,000 cases only)

Note: The black vertical lines show the means for each courthouse. Only the main eleven courthouses are shown and cases with no money judgment information are excluded. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

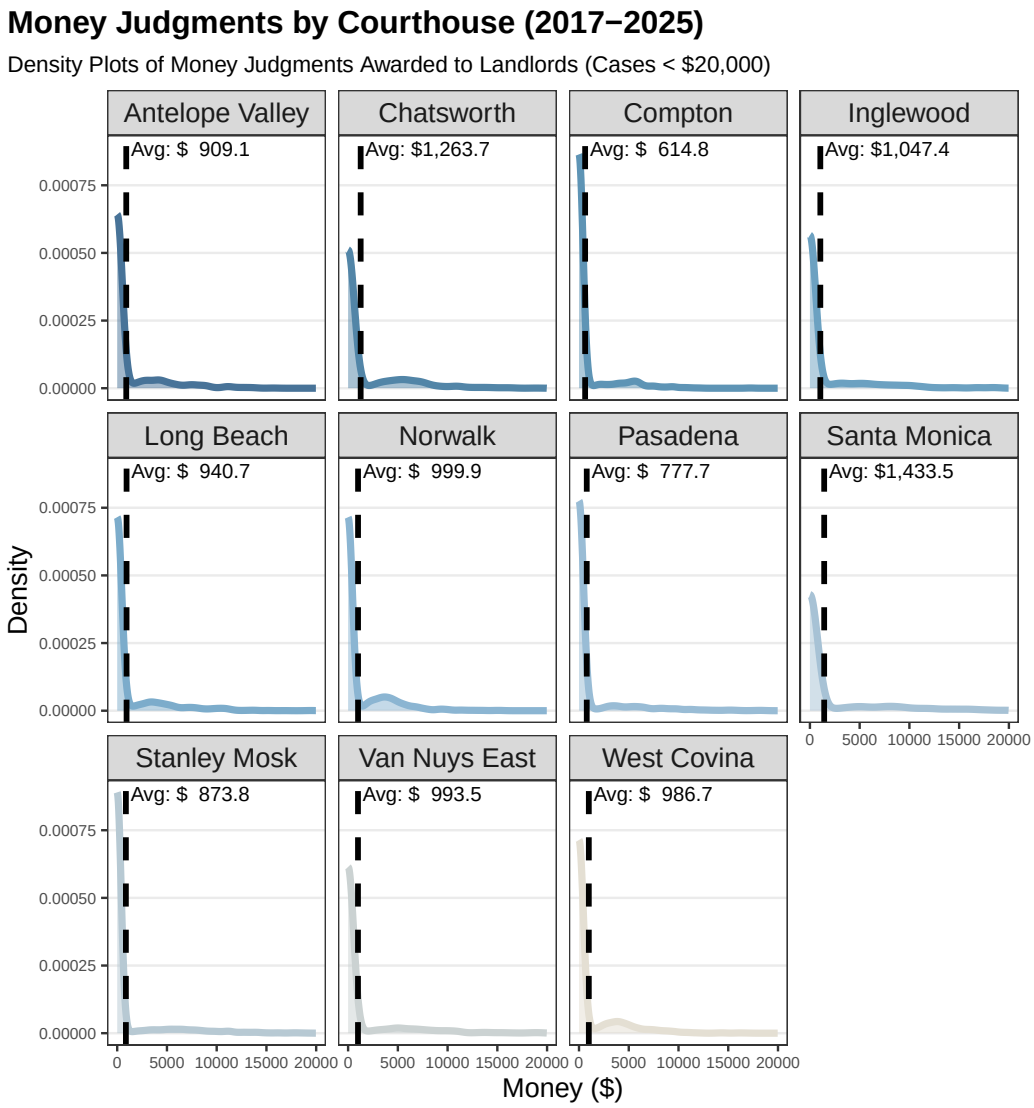

Figure B.12: Money Judgments by Courthouse (< $20,000 cases only)

Note: The black vertical lines show the means for each courthouse. Only the main eleven courthouses are shown and cases with no money judgment information are excluded. Analysis is restricted to residential default evictions from after Sep. 1, 2017 here.

B.6 Data Overview

Appendix B.6 gives an overview of the eviction data (from the Eviction Lab and data we collected) and CA statewide trends.

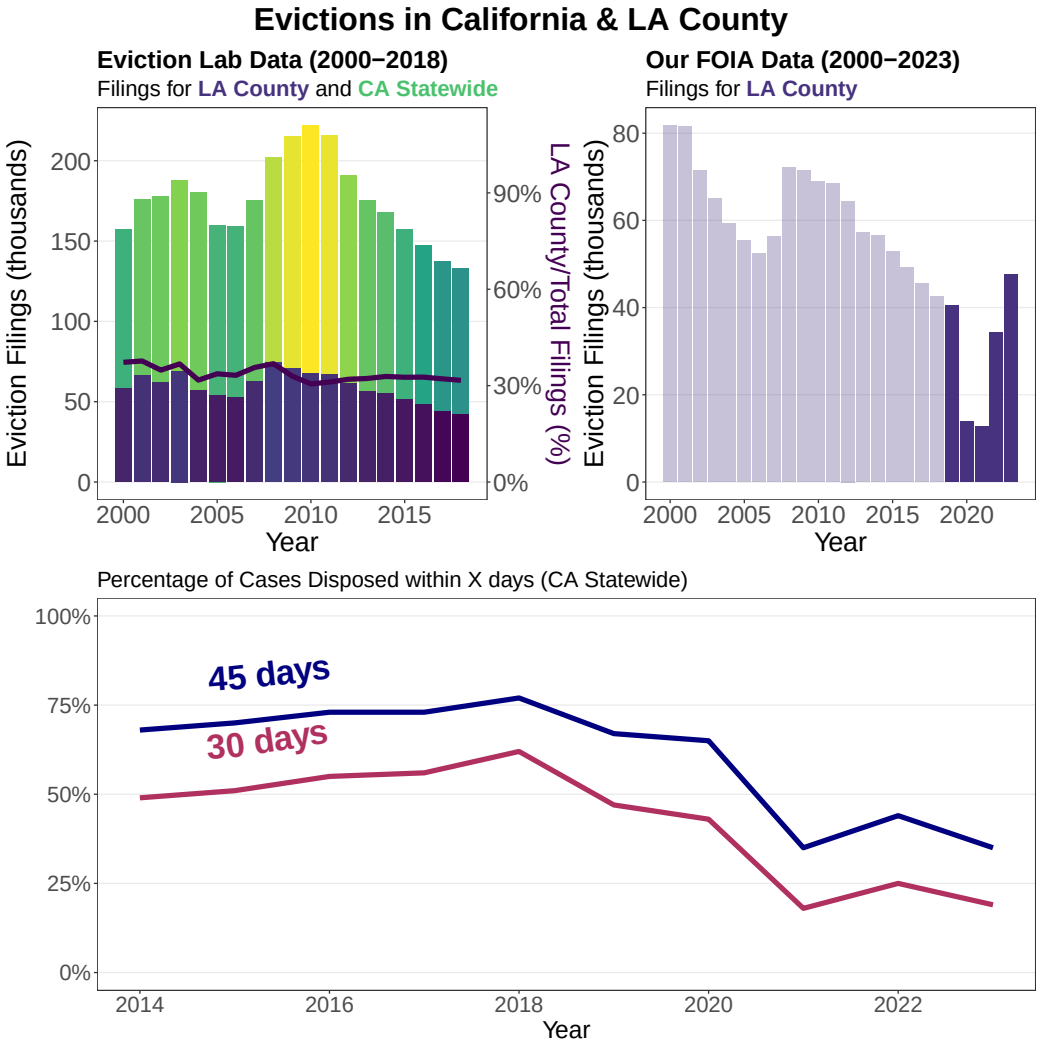

Figure B.13: Eviction Filing Volume (CA & LA County) and Disposition Times (CA)

Note: The top left panel shows the modeled filing estimates from the Eviction Lab (Gromis et al., 2022). The top right panel shows the aggregate filing data we collected for LA County (2018-2023 emphasized). The bottom panel shows the statewide percentage of eviction cases disposed in 45 days (blue) and 30 days (red) from the 2024 Court Statistics Report.

B.7 Procedures in Docket Records

Appendix B.7 describes the docket procedural variables and further court information.

Table B.3: Procedural Events in Docket Records (Collected by Regex)

Category	Procedural Event
Initial Filing & Case Setup	Complaint
	Civil Case Cover Sheet
	Supplemental Unlawful Detainer Cover Sheet
	Summons on Complaint
	Notice of Case Assignment (Limited Civil)
	Property Owner / Landlord Only Hearing Notice
	Online Dispute Resolution Notice
Service of Process	Proof of Personal Service
	Proof of Substituted Service
	Proof of Service by Mail
	Notice of Unlawful Detainer Mailed
	Show Cause for Failure to File Proof of Service
Pleadings & Motions	Answer
	Demand for Jury Trial
	Request or Counter-Request to Set Case for Trial
	Order on Court Fee Waiver
Trial Preparation	Notice of Unlawful Detainer Trial
	Trial Brief
	Jury Instructions
	Witness List
	Exhibit List
	Statement of the Case
	Meet-and-Confer on Joint Trial Documents
Hearings & Other Orders	Jury Trial Scheduled
	Minute Order
	Certificate of Mailing (Jury Trial)
	Notice of Ruling
	Ex Parte App. to Stay Execution of Judgment
	Ex Parte Hearing
	Motion to Set Aside / Vacate Judgment
	Hearing on Motion to Set Aside / Vacate Judgment
Judgments & Defaults	Judgment – Unlawful Detainer
	Notice of Entry of Judgment
	Court Orders Judgment Entered
	Request for Entry of Default / Judgment
	Default Entered
Enforcement & Writs	Application for Issuance of Writ of Execution
	Writ of Possession
	Writ Return

Procedures by Courthouse

The two tables that follow for procedural averages by courthouse are broken up into two tables for viewability.

	Antelope	Chatsworth	Compton	Inglewood	Long Bch	Norwalk
complaint	1.000	1.000	1.000	1.000	1.000	1.000
civil_case_cover_sheet	1.000	1.000	0.999	1.000	0.999	1.000
supplemental_ud_cover_sheet	0.350	0.385	0.314	0.341	0.352	0.302
summons_on_complaint	0.998	1.000	1.000	1.000	0.999	0.999
notice_case_assignment_limited	1.000	1.000	0.911	1.000	1.000	1.000
landlord_only_hearing_notice	1.000	1.000	1.000	1.000	1.000	1.000
odr_ud_notice	0.413	0.390	0.340	0.357	0.359	0.333
pos_personal	0.561	0.497	0.671	0.619	0.611	0.689
pos_substituted	0.705	0.695	0.716	0.777	0.721	0.769
pos_mail	0.022	0.017	0.053	0.024	0.028	0.017
notice_ud_mailed	0.999	1.000	1.000	0.999	1.000	1.000
osc_fail_pos	1.000	1.000	1.000	1.000	1.000	1.000
answer	0.019	0.030	0.035	0.019	0.022	0.009
demand_jury	0.002	0.004	0.003	0.001	0.001	0.001
request_set_for_trial	0.006	0.009	0.010	0.046	0.005	0.027
fee_waiver_order	0.122	0.113	0.255	0.147	0.101	0.141
notice_ud_trial	0.004	0.009	0.008	0.009	0.005	0.006
trial_brief	0.001	0.001	0.000	0.000	0.000	0.000
jury_instructions	0.000	0.000	0.000	0.000	0.000	0.000
witness_list	0.002	0.001	0.001	0.000	0.000	0.000
exhibit_list	0.002	0.001	0.001	0.001	0.000	0.000
statement_of_case	0.000	0.000	0.001	0.000	0.000	0.000
decl_meet_confer_joint_docs	0.000	0.000	0.000	0.000	0.000	0.000
jury_trial_scheduled	0.006	0.013	0.013	0.015	0.007	0.013
minute_order	0.116	0.103	0.155	0.121	0.083	0.092
cert_of_mailing_jury	0.001	0.000	0.001	0.000	0.000	0.001
notice_of_ruling	0.010	0.027	0.020	0.030	0.018	0.025
ex_parte_stay_exec_app	0.043	0.026	0.081	0.052	0.029	0.023
ex_parte_hearing	0.035	0.029	0.043	0.010	0.021	0.034
motion_set_aside_vacate	0.026	0.015	0.011	0.017	0.011	0.009
hearing_motion_set_aside	0.026	0.016	0.012	0.016	0.011	0.010
judgment_ud	0.952	0.975	0.958	0.969	0.971	0.970
notice_entry_judgment	0.060	0.019	0.046	0.016	0.010	0.060
court_orders_judgment_entered	0.000	0.001	0.000	0.001	0.000	0.000
request_entry_default	1.000	0.999	0.999	0.998	1.000	0.998
default_entered	0.999	0.999	0.997	0.987	0.998	0.998
app_writ_exec_pos_sale	0.169	0.218	0.265	0.180	0.196	0.187
writ_possession	0.858	0.938	0.861	0.833	0.913	0.915
writ_return	0.806	0.707	0.532	0.555	0.778	0.679

Table B.4: Procedure Averages by Courthouse

Note: Columns give procedure averages by courthouse. For rows with variation, the last column is a likelihood ratio test p-value from logistic regressions of procedures on courthouses, whose null is procedure means in each row are equal (significant p-values at 0.05 level shown in red).

Pasadena	Santa Mon.	Mosk	Van Nuys E.	W. Covina	p-value	p-value (controls)
1.000	1.000	1.000	1.000	1.000		
0.999	0.999	1.000	0.999	0.999	0.039	0.060
0.441	0.255	0.469	0.328	0.327	<0.001	<0.001
1.000	1.000	0.999	1.000	1.000	0.060	0.101
1.000	1.000	1.000	1.000	1.000	<0.001	<0.001
1.000	1.000	1.000	1.000	1.000		
0.463	0.259	0.480	0.328	0.344	<0.001	<0.001
0.517	0.503	0.567	0.552	0.643	<0.001	<0.001
0.673	0.621	0.724	0.696	0.700	<0.001	<0.001
0.026	0.054	0.033	0.044	0.012	<0.001	<0.001
1.000	1.000	1.000	1.000	1.000	0.513	0.533
1.000	0.999	1.000	1.000	1.000	0.550	0.623
0.019	0.023	0.037	0.037	0.007	<0.001	<0.001
0.002	0.001	0.009	0.004	0.000	<0.001	<0.001
0.005	0.117	0.074	0.093	0.003	<0.001	<0.001
0.121	0.088	0.142	0.103	0.098	<0.001	<0.001
0.005	0.010	0.005	0.006	0.003	0.102	0.099
0.000	0.001	0.000	0.000	0.000	0.441	0.951
0.000	0.001	0.000	0.000	0.000	0.486	0.523
0.000	0.001	0.001	0.000	0.000	0.199	0.280
0.000	0.002	0.001	0.000	0.000	0.272	0.276
0.000	0.001	0.000	0.001	0.001	0.234	0.200
0.000	0.000	0.000	0.000	0.000		
0.006	0.011	0.007	0.007	0.004	<0.001	<0.001
0.116	0.127	0.109	0.086	0.081	<0.001	<0.001
0.000	0.000	0.000	0.000	0.000	0.113	0.113
0.016	0.039	0.038	0.033	0.017	<0.001	<0.001
0.039	0.024	0.027	0.022	0.021	<0.001	<0.001
0.042	0.023	0.031	0.031	0.024	<0.001	<0.001
0.010	0.014	0.031	0.021	0.017	<0.001	<0.001
0.011	0.014	0.031	0.023	0.016	<0.001	<0.001
0.980	0.974	0.978	0.976	0.968	<0.001	<0.001
0.259	0.018	0.051	0.097	0.010	<0.001	<0.001
0.000	0.000	0.000	0.000	0.000	0.755	0.792
1.000	0.998	1.000	1.000	0.999	0.145	0.206
0.996	0.997	1.000	0.999	0.997	<0.001	<0.001
0.266	0.164	0.206	0.156	0.192	<0.001	<0.001
0.821	0.807	0.900	0.823	0.912	<0.001	<0.001
0.701	0.404	0.630	0.664	0.831	<0.001	<0.001

Table B.5: Procedure Averages by Courthouse (cont.)

Note: Columns give procedure averages by courthouse. For rows with variation, the last column is a likelihood ratio test p-value from logistic regressions of procedures on courthouses, whose null is procedure means in each row are equal (significant p-values at 0.05 level shown in red).

Hearing Start Times

Per the LASC Fourth Amended Standing Order, hearing times in the 11 courthouses are non-uniform. Unfortunately, however, the available docket records do not contain information on the precise times when cases are heard, or the case order on a given day. Nevertheless, per the Order the hearing start times for eviction proceedings are as follows: Antelope Valley on Monday and Wednesday at 8:30am; Chatsworth on Mondays and Wednesdays at 8:30am; Compton on Mondays and Wednesdays at 8:30am; Inglewood on Mondays to Fridays at 1:30pm; Long Beach on Mondays and Wednesdays at 1:30pm; Norwalk on Mondays to Fridays at 9:30am; Pasadena on Mondays to Fridays at 1:30pm; Santa Monica on Mondays and Wednesdays at 8:30am, 10:00am, or 3:00pm; Van Nuys East on Fridays at 8:30am and 10:00am; West Covina on Mondays to Fridays at 8:30am; and Stanley Mosk on Mondays at 1:30pm, Tuesdays at 1:30pm, Wednesdays at 8:30am, or Thursdays at 8:30am. Hearing times for ex parte applications differ and are generally heard Mondays through Fridays. Because of the uncertainty in hearing times, it is important to note that commute-to-court time estimates may suffer from measurement error, which is one reason to prefer using distance-to-court metrics as a proxy for the tenant cost of getting to court. However, all these metrics are highly correlated and, in companion work, differences from using commute-to-court times versus distance-to-court are minimal. See Estes (2025) for further details.

Parking Fees

Here, we describe the parking fees we assume by courthouse for the generalized cost measure, which we use as an alternative to distance-to-court as a measure of estimating tenant costs of getting to court. The costs for the eleven courthouses used in computing our metric are as follows: Antelope Valley Courthouse is free (lot near courthouse);[5] Chatsworth Courthouse is $7 (9375 Penfield Avenue, Los Angeles, CA 91311);[6] Compton Courthouse is $7 (422 South Acacia Avenue, Compton, CA 90220);[7] Inglewood Courthouse is $8 (1 East Regent Street, Inglewood, CA 90301);[8] Long Beach Courthouse is $16.49 (101 S. Magnolia Ave Long Beach, CA 90802);[9]

[5] https://www.lacourt.ca.gov/courthouse/info/atp

[6] https://abmparking.com/facilities/chatsworth-court-lot/

[7] https://www.lacourt.ca.gov/courthouse/info/com and https://parkchirp.com/facilities/compton-courthouse-garage/

[8] https://www.lacourt.ca.gov/courthouse/info/ing and https://abmparking.com/facilities/inglewood-courthouse-parking/

[9] https://www.lacourt.ca.gov/courthouse/info/lb and https://en.parkopedia.com/parking/garage/superior_court_of_california/90802/long_beach/?arriving=202604041500&leaving=

Norwalk Courthouse is $16 (Norwalk Civic Center Avenida Manuel Salinas Los Angeles Norwalk, CA 90650);[10] Pasadena Courthouse is $10 (240 Ramona Street, Pasadena, CA 91107);[11] Santa Monica Courthouse is $7 (333 Civic Center Dr Santa Monica, 90401);[12] Van Nuys East Courthouse is $10 (6170 Sylmar Avenue, Los Angeles, CA 91401);[13] West Covina Courthouse is free (lot near courthouse);[14] and Stanley Mosk Courthouse is $20 at several lots near the courthouse.[15]

Procedural PCA Histograms

Procedure index histograms are shown below using different sets of procedures. The captions in each plot detail the procedures used to construct the index and the largest component procedures of each index.

202604041700

[10] https://www.lacourt.ca.gov/courthouse/info/se and https://en.parkopedia.com/parking/garage/norwalk_civic_center/90650/norwalk/

[11] https://www.lacourt.ca.gov/courthouse/info/pas and https://abmparking.com/facilities/pasadena-court-garage/. Since hearing start times at Pasadena are 1:30pm, we assume the pre-2pm $10 rate applies instead of the post-2pm $5 rate.

[12] https://www.lacourt.ca.gov/courthouse/info/sm and https://www.santamonica.gov/places/parking-lots/civic-center-parking-structure. Since we assume 3.5 hours spent at the courthouse, this corresponds to a $7 rate. Note that the daily maximum rate is $14.

[13] https://www.lacourt.ca.gov/courthouse/info/nw and https://abmparking.com/facilities/van-nuys-court/. Since hearing start times at Van Nuys East are in the mornings, we assume the morning rate (6am–2pm) of $10 applies.

[14] https://www.lacourt.ca.gov/courthouse/info/cit

[15] https://www.lacourt.ca.gov/courthouse/info/la. Per Parkopedia, there are several lots near the Stanley Mosk Courthouse with $20 rates, including: Grand Park (PCAM Lot #10), 145 N Broadway, Downtown, Los Angeles, CA 90012; Music Center Garage 111 S Grand Avenue, Downtown, Los Angeles, CA 90012; and The Music Center, 135 N Grand Ave, Downtown, Los Angeles, CA 90012. The first lot is near the LA County Law Library, which is across the street from the Stanley Mosk Courthouse.

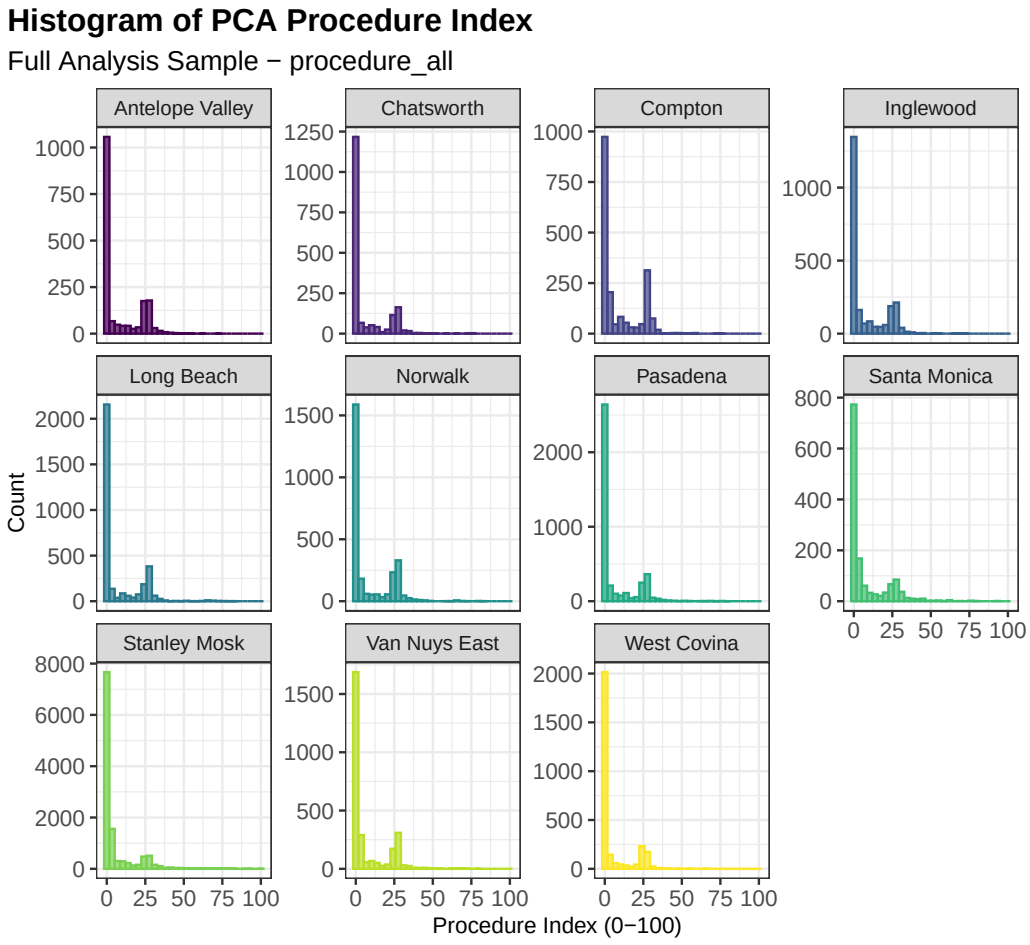

Figure B.14: PCA Procedure Index Histograms by Court (procedure all)

Note: Procedure index (standardized first principal component values) shown in histograms for each courthouse. This PCA procedure index is constructed using the following binary procedure variables: complaint; civil_case_cover_sheet; supplemental_ud_cover_sheet; summons_on_complaint; notice_case_assignment_limited; landlord_only_hearing_notice; odr_ud_notice; pos_personal; pos_substituted; pos_mail; notice_ud_mailed; osc_fail_pos; answer; demand_jury; request_set_for_trial; fee_waiver_order; notice_ud_trial; trial_brief; jury_instructions; witness_list; exhibit_list; statement_of_case; decl_meet_confer_joint_docs; jury_trial_scheduled; minute_order; cert_of_mailing_jury; notice_of_ruling; ex_parte_stay_exec_app; ex_parte_hearing; motion_set_aside_vacate; hearing_motion_set_aside. The five variables with highest weight in this procedure index specification are: minute_order (0.0875), fee_waiver_order (0.0804), motion_set_aside_vacate (0.0802), hearing_motion_set_aside (0.0802), and ex_parte_hearing (0.0781). Variance in these procedures explained by first principal component is 0.1478.

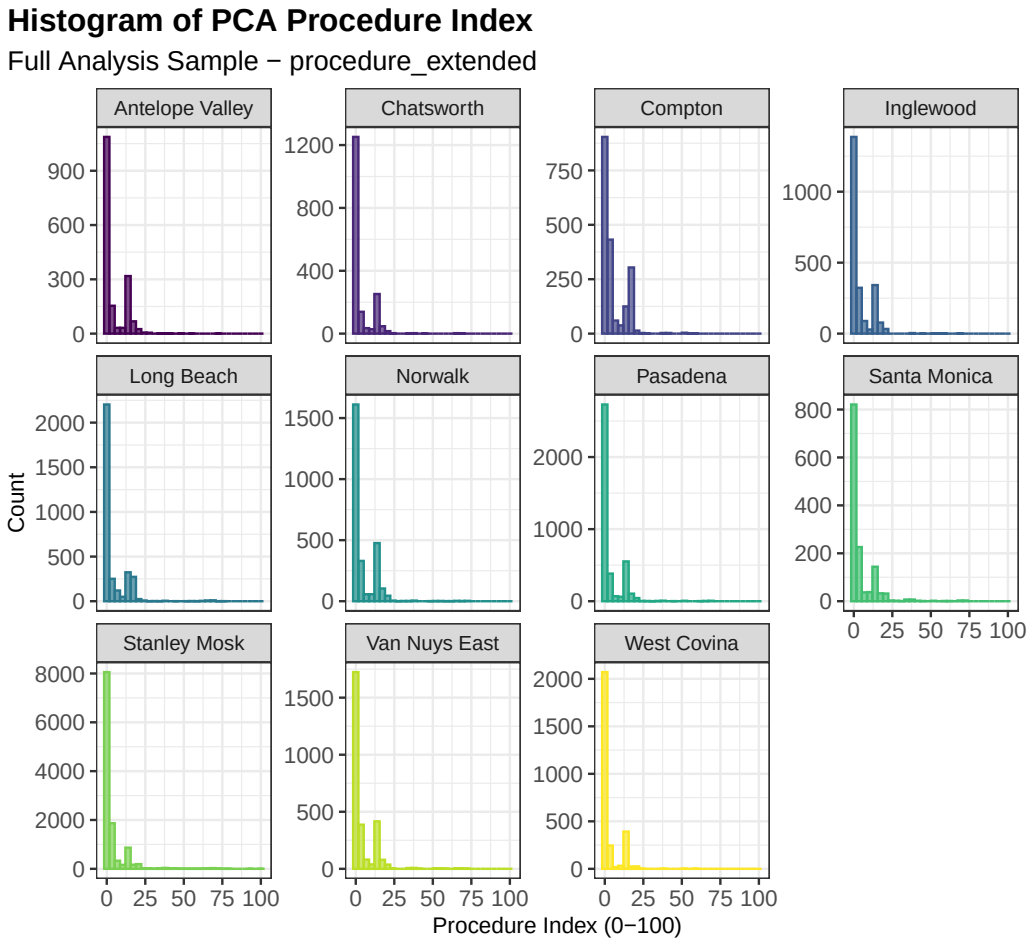

Figure B.15: PCA Procedure Index Histograms by Court (procedure extended)

Note: Procedure index (standardized first principal component values) shown in histograms for each courthouse. This PCA procedure index is constructed using the following binary procedure variables: complaint; civil_case_cover_sheet; supplemental_ud_cover_sheet; summons_on_complaint; notice_case_assignment_limited; landlord_only_hearing_notice; odr_ud_notice; pos_personal; pos_substituted; pos_mail; notice_ud_mailed; osc_fail_pos; answer; demand_jury; request_set_for_trial; fee_waiver_order; notice_ud_trial; trial_brief; jury_instructions; witness_list; exhibit_list; statement_of_case; decl_meet_confer_joint_docs. The five variables with highest weight in this procedure index specification are: witness_list (0.1678), jury_instructions (0.1567), exhibit_list (0.1447), statement_of_case (0.1341), and demand_jury (0.0708). Variance in these procedures explained by first principal component is 0.1486.

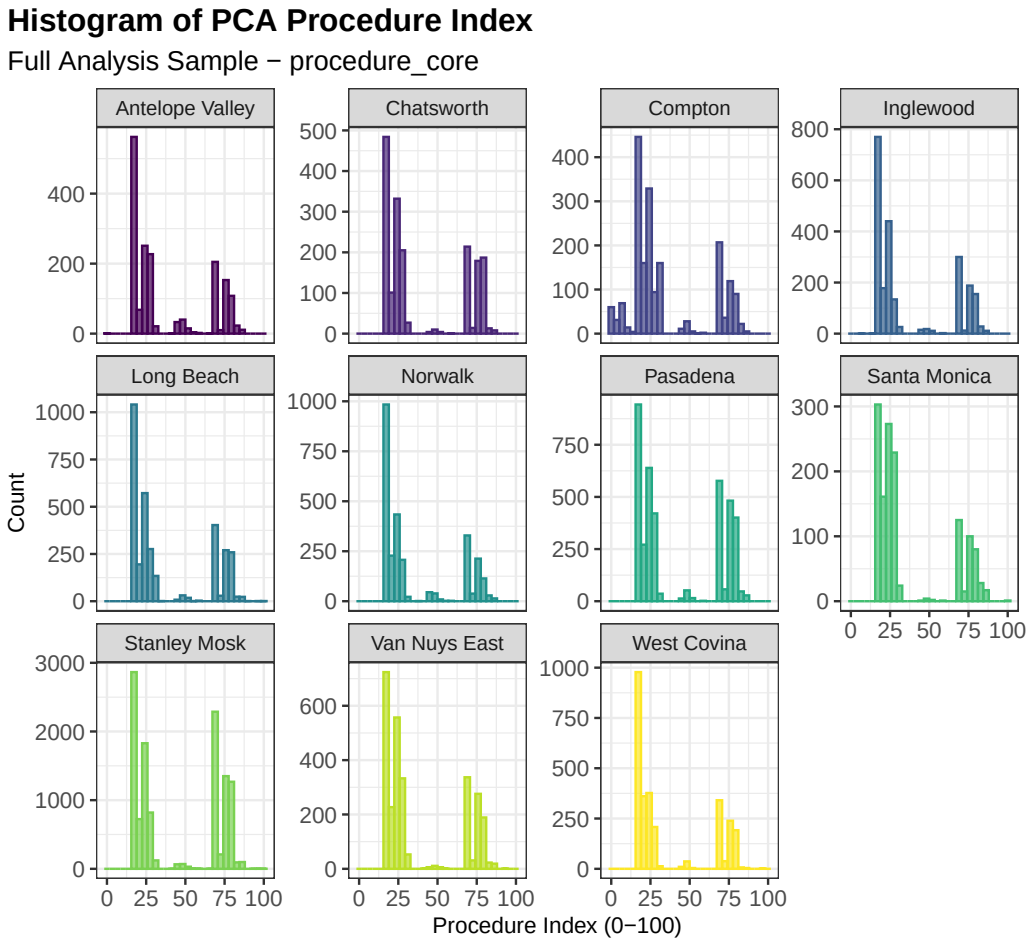

Figure B.16: PCA Procedure Index Histograms by Court (procedure core)

Note: Procedure index (standardized first principal component values) shown in histograms for each courthouse. This PCA procedure index is constructed using the following binary procedure variables: complaint; civil_case_cover_sheet; supplemental_ud_cover_sheet; summons_on_complaint; notice_case_assignment_limited; landlord_only_hearing_notice; odr_ud_notice; pos_personal; pos_substituted; pos_mail; notice_ud_mailed; osc_fail_pos. The five variables with highest weight in this procedure index specification are: supplemental_ud_cover_sheet (0.3760), odr_ud_notice (0.3753), pos_personal (0.1058), pos_ substituted (0.0471), and notice_case_assignment_limited (0.0392). Variance in these procedures explained by first principal component is 0.2010.

B.8 Misfiled Cases

Appendix B.8 discusses possible explanations for the observed phenomenon of misfiled cases. Misfiling is when a docket record's observed courthouse does not match the proper filing courthouse from the LA County Zip Code assignment rule (Table 2.1). Because the misfiling classification is map-dependent, we proceed conservatively. Specifically, we assume that, in each neighborhood-zip pair, the majority of cases are filed in the correct courthouse. Even under this assumption, approximately 2.3% of the eviction cases are misfiled (2018, 2019, 2023). By comparison, the baseline map created (see Appendix B.1) had 5.5% of cases classified as misfiled (2018, 2019, 2023). As discussed in-text, we restrict analyses to correctly filed cases only in the main analyses.

Here we explore three possibilities for misfiling. First, the courthouse assignment map could be wrong.[16] Second, cases might be incorrectly geo-located: the Mapbox API could be returning incorrect lat-long coordinates for certain addresses. Third, landlords may knowingly or unknowingly file cases in the incorrect courthouse.[17]

The first possibility is unlikely in regions with observed eviction cases. Although we cannot fully rule out that the map is incorrect in areas with no observed eviction cases, in regions with observed eviction records it is often not possible to draw a map that clearly separates the observed eviction cases into consistent, continguous filing regions. For example, in the left panel of Appendix Figure B.17 there are eviction cases observed in Hyde Park (Zip 90043) from three different courthouses. There is no contiguous map separating Inglewood and Santa Monica misfiled cases from the correctly filed Stanley Mosk cases because misfiled cases are stacked "on top of" correctly filed cases.

[16] This could occur because neighborhoods and zip codes lack clear, defined geographical boundaries. Indeed, despite using the most current LA County geospatial data, there are some regions in the court filing table that do not correspond to the names given to particular areas of LA County by the assignment rule.

[17] A fourth possibility is that clerks or judges refiled the cases to change the venue. We found no instances of refiling among the cases we label misfiled.

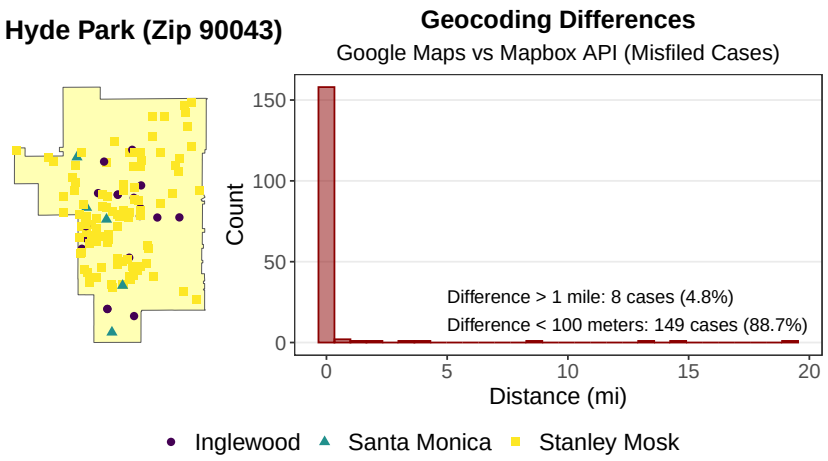

Figure B.17: Misfiling: Zip 90043 Example & Geolocating Check

Note: In the left panel, the background color is the most observed filing courthouse in the region. Each dot is a building with observed default eviction cases, which is colored and shaped by the observed courthouse in the eviction records. In the right panel, we manually checked 10% of the cases we label as misfiled with Google Maps. The histogram of the distance from the Google Maps to the Mapbox API coordinates are shown above.

The second possibility is an error in geo-locating these cases. We manually checked a random subset of the misfiled cases for geo-location errors. We compared the lat-long Mapbox locations with the Google Maps lat-long locations. We assume for this exercise that the Google Maps coordinates are "ground truth," although in some cases there may be errors in Google Maps as well. In the right panel of Appendix Figure B.17, we find that the vast majority (89%) of the geocoded coordinates from Mapbox are very close (< 100 meters) to the Google Maps coordinates. Only eight cases (4.8%) of the manual checks were very clearly wrong, by which we mean that the Google Maps coordinates were more than 1 mile from the location of the Mapbox API coordinates.

Finally, we discuss a third possible explanation for misfiling: landlords file in the incorrect courthouse. Misfiling is not necessarily nefarious, as it could be due to confusion or lack of information. Indeed, in some regions with misfiled cases the assignment rule (Table 2.1) is relatively complicated.[18] Alternatively, landlords may strategically disadvantage tenants by misfiling in distant courthouses. Using the docket records, we test whether landlords strategically disadvantage tenants by filing in more distant courthouses. To test whether landlords are strategic, we compare two distance histograms. In Appendix Figure B.18, the red histogram is the distance

[18]Meaning the assignment rule does not pick out one unique courthouse for the zip code.

from the tenant address to the assigned courthouse, whereas the blue histogram is the distance from the tenant address to the observed filing courthouse.

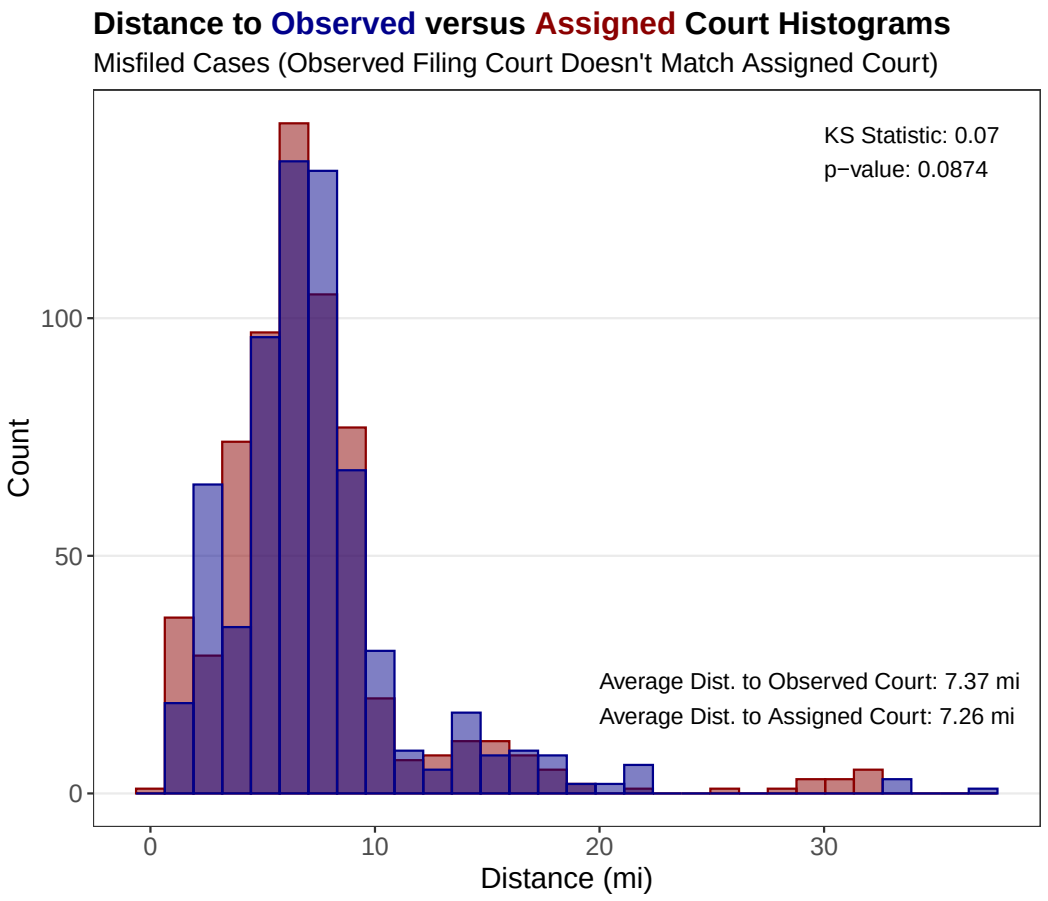

Figure B.18: Misfiled Cases: Distance to Observed vs Assigned Court

Note: The blue histogram is the distance to observed court (among misfiled cases) and the red histogram is the distance to the assigned court (among misfiled cases). The KS statistic can be used to test the null hypothesis that the distances are drawn from the same distribution.

Although the average distance to the observed courthouse is 0.11 miles further for tenants, we cannot statistically distinguish between the two distributions ($p \approx 0.0874$) using the standard KS test. In other words, there is no statistical evidence of landlord strategic filing to disadvantage tenants among the cases we here label misfiled.

It has long been recognized that clear legal rules lower legal costs and reduce uncertainty.[19] An unclear procedural rule like the LA County eviction assignment

[19] See, e.g., Kaplow (1992).

mechanism can be clarified or refashioned. Although the optimal assignment rule is beyond the scope of this paper, we have identified some neighborhood-zip pairs (e.g., Hyde Park in Appendix Figure B.17) where there is, at least in practice, assignment uncertainty.

B.9 Covariate Continuity Tests with `rdrobust`

Appendix B.9 includes covariate continuity checks for several pre-determined covariates across each courthouse pair spatial boundary. Using the pre-determined covariates as the outcome variables, we look for null treatment effects across the boundaries for the following variables: property assessment variables from LA County Assessor datasets (building total value and taxable value); demographic variables (ACS-imputed female, white variables and predicted race/gender variables from docket record names using `wru` package); income variables (ACS-imputed median income and rent); and representation variables (tenant and plaintiff landlord representation rates, from docket records). The `rdrobust` treatment effect estimates on each pre-determined covariate, 95% confidence interval (CI), and robust p-values are recorded for each courthouse pair below.

Table B.6: Chatsworth and Van Nuys East: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	15.85	(9.80, 21.89)	<0.001
Taxable Value (millions)	16.53	(10.45, 22.60)	<0.001
Female	-167.11	(-264.66, -69.56)	<0.001
White	2.13	(-112.53, 116.79)	0.971
Median Income	18686.45	(12731.64, 24641.27)	<0.001
Predicted White	0.01	(-0.06, 0.08)	0.759
Predicted Female	0.13	(-0.04, 0.30)	0.137
Predicted Black	0.02	(-0.02, 0.07)	0.361
Predicted Hispanic	-0.04	(-0.16, 0.08)	0.490
Rent	-45.82	(-179.91, 88.28)	0.503
Tenant Representation	0.02	(-0.03, 0.08)	0.374
Plaintiff Representation	0.05	(-0.01, 0.11)	0.123

Table B.7: Compton and Norwalk: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	-8.85	(-11.91, -5.78)	<0.001
Taxable Value (millions)	3.24	(0.44, 6.04)	0.023
Female	165.37	(73.07, 257.66)	<0.001
White	6.01	(-90.36, 102.38)	0.903
Median Income	4252.52	(21.65, 8483.40)	0.049
Predicted White	0.13	(0.06, 0.20)	<0.001
Predicted Female	-0.02	(-0.22, 0.18)	0.846
Predicted Black	0.11	(0.05, 0.18)	<0.001
Predicted Hispanic	-0.25	(-0.38, -0.12)	<0.001
Rent	414.12	(343.81, 484.44)	<0.001
Tenant Representation	-1.56e-03	(-0.05, 0.05)	0.951
Plaintiff Representation	0.17	(0.05, 0.29)	0.005

Table B.8: Compton and Stanley Mosk: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	-1.54	(-2.70, -0.38)	0.009
Taxable Value (millions)	0.28	(9.77e-03, 0.55)	0.042
Female	-100.14	(-216.01, 15.72)	0.090
White	21.11	(-74.16, 116.38)	0.664
Median Income	-3995.03	(-9827.54, 1837.47)	0.179
Predicted White	0.03	(-0.06, 0.11)	0.558
Predicted Female	0.15	(-0.07, 0.36)	0.180
Predicted Black	0.02	(-0.06, 0.10)	0.618
Predicted Hispanic	-0.07	(-0.22, 0.08)	0.341
Rent	9.35	(-67.25, 85.95)	0.811
Tenant Representation	-0.04	(-0.14, 0.06)	0.477
Plaintiff Representation	-0.31	(-0.46, -0.16)	<0.001

Table B.9: Inglewood and Long Beach: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	1.88	(-1.50, 5.26)	0.276
Taxable Value (millions)	3.89	(0.31, 7.46)	0.033
Female	-133.67	(-378.56, 111.23)	0.285
White	-87.27	(-409.91, 235.38)	0.596
Median Income	11853.22	(-11111.09, 34817.52)	0.312
Predicted White	-0.06	(-0.25, 0.12)	0.498
Predicted Female	-0.33	(-0.65, -7.48e-03)	0.045
Predicted Black	0.05	(-0.06, 0.16)	0.395
Predicted Hispanic	-0.02	(-0.28, 0.23)	0.850
Rent	-283.09	(-555.51, -10.67)	0.042
Tenant Representation	0.03	(-0.08, 0.14)	0.557
Plaintiff Representation	-0.04	(-0.18, 0.10)	0.590

Table B.10: Pasadena and Stanley Mosk: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	54.11	(43.31, 64.90)	<0.001
Taxable Value (millions)	56.11	(45.33, 66.89)	<0.001
Female	-91.89	(-189.05, 5.28)	0.064
White	293.45	(178.03, 408.86)	<0.001
Median Income	10747.94	(4044.52, 17451.35)	0.002
Predicted White	0.07	(-0.01, 0.15)	0.095
Predicted Female	-0.14	(-0.26, -0.02)	0.027
Predicted Black	0.02	(-0.03, 0.06)	0.500
Predicted Hispanic	-0.14	(-0.24, -0.05)	0.003
Rent	1021.71	(939.61, 1103.81)	<0.001
Tenant Representation	0.04	(-0.02, 0.09)	0.196
Plaintiff Representation	0.05	(-0.03, 0.12)	0.200

Table B.11: Pasadena and West Covina: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	-7.08	(-11.85, -2.31)	0.004
Taxable Value (millions)	-5.12	(-9.70, -0.55)	0.028
Female	47.75	(-90.47, 185.97)	0.498
White	105.85	(-98.25, 309.94)	0.309
Median Income	10772.42	(-8098.26, 29643.10)	0.263
Predicted White	-0.05	(-0.15, 0.04)	0.294
Predicted Female	-0.08	(-0.39, 0.23)	0.617
Predicted Black	-8.17e-03	(-0.06, 0.04)	0.737
Predicted Hispanic	-0.04	(-0.25, 0.18)	0.740
Rent	-3.15	(-252.31, 246.01)	0.980
Tenant Representation	0.08	(-0.06, 0.21)	0.259
Plaintiff Representation	0.09	(-0.03, 0.20)	0.147

Table B.12: Santa Monica and Stanley Mosk: Covariate Continuity Checks

Covariate	Robust Estimate	Robust CI	Robust p-value
Total Value (millions)	-8.95	(-12.81, -5.09)	<0.001
Taxable Value (millions)	-9.14	(-14.95, -3.33)	0.002
Female	88.34	(14.48, 162.19)	0.019
White	37.54	(-62.36, 137.44)	0.461
Median Income	-15918.08	(-23611.12, -8225.05)	<0.001
Predicted White	8.26e-03	(-0.05, 0.06)	0.762
Predicted Female	-0.01	(-0.15, 0.12)	0.863
Predicted Black	0.05	(6.05e-03, 0.09)	0.024
Predicted Hispanic	-0.01	(-0.09, 0.06)	0.707
Rent	-198.54	(-329.45, -67.62)	0.003
Tenant Representation	-0.04	(-0.08, 4.30e-03)	0.080
Plaintiff Representation	-0.04	(-0.08, -3.60e-03)	0.032

Finally, the p-values and bandwidths from McCrary tests on the density of the running variable are shown below, when using the baseline specification sample. Note that these tests use all observations near a courthouse pair spatial boundary. Jackknife p-values in right-most column.

Table B.13: McCrary / rddensity Tests by Courthouse Pair

Courthouse Pair	Bandwidth Left	Bandwidth Right	p-value
Chatsworth and Van Nuys East	930.014	802.089	<0.001
Compton and Norwalk	1535.050	1512.981	0.855
Compton and Stanley Mosk	869.820	1078.225	0.049
Inglewood and Long Beach	961.104	1060.890	<0.001
Pasadena and Stanley Mosk	843.842	1138.083	<0.001
Pasadena and West Covina	1550.607	1157.878	<0.001
Santa Monica and Stanley Mosk	597.778	678.735	0.049

B.10 Additional Robustness Checks

Appendix B.10 includes additional robustness checks for the main RDD LATE estimation and the mechanism model.

LATE Estimates

This sub-section provides additional robustness checks on the LATE estimates. First, we assess the sensitivity of the LATE estimates to the bandwidth. The dashed vertical lines in each panel are the rdrobust selected bandwidth, whose results are reported in-text as the "Baseline" specification and are reproduced in the lower-right corner of Appendix Figure B.19.

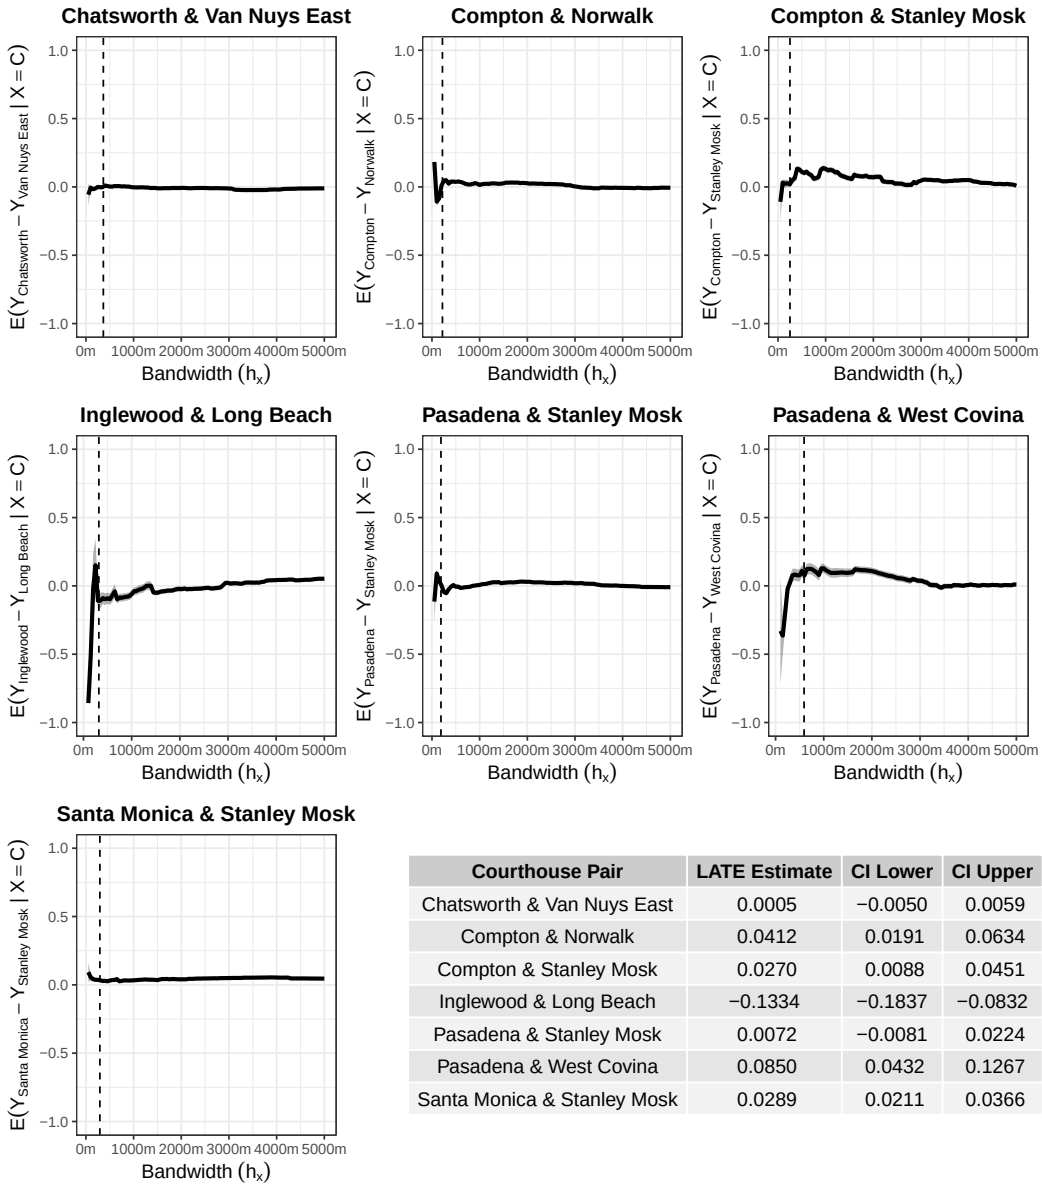

Courthouse Pair	LATE Estimate	CI Lower	CI Upper
Chatsworth & Van Nuys East	0.0005	−0.0050	0.0059
Compton & Norwalk	0.0412	0.0191	0.0634
Compton & Stanley Mosk	0.0270	0.0088	0.0451
Inglewood & Long Beach	−0.1334	−0.1837	−0.0832
Pasadena & Stanley Mosk	0.0072	−0.0081	0.0224
Pasadena & West Covina	0.0850	0.0432	0.1267
Santa Monica & Stanley Mosk	0.0289	0.0211	0.0366

Figure B.19: LATE Estimates by Courthouse Pair

Note: The LATE estimates ($\hat{\tau}_C$) at each courthouse pair boundary C are plotted as a function of bandwidth h_x. For each bandwidth h_x, the local linear estimate is plotted along with the CIs (gray ribbon). The table again gives the `rdrobust` package estimates of the LATE, along with the robust CIs. The dashed vertical line is the `rdrobust` optimal bandwidth.

Next, we briefly discuss our similarity score matching procedure to do the "Same Address" (Column (3)) robustness check. The estimates are obtained from exclud-

ing cases where the Jaro-Winkler (JW) string similarity score (range 0.0-1.0), which measures differences between the observed address and the assessor address, falls below a similarity threshold (0.88 here). To compute similarity scores we first standardize addresses by turning the assessor and docket addresses to lower case and standardizing certain street naming conventions (e.g., "123 Main Street" becomes "123 Main St"). Then we use the `stringdist::stringsim(method="jw",p=0.1)` function with `method="jw", p=0.1` to compute similarity scores, which uses the Jaro-Winkler method with an increased score for matching characters near the beginning of the string, which gives a higher similarity score if the street numbers match at beginning of docket address and the assessor address. For example, using this procedure the docket address "38059 High County Road Palmdale CA 93551" and the assessor address "38059 High Country Rd Palmdale CA 93551" are recorded as matching addresses with a score of 0.97.

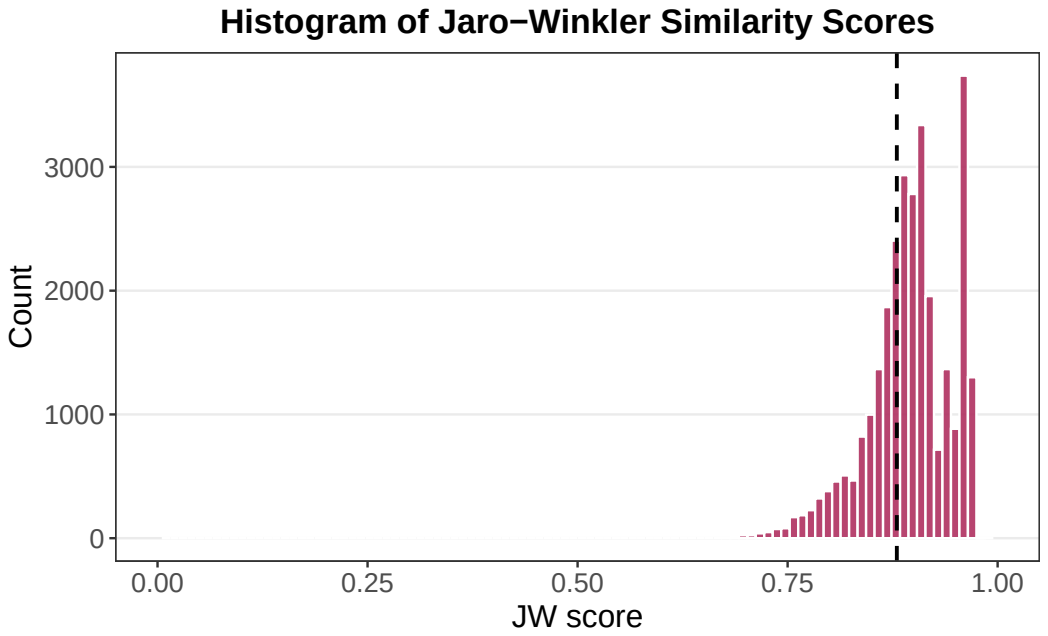

Figure B.20: JW Similarity Scores

Note: The threshold (0.88) for the robustness check is shown as the dashed black vertical line.

Additional Outcome RDDs: LATE Estimates

Finally, we show additional LATE estimates for further outcome variables.

	(1) Total Time	(2) Time-to-Default (E)	(3) Time-to-Default (L)	(4) Money Judgment
Chatsworth and Van Nuys East				
Mean (1st Court)	270.08	29.61	31.08	1207.35
Mean (2nd Court)	174.83	27.02	29.66	12823.35
Treatment Effect	81.83	1.68	-0.36	-54242.86
95% CI	(-13.18, 176.84)	(-1.38, 4.75)	(-4.56, 3.84)	(-152249.58, 43763.87)
p-value	0.09	0.28	0.87	0.28
Compton and Norwalk				
Mean (1st Court)	137.95	27.41	28.69	1506.83
Mean (2nd Court)	394.69	27.30	27.79	685.10
Treatment Effect	-155.01	-4.21	-3.51	1466.38
95% CI	(-391.30, 81.29)	(-8.44, 0.02)	(-9.77, 2.74)	(331.70, 2601.06)
p-value	0.20	0.05	0.27	0.01
Compton and Stanley Mosk				
Mean (1st Court)	121.54	25.91	26.46	417.86
Mean (2nd Court)	142.38	26.20	27.38	209.54
Treatment Effect	-46.87	2.63	2.08	-160.30
95% CI	(-120.86, 27.12)	(-2.82, 8.08)	(-3.75, 7.91)	(-862.04, 541.44)
p-value	0.21	0.34	0.48	0.65
Inglewood and Long Beach				
Mean (1st Court)	266.55	25.77	27.13	838.90
Mean (2nd Court)	120.04	31.25	31.45	818.59
Treatment Effect	-60.74	0.15	2.72	-1060.75
95% CI	(-265.15, 143.67)	(-9.19, 9.49)	(-7.38, 12.83)	(-2751.98, 630.48)
p-value	0.56	0.97	0.60	0.22
Pasadena and Stanley Mosk				
Mean (1st Court)	172.38	29.59	33.62	432.51
Mean (2nd Court)	203.33	26.29	29.60	768.55
Treatment Effect	-33.65	-0.09	0.71	403.55
95% CI	(-91.68, 24.37)	(-2.90, 2.72)	(-4.70, 6.13)	(-355.42, 1162.53)
p-value	0.26	0.95	0.80	0.30
Pasadena and West Covina				
Mean (1st Court)	293.49	26.84	27.13	1475.27
Mean (2nd Court)	160.17	26.30	28.48	857.56
Treatment Effect	243.42	-4.57	-4.10	-706.73
95% CI	(-241.24, 728.09)	(-10.86, 1.71)	(-10.93, 2.74)	(-2343.50, 930.03)
p-value	0.32	0.15	0.24	0.40
Santa Monica and Stanley Mosk				
Mean (1st Court)	125.11	24.70	26.69	1104.45
Mean (2nd Court)	171.09	27.54	30.18	1363.86
Treatment Effect	-6.86	-5.01	-3.32	-123.51
95% CI	(-76.60, 62.88)	(-7.66, -2.37)	(-8.11, 1.47)	(-1035.12, 788.10)
p-value	0.85	0.00	0.17	0.79

Table B.14: Additional Outcome LATE Estimates

Note: Column (1) uses total case time as the outcome variable. Column (2) use the early-estimated time-to-default measure as the outcome variable. Column (3) uses the late-estimated time-to-default measure as the outcome variable. Column (4) uses total money judgments as the outcome variable. The court means and treatment effects are in days for Columns (1)-(3) and dollars for Column (4).

	(1) Total Time	(2) Time-to-Default (E)	(3) Time-to-Default (L)	(4) Money Judgment
Chatsworth and Van Nuys East				
Bandwidth	978.40	766.04	860.24	2207.82
Number Obs.	2294	1706	1706	1661
Mean (1st Court)	271.40	29.38	30.62	1267.67
Mean (2nd Court)	176.12	27.11	29.92	19761.78
Treatment Effect	92.20	2.33	-0.63	-72042.39
95% CI	(-4.44, 188.83)	(-0.76, 5.43)	(-5.17, 3.92)	(-186854.89, 42770.11)
p-value	0.06	0.14	0.79	0.22
Bonferroni p-value	0.43	0.97	1.00	1.00
Compton and Norwalk				
Bandwidth	1655.79	1114.78	1290.58	1048.52
Number Obs.	2038	1626	1626	1547
Mean (1st Court)	146.99	27.54	28.80	1506.83
Mean (2nd Court)	391.07	27.65	28.07	637.33
Treatment Effect	-169.26	-4.52	-3.75	1312.68
95% CI	(-398.11, 59.59)	(-8.82, -0.22)	(-9.82, 2.31)	(177.71, 2447.65)
p-value	0.15	0.04	0.23	0.02
Bonferroni p-value	1.00	0.27	1.00	0.16
Compton and Stanley Mosk				
Bandwidth	625.63	1099.50	873.66	1086.82
Number Obs.	1905	1474	1474	1374
Mean (1st Court)	117.33	26.74	26.37	391.98
Mean (2nd Court)	156.95	26.01	27.85	206.55
Treatment Effect	-66.73	-0.54	0.56	-113.55
95% CI	(-158.43, 24.97)	(-6.00, 4.91)	(-5.38, 6.49)	(-846.67, 619.57)
p-value	0.15	0.85	0.85	0.76
Bonferroni p-value	1.00	1.00	1.00	1.00
Inglewood and Long Beach				
Bandwidth	895.63	795.68	739.86	992.57
Number Obs.	799	621	621	593
Mean (1st Court)	254.99	26.17	27.30	836.97
Mean (2nd Court)	117.30	31.73	31.78	951.85
Treatment Effect	-12.02	1.60	2.75	-1022.51
95% CI	(-204.04, 180.01)	(-8.21, 11.41)	(-7.76, 13.26)	(-2708.51, 663.49)
p-value	0.90	0.75	0.61	0.23
Bonferroni p-value	1.00	1.00	1.00	1.00
Pasadena and Stanley Mosk				
Bandwidth	1230.69	1328.42	1063.91	739.12
Number Obs.	4899	3551	3549	3187
Mean (1st Court)	166.89	29.75	33.70	427.71
Mean (2nd Court)	199.64	26.40	29.61	803.29
Treatment Effect	-35.42	-0.18	1.91	399.38
95% CI	(-89.13, 18.30)	(-2.96, 2.59)	(-3.66, 7.47)	(-359.33, 1158.08)
p-value	0.20	0.90	0.50	0.30
Bonferroni p-value	1.00	1.00	1.00	1.00
Pasadena and West Covina				
Bandwidth	1747.81	896.54	972.60	890.08
Number Obs.	1020	796	796	736
Mean (1st Court)	274.00	26.81	27.02	1215.22
Mean (2nd Court)	143.85	26.47	28.48	932.49
Treatment Effect	224.52	-4.83	-4.89	-1328.63
95% CI	(-202.93, 651.97)	(-11.77, 2.12)	(-11.88, 2.10)	(-2511.23, -146.04)
p-value	0.30	0.17	0.17	0.03
Bonferroni p-value	1.00	1.00	1.00	0.19
Santa Monica and Stanley Mosk				
Bandwidth	559.26	1147.30	1014.76	1015.48
Number Obs.	5310	3802	3801	3558
Mean (1st Court)	121.64	24.74	25.84	1048.48
Mean (2nd Court)	171.26	27.76	30.66	1342.69
Treatment Effect	-22.13	-5.05	-4.57	-9.32
95% CI	(-81.18, 36.93)	(-7.98, -2.12)	(-8.25, -0.89)	(-893.59, 874.95)
p-value	0.46	0.00	0.02	0.98
Bonferroni p-value	1.00	0.01	0.11	1.00

Table B.15: Additional Outcome LATE Estimates (w/Covariates)

B.11 Mechanism Model

Appendix B.11 includes additional information on the mechanism model, including a simulation example showing how parameter estimates are recovered and additional mechanism model results plots.

Mechanism Model Simulation

This sub-section simulates a spatial RDD mechanism model compatible with the assumptions made in-text. The goal is to show how heterogeneous segment-specific gaps may be drawn along a boundary, yet common mechanism coefficients like β_D (distance-to-court mechanism coefficient) and β_P (procedure mechanism coefficient) may be recovered. Here, we use A and B to denote two mechanisms in general, but above we are considering the tenant cost mechanism (measured by distance-to-court or generalized costs) and the procedure mechanism.

Suppose there is a single boundary with 50 segments s, with units on each side of the boundary having running variable $X_{is} \sim U[-1, 1]$ and treatment being $W_{is} = \mathbf{1}\{X_{is} \geq 0\}$. Suppose there are two mechanisms A and B of interest and outcome Y, each of which may have segment-specific intercepts and slopes. Accordingly, baseline levels and trends (in X) may vary (continuously) across spatial segments s. Define the following equations for the simulation:

$$\begin{aligned} A_{is} &= \alpha^A_{0s} + \alpha^A_{1s} X_{is} + \Delta A_s * W_{is} + \epsilon^A_{is} \\ B_{is} &= \alpha^B_{0s} + \alpha^B_{1s} X_{is} + \Delta B_s * W_{is} + \epsilon^B_{is} \\ Y_{is} &= \alpha^Y_{0s} + \alpha^Y_{1s} X_{is} + \Delta Y_s * W_{is} + \epsilon^Y_{is} \end{aligned}$$

The identifying assumption imposed on the model is in how outcome jumps are generated:

$$\Delta Y_s = \beta_A \Delta A_s + \beta_B \Delta B_s + \eta_s$$

which allows for heterogeneity in levels and slopes across segments, but restricts jumps or discontinuities in the outcome variable Y to be caused only by jumps in the mechanisms.

In the simulation that follows, we set $\beta_A = 0.45$ and $\beta_B = -0.35$ with 50 boundary segments and 250 observations per segment, which is similar to the 50 boundary segment specifications run in the actual estimation process. Segment-specific mechanism jumps are drawn from:

$$\Delta A_s \sim N(0.6, 0.35^2) \quad \text{and} \quad \Delta B_s \sim N(0.0, 0.45^2)$$

Segment-specific intercepts and slopes (which vary with X in the DGPs specified above) are drawn iid from normal distributions:

$$\begin{aligned}
\alpha^{A}_{0s} &\sim N(0.5, 0.5^2) \quad \text{and} \quad \alpha^{A}_{1s} \sim N(0.2, 0.3^2) \\
\alpha^{B}_{0s} &\sim N(0.0, 0.7^2) \quad \text{and} \quad \alpha^{B}_{1s} \sim N(-0.1, 0.3^2) \\
\alpha^{Y}_{0s} &\sim N(0.2, 0.8^2) \quad \text{and} \quad \alpha^{Y}_{1s} \sim N(0.15, 0.35^2)
\end{aligned}$$

Epsilon errors in the DGPs are $N(0, 0.15^2)$ for A and B and $N(0, 0.20^2)$ for Y, with further sampling errors $\eta_s \sim N(0, 0.03^2)$. And, as noted above, the outcome jumps are generated according to $\Delta Y_s = \beta_A \Delta A_s + \beta_B \Delta B_s + \eta_s$, so there is additional error in the outcome jump measure (η_s mean zero and independent of the mechanisms).

Using simulated data, we proceed to estimate β_A and β_B in two stages. In the first stage, we estimate segment-specific discontinuities in the mechanism measures around each segment s using bandwidth 0.5. In the second stage, we estimate the β coefficients by regressing across segments the estimated outcome Y gaps on estimated mechanism gaps from stage 1.

The simulation plot below shows how identification from segment variation is possible, even though segments themselves have heterogeneous level and slopes (as a continuous function of X, that is). This means we can recover mechanism coefficients even when levels and slopes are confounded across space, provided the identifying assumptions are satisfied (continuity plus shared mechanism coefficients). Most importantly, this means the outcome gaps must be generated by the structural equation, i.e., only gaps in the two mechanisms cause gaps in the outcome, and other spatial heterogeneity must be a continuous function of X.

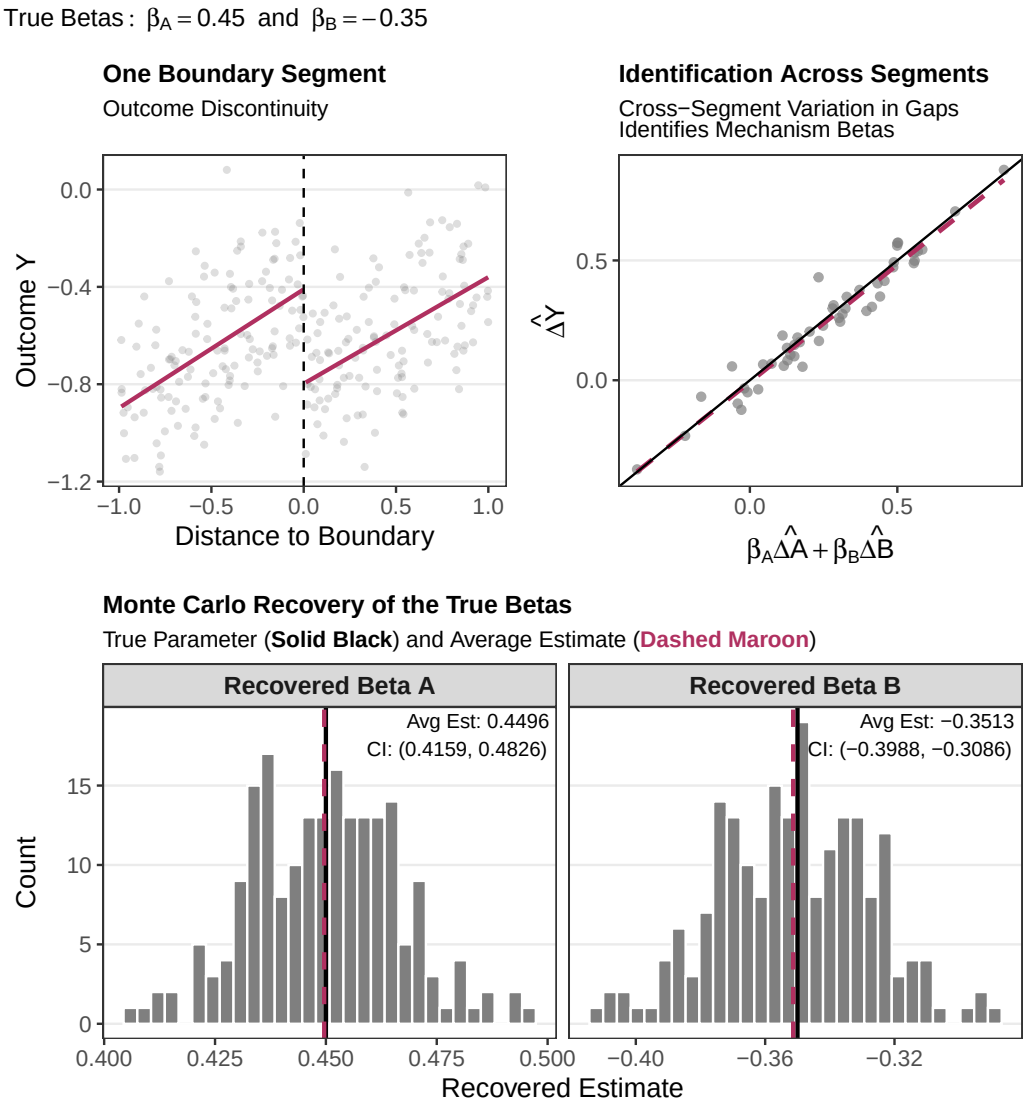

Figure B.21: Mechanism Simulation

Note: The simulation uses 50 boundary segments with segment-specific levels and slopes, but common structural coefficients where discontinuities in mechanisms causes discontinuities in outcomes. The top-left panel shows an example segment-specific RDD. The top-right panel illustrates how cross-segment variation in mechanism gaps maps into estimated outcome gaps when the betas are set to true values. Finally, the bottom panels show MC simulation estimate histograms recovering the mechanism beta coefficients: the true coefficients are plotted as solid black vertical lines, whereas the dashed maroon vertical lines are average recovered estimates.

Reduced Form Results

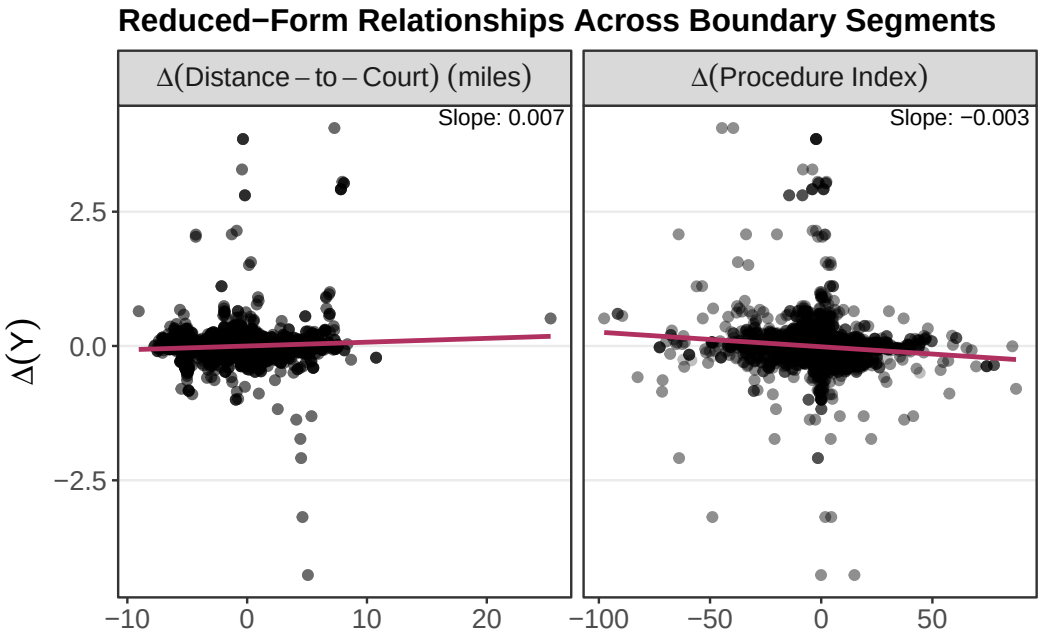

Figure B.22: Reduced Form Relationships Across Boundary Segments

Note: In the right panel, local linear estimated gaps outside the possible range (-200 to 200) are omitted.

Generalized Costs

Additional plots on the generalized cost metric as a measure of the tenant cost of getting to court are included below. The first plot shows the positive relationship between generalized cost variable and the distance-to-court variable. The second plot shows generalized costs over time.

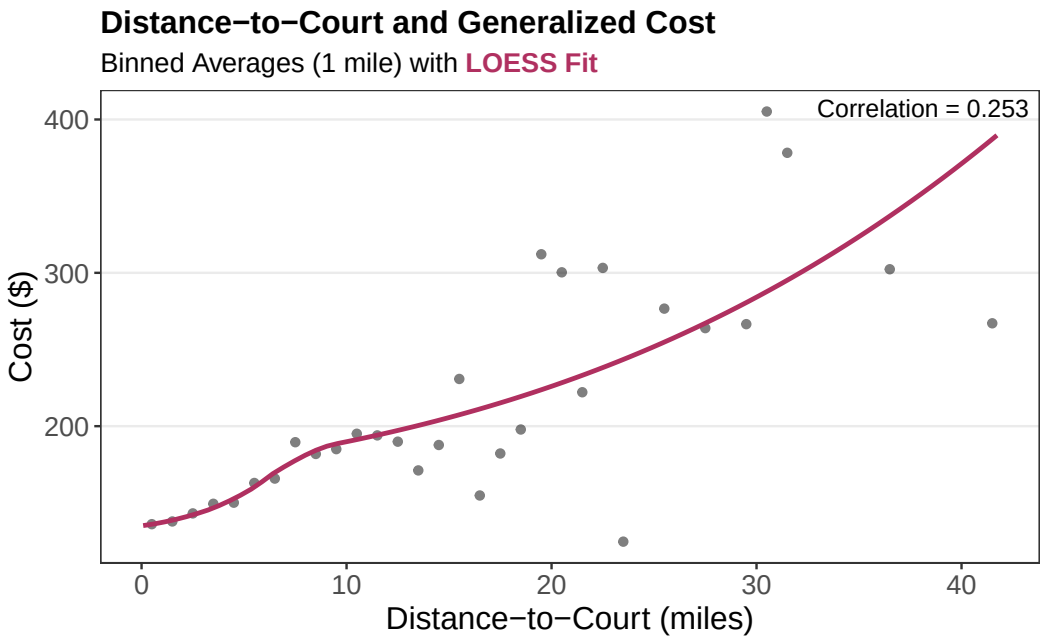

Figure B.23: Distance-to-Court and Generalized Cost Relationship

Note: The 1-mile binned averages are shown as points, with a LOESS line of fit in red. The correlation between the measures is shown in the upper-right. All nominal costs are inflation-adjusted to December 2024 dollars using the FRED CPI-U series (CPIAUCSL).

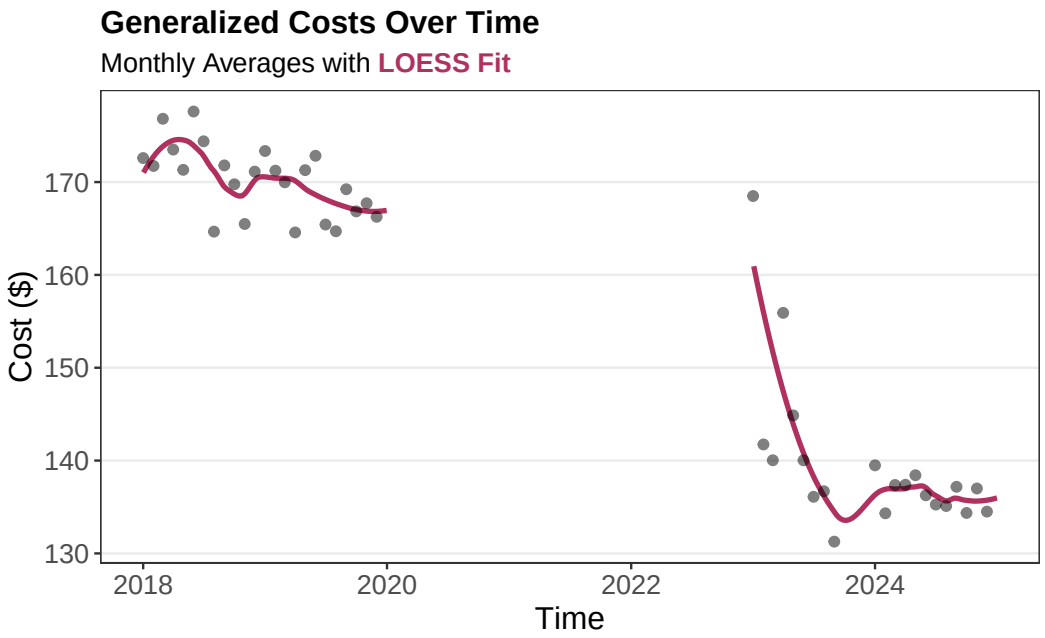

Figure B.24: Generalized Costs Over Time

Note: Monthly average inflation-adjusted costs are shown as points, with a LOESS fit in maroon. Data from the COVID-19 moratoria are unavailable. All nominal costs are inflation-adjusted to December 2024 dollars using the FRED CPI-U series (CPIAUCSL).

Mechanism Model Estimates: Summary Histograms

This section includes histograms summarizing beta coefficient estimates across specifications. Estimated coefficient histograms for two primary outcome measures (case defaults vs other case outcomes; and imputed defaults) are included in each plot. Average estimates (plotted as vertical black dashed line in each panel), median estimates, and the interquartile range (IQR) are reported in each panel. Some of the plots report coefficient estimates from using the generalized cost measure as our tenant cost variable instead of the distance-to-court measure alone.

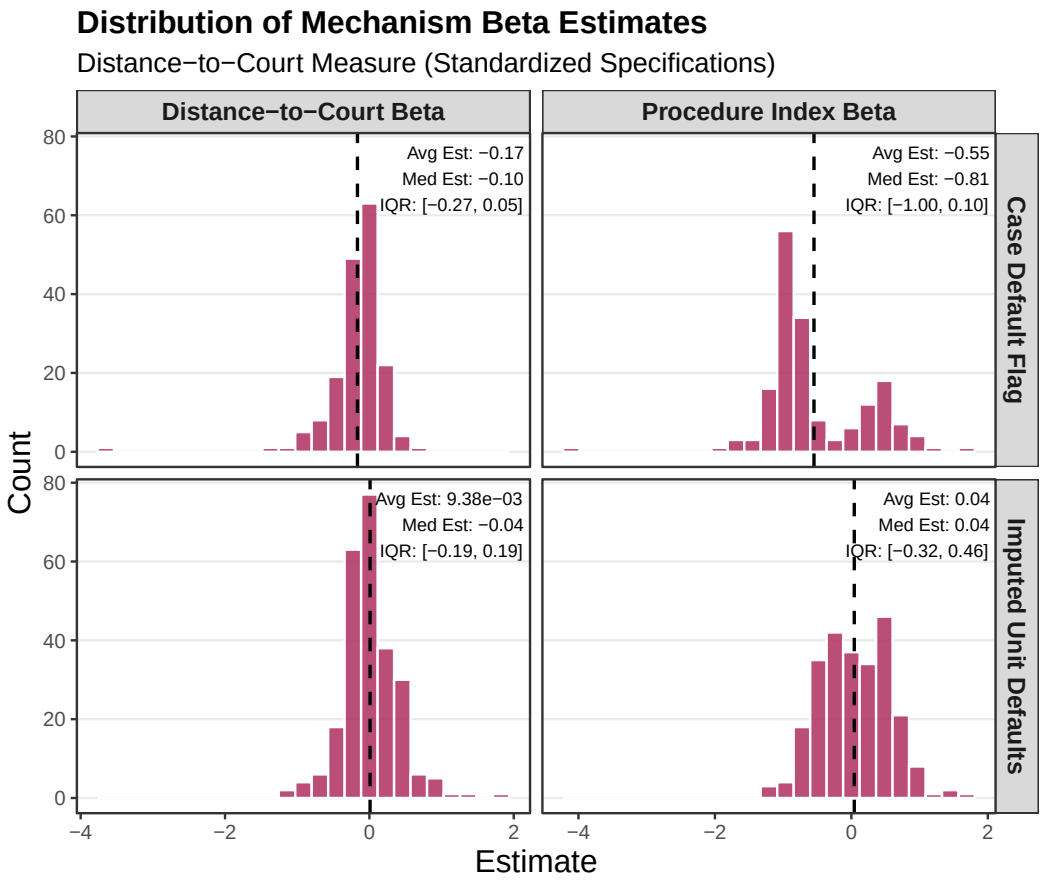

Figure B.25: Distribution of Mechanism Beta Estimates (Distance, Standardized)

Note: Each panel reports standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the distance-to-court variable.

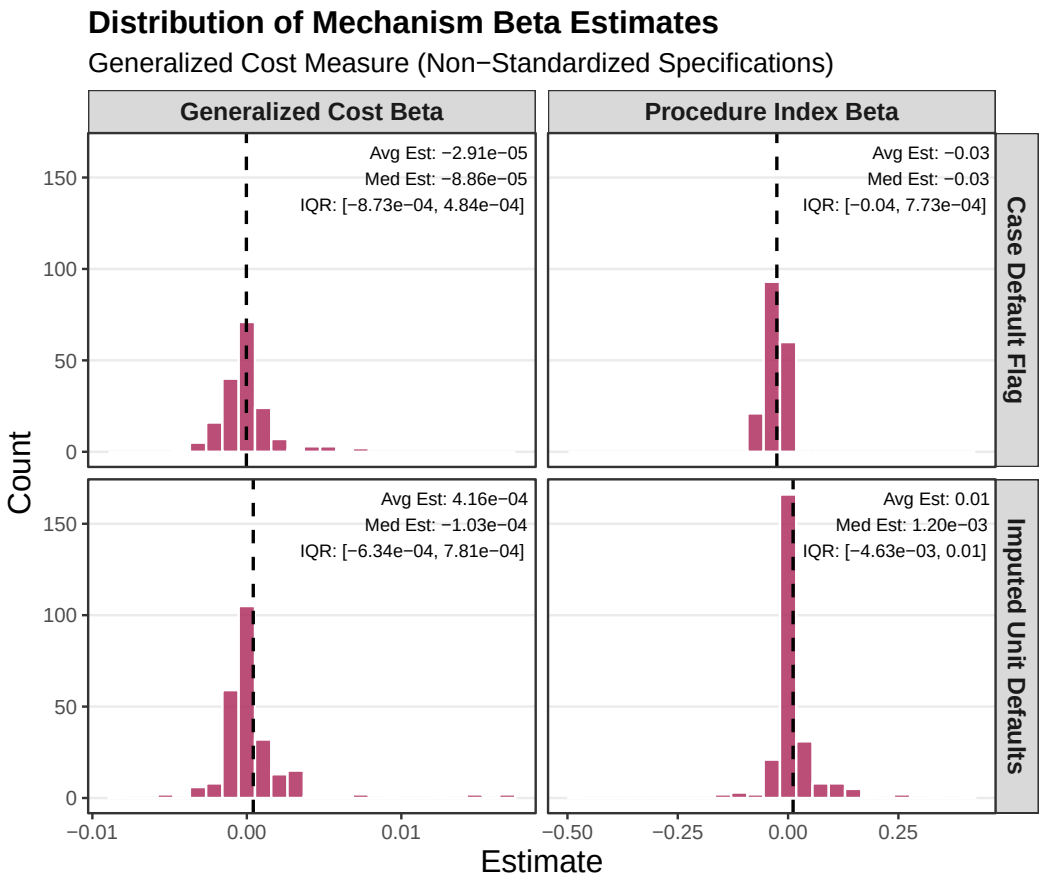

Figure B.26: Distribution of Mechanism Beta Estimates (Generalized Cost, Non-Standardized)

Note: Each panel reports non-standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the generalized cost variable.

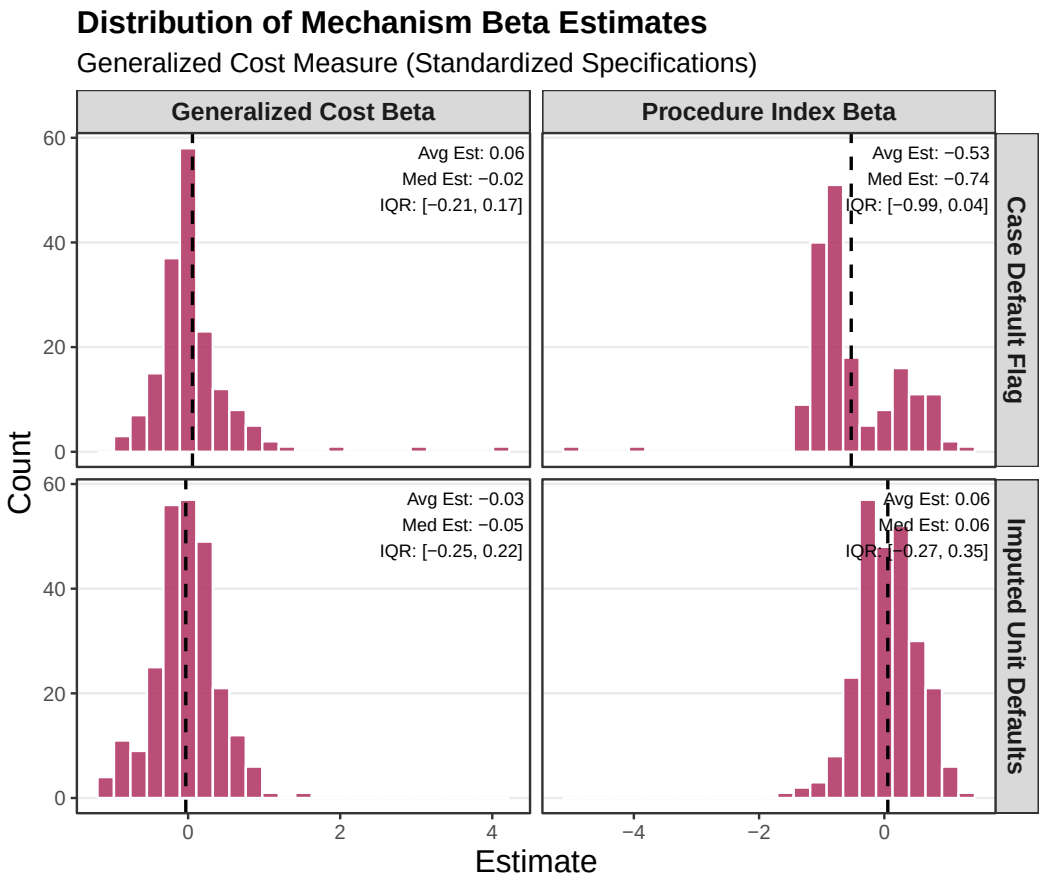

Figure B.27: Distribution of Mechanism Beta Estimates (Generalized Cost, Standardized)

Note: Each panel reports standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the generalized cost variable.

Mechanism Model Unweighted Estimates: Summary Histograms

Next, we report histograms of mechanism beta coefficients across models using unweighted linear regressions. The main specifications use a weighted regression, where the weighting comes from the number of observations within the bandwidth near a boundary segment: segments with fewer observations are noisier, so the weighted regressions upweight boundary segments with more observations (larger sample size). However, we also report beta estimates across specifications from unweighted regressions, where each boundary segment contributes equally to the estimation irrespective of local sample sizes.

Four histograms are reported as before, along two dimensions: (1) standardized vs non-standardized estimates, (2) using two tenant cost measures (distance-to-court and generalized cost). Compare, for example, the unweighted and weighted standardized histograms using the distance-to-court tenant cost measure: the average and median estimates in the four panels are roughly unchanged, although the unweighted estimates on the distance-to-court beta (β_D) is slightly larger in magnitude and noisier (wider IQR) in the imputed unit default specifications. Unweighted estimates using the generalized cost measure are similar: the generalized cost betas are different in the imputed unit default specifications (negative to positive), but remain small in standardized terms with larger uncertainty (wider IQR). Weighted vs unweighted regressions in general do not seem to meaningfully affect the case default flag specifications.

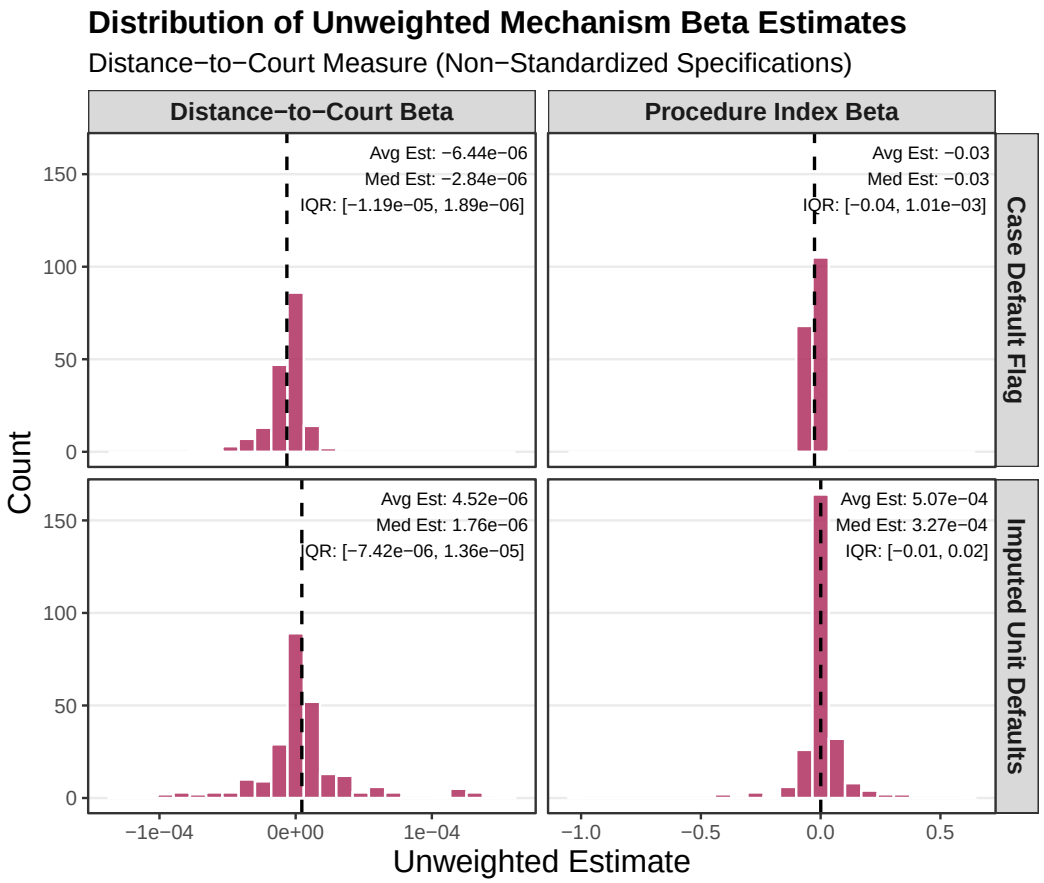

Figure B.28: Distribution of Unweighted Mechanism Beta Estimates (Distance-to-Court, Non-Standardized)

Note: Each panel reports non-standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the distance-to-court variable. Estimates are from unweighted linear regressions of boundary segment outcome gaps on boundary segment mechanism gaps.

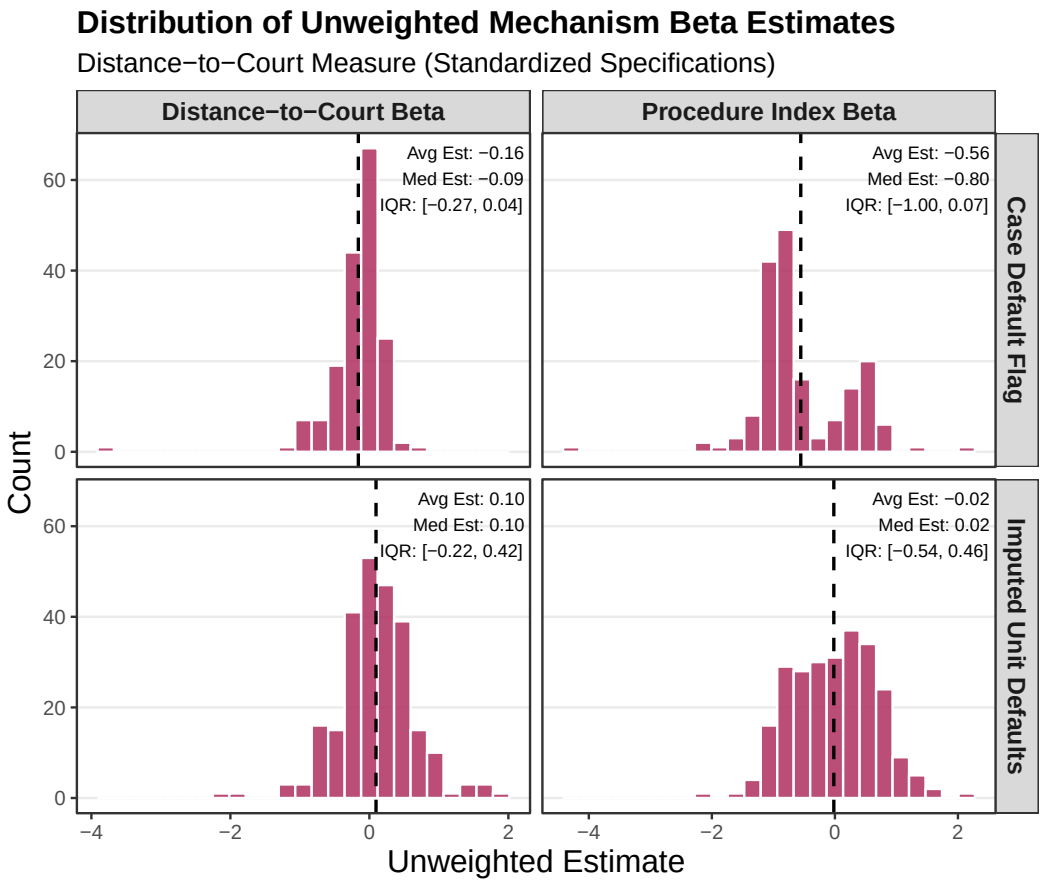

Figure B.29: Distribution of Unweighted Mechanism Beta Estimates (Distance-to-Court, Standardized)

Note: Each panel reports standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the distance-to-court variable. Estimates are from unweighted linear regressions of boundary segment outcome gaps on boundary segment mechanism gaps.

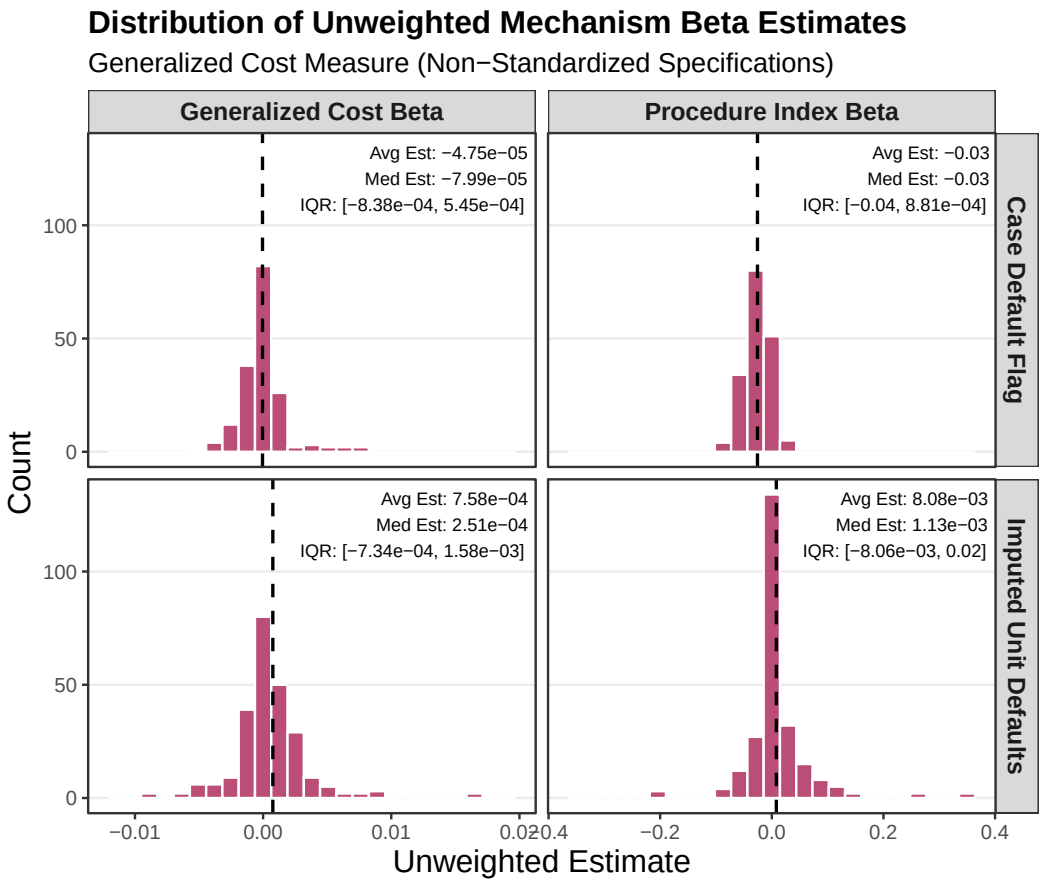

Figure B.30: Distribution of Unweighted Mechanism Beta Estimates (Generalized Cost, Non-Standardized)

Note: Each panel reports non-standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the generalized cost variable. Estimates are from unweighted linear regressions of boundary segment outcome gaps on boundary segment mechanism gaps.

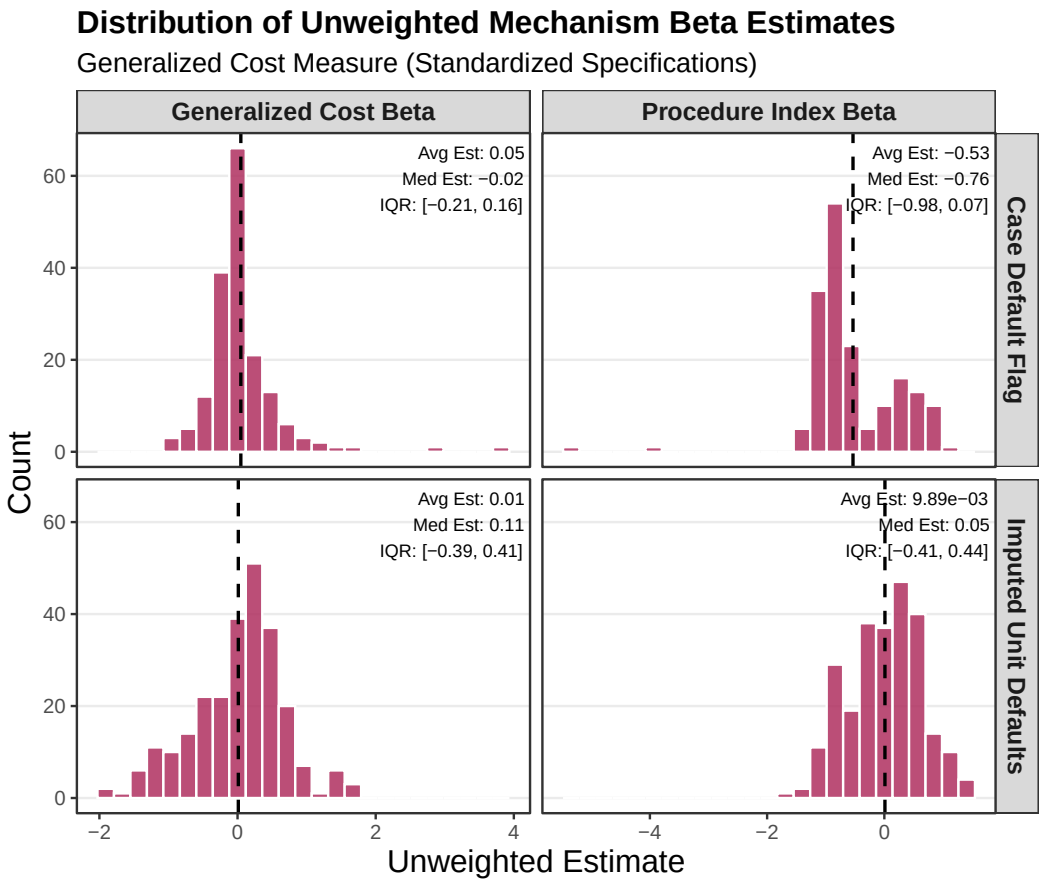

Figure B.31: Distribution of Unweighted Mechanism Beta Estimates (Generalized Cost, Standardized)

Note: Each panel reports standardized beta coefficient estimates using two outcome measures (case defaults flags and imputed defaults). Average estimate, median estimate, and IQR are reported in each panel, with the average plotted as the black dashed vertical line. The tenant cost measure used here is the generalized cost variable. Estimates are from unweighted linear regressions of boundary segment outcome gaps on boundary segment mechanism gaps.

Mechanism Model Estimates by Courthouse Pair

Finally, we include detailed tables of mechanism model estimates by courthouse pairs. We include average coefficient estimates from varying specification settings, including: the procedure measure; the outcome measure; the first stage estimation method for the mechanism gaps (local linear vs difference-in-means estimators); bandwidth settings around each segment s; and the number of boundary points sampled along each courthouse pair boundary. Average estimates, median estimates, and median confidence interval widths are included across specifications for both the

distance-to-court measure of tenant costs and the generalized cost measure discussed in text.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	-6.370e-06	7.810e-07	6.497e-06
procedure_core	Distance-to-Court Beta	-4.031e-06	5.951e-07	1.123e-05
procedure_extended	Distance-to-Court Beta	-7.983e-06	9.360e-07	7.505e-06
procedure_all	Procedure Index Beta	-3.189e-03	-2.173e-03	0.018
procedure_core	Procedure Index Beta	5.465e-04	-6.324e-04	4.225e-03
procedure_extended	Procedure Index Beta	0.011	3.408e-03	0.050
Outcome Spec				
case_default_flag	Distance-to-Court Beta	4.115e-06	4.126e-06	1.646e-05
imputed_unit_defaults	Distance-to-Court Beta	-1.052e-05	-3.216e-07	5.577e-06
case_default_flag	Procedure Index Beta	-0.025	-0.031	0.020
imputed_unit_defaults	Procedure Index Beta	0.015	-1.079e-04	0.020
First-Stage Method				
local_linear	Distance-to-Court Beta	-1.406e-05	-1.924e-06	1.810e-05
means	Distance-to-Court Beta	1.800e-06	1.010e-06	5.531e-06
local_linear	Procedure Index Beta	0.013	-3.001e-05	0.028
means	Procedure Index Beta	-7.370e-03	-8.283e-04	0.017
Local Bandwidth				
0.5 km	Distance-to-Court Beta	-7.523e-06	-2.490e-06	1.197e-05
1 km	Distance-to-Court Beta	-2.153e-05	-1.061e-06	1.123e-05
2 km	Distance-to-Court Beta	1.807e-06	9.691e-07	7.007e-06
0.5 km	Procedure Index Beta	0.012	7.530e-03	0.019
1 km	Procedure Index Beta	0.031	2.913e-03	0.036
2 km	Procedure Index Beta	-0.013	-2.002e-03	0.018
Boundary Points				
10	Distance-to-Court Beta	1.562e-06	9.240e-07	8.573e-06
20	Distance-to-Court Beta	-1.865e-05	6.349e-07	7.247e-06
50	Distance-to-Court Beta	-2.504e-06	6.670e-07	5.700e-06
10	Procedure Index Beta	-9.985e-03	-1.470e-03	0.028
20	Procedure Index Beta	0.011	1.179e-03	0.018
50	Procedure Index Beta	6.667e-03	-2.704e-04	0.016

Table B.16: Chatsworth and Van Nuys East (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	-3.427e-05	4.273e-05	7.350e-04
procedure_core	Generalized Cost Beta	-5.147e-04	-1.416e-04	7.311e-04
procedure_extended	Generalized Cost Beta	-1.823e-04	-3.171e-05	8.365e-04
procedure_all	Procedure Index Beta	-9.197e-03	-5.685e-03	0.017
procedure_core	Procedure Index Beta	-3.211e-04	-4.131e-04	4.280e-03
procedure_extended	Procedure Index Beta	-2.972e-03	-7.195e-03	0.049
Outcome Spec				
case_default_flag	Generalized Cost Beta	-1.816e-04	-1.583e-04	1.557e-03
imputed_unit_defaults	Generalized Cost Beta	-2.704e-04	-2.074e-05	4.728e-04
case_default_flag	Procedure Index Beta	-0.024	-0.027	0.020
imputed_unit_defaults	Procedure Index Beta	4.206e-03	-3.824e-04	0.018
First-Stage Method				
local_linear	Generalized Cost Beta	-4.538e-04	-1.385e-04	1.557e-03
means	Generalized Cost Beta	-3.375e-05	-4.192e-06	3.766e-04
local_linear	Procedure Index Beta	1.100e-03	-6.986e-04	0.026
means	Procedure Index Beta	-9.427e-03	-8.257e-04	0.016
Local Bandwidth				
0.5 km	Generalized Cost Beta	-2.074e-04	-1.176e-04	2.141e-03
1 km	Generalized Cost Beta	-7.049e-04	-1.736e-04	7.800e-04
2 km	Generalized Cost Beta	-1.925e-05	5.260e-05	6.782e-04
0.5 km	Procedure Index Beta	4.452e-03	3.253e-03	0.017
1 km	Procedure Index Beta	0.010	-3.844e-04	0.031
2 km	Procedure Index Beta	-0.013	-3.087e-03	0.016
Boundary Points				
10	Generalized Cost Beta	-1.133e-04	-6.692e-05	7.096e-04
20	Generalized Cost Beta	-6.631e-04	-1.587e-04	6.701e-04
50	Generalized Cost Beta	-2.708e-05	8.264e-06	7.486e-04
10	Procedure Index Beta	-0.011	-1.505e-03	0.028
20	Procedure Index Beta	-8.606e-03	-8.257e-04	0.020
50	Procedure Index Beta	4.063e-03	-3.824e-04	0.015

Table B.17: Chatsworth and Van Nuys East (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	5.930e-06	-2.210e-06	1.991e-05
procedure_core	Distance-to-Court Beta	-3.548e-06	-4.240e-06	2.848e-05
procedure_extended	Distance-to-Court Beta	2.784e-06	-2.705e-06	2.143e-05
procedure_all	Procedure Index Beta	9.850e-03	-0.016	0.035
procedure_core	Procedure Index Beta	-2.103e-03	-1.032e-03	8.495e-03
procedure_extended	Procedure Index Beta	0.016	-0.027	0.075
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-1.039e-05	-8.422e-06	2.630e-05
imputed_unit_defaults	Distance-to-Court Beta	1.038e-05	1.505e-06	1.677e-05
case_default_flag	Procedure Index Beta	-0.034	-0.035	0.035
imputed_unit_defaults	Procedure Index Beta	0.038	3.613e-03	0.037
First-Stage Method				
local_linear	Distance-to-Court Beta	5.617e-06	-2.115e-06	3.241e-05
means	Distance-to-Court Beta	-2.172e-06	-2.695e-06	1.690e-05
local_linear	Procedure Index Beta	3.953e-03	-2.100e-03	0.043
means	Procedure Index Beta	0.012	-3.932e-03	0.031
Local Bandwidth				
0.5 km	Distance-to-Court Beta	1.569e-05	-9.841e-06	1.295e-04
1 km	Distance-to-Court Beta	7.561e-06	4.906e-06	3.101e-05
2 km	Distance-to-Court Beta	-5.470e-06	-3.157e-06	1.131e-05
0.5 km	Procedure Index Beta	0.086	-1.158e-03	0.079
1 km	Procedure Index Beta	7.913e-04	-7.314e-03	0.051
2 km	Procedure Index Beta	7.832e-04	-1.032e-03	0.029
Boundary Points				
10	Distance-to-Court Beta	2.696e-06	-8.236e-07	1.794e-05
20	Distance-to-Court Beta	-6.095e-07	-2.695e-06	2.173e-05
50	Distance-to-Court Beta	3.003e-06	-3.437e-06	2.547e-05
10	Procedure Index Beta	0.015	-1.814e-03	0.034
20	Procedure Index Beta	-3.502e-03	-7.642e-03	0.039
50	Procedure Index Beta	0.013	2.902e-04	0.035

Table B.18: Compton and Norwalk (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	5.313e-04	5.061e-04	2.516e-03
procedure_core	Generalized Cost Beta	-2.667e-04	-1.929e-04	3.526e-03
procedure_extended	Generalized Cost Beta	2.142e-04	4.446e-04	2.368e-03
procedure_all	Procedure Index Beta	0.011	8.298e-04	0.043
procedure_core	Procedure Index Beta	-2.580e-03	-1.327e-03	0.011
procedure_extended	Procedure Index Beta	0.018	-3.337e-03	0.077
Outcome Spec				
case_default_flag	Generalized Cost Beta	-1.365e-03	-1.087e-03	4.443e-03
imputed_unit_defaults	Generalized Cost Beta	1.248e-03	9.861e-04	1.862e-03
case_default_flag	Procedure Index Beta	-0.033	-0.035	0.042
imputed_unit_defaults	Procedure Index Beta	0.039	5.818e-03	0.042
First-Stage Method				
local_linear	Generalized Cost Beta	2.687e-04	4.157e-04	3.671e-03
means	Generalized Cost Beta	5.056e-05	-7.328e-05	2.231e-03
local_linear	Procedure Index Beta	3.251e-03	-1.995e-03	0.048
means	Procedure Index Beta	0.014	1.015e-03	0.037
Local Bandwidth				
0.5 km	Generalized Cost Beta	8.900e-04	-2.676e-04	5.205e-03
1 km	Generalized Cost Beta	6.698e-04	1.079e-03	3.497e-03
2 km	Generalized Cost Beta	-3.872e-04	-2.697e-04	1.756e-03
0.5 km	Procedure Index Beta	0.084	2.790e-03	0.075
1 km	Procedure Index Beta	-4.768e-04	-7.153e-03	0.058
2 km	Procedure Index Beta	4.090e-03	-5.048e-05	0.026
Boundary Points				
10	Generalized Cost Beta	3.426e-04	4.157e-04	2.740e-03
20	Generalized Cost Beta	-1.481e-04	-7.328e-05	2.476e-03
50	Generalized Cost Beta	2.960e-04	-1.929e-04	3.362e-03
10	Procedure Index Beta	0.019	-7.179e-04	0.039
20	Procedure Index Beta	-2.427e-03	-8.135e-03	0.044
50	Procedure Index Beta	0.012	1.666e-03	0.042

Table B.19: Compton and Norwalk (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	1.046e-05	-2.818e-06	2.552e-05
procedure_core	Distance-to-Court Beta	-2.553e-06	-3.663e-06	4.900e-05
procedure_extended	Distance-to-Court Beta	7.148e-06	-2.869e-06	2.751e-05
procedure_all	Procedure Index Beta	-6.308e-03	-0.013	0.031
procedure_core	Procedure Index Beta	4.508e-03	5.601e-03	0.019
procedure_extended	Procedure Index Beta	-1.006e-03	-0.025	0.064
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-5.087e-06	-4.391e-06	4.014e-05
imputed_unit_defaults	Distance-to-Court Beta	1.400e-05	-1.258e-06	1.725e-05
case_default_flag	Procedure Index Beta	-0.021	-0.026	0.036
imputed_unit_defaults	Procedure Index Beta	0.017	5.608e-03	0.032
First-Stage Method				
local_linear	Distance-to-Court Beta	1.252e-05	-3.210e-06	3.458e-05
means	Distance-to-Court Beta	-2.485e-06	-2.843e-06	2.190e-05
local_linear	Procedure Index Beta	5.582e-03	-2.540e-03	0.035
means	Procedure Index Beta	-7.453e-03	1.416e-03	0.032
Local Bandwidth				
0.5 km	Distance-to-Court Beta	4.837e-05	5.065e-06	5.702e-05
1 km	Distance-to-Court Beta	-7.675e-06	-5.565e-06	1.874e-05
2 km	Distance-to-Court Beta	-1.720e-06	-2.698e-06	3.236e-05
0.5 km	Procedure Index Beta	0.042	0.017	0.033
1 km	Procedure Index Beta	-0.010	-3.402e-03	0.026
2 km	Procedure Index Beta	-0.010	-8.163e-03	0.044
Boundary Points				
10	Distance-to-Court Beta	4.142e-06	-3.475e-06	2.736e-05
20	Distance-to-Court Beta	6.152e-06	-3.070e-06	2.549e-05
50	Distance-to-Court Beta	4.611e-06	-2.249e-06	2.846e-05
10	Procedure Index Beta	2.303e-03	2.165e-06	0.034
20	Procedure Index Beta	-1.063e-03	-1.415e-03	0.033
50	Procedure Index Beta	-3.507e-03	5.017e-05	0.033

Table B.20: Compton and Stanley Mosk (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	2.252e-04	4.042e-06	3.984e-03
procedure_core	Generalized Cost Beta	-1.088e-03	-8.658e-04	7.767e-03
procedure_extended	Generalized Cost Beta	4.978e-05	-2.432e-04	4.293e-03
procedure_all	Procedure Index Beta	-7.126e-03	-0.017	0.030
procedure_core	Procedure Index Beta	3.643e-03	5.526e-03	0.021
procedure_extended	Procedure Index Beta	3.972e-03	-0.030	0.055
Outcome Spec				
case_default_flag	Generalized Cost Beta	-5.992e-04	-6.209e-04	6.800e-03
imputed_unit_defaults	Generalized Cost Beta	2.053e-05	-2.620e-04	2.386e-03
case_default_flag	Procedure Index Beta	-0.024	-0.031	0.043
imputed_unit_defaults	Procedure Index Beta	0.022	8.342e-03	0.034
First-Stage Method				
local_linear	Generalized Cost Beta	-1.141e-04	-1.552e-04	5.318e-03
means	Generalized Cost Beta	-4.281e-04	-6.638e-04	3.274e-03
local_linear	Procedure Index Beta	5.961e-03	-2.756e-03	0.051
means	Procedure Index Beta	-5.635e-03	5.703e-03	0.030
Local Bandwidth				
0.5 km	Generalized Cost Beta	5.600e-04	3.933e-04	7.536e-03
1 km	Generalized Cost Beta	-9.014e-04	-6.780e-04	4.588e-03
2 km	Generalized Cost Beta	-1.100e-04	-3.402e-04	4.151e-03
0.5 km	Procedure Index Beta	0.046	0.011	0.047
1 km	Procedure Index Beta	-0.012	-3.106e-03	0.027
2 km	Procedure Index Beta	-8.002e-03	1.415e-03	0.045
Boundary Points				
10	Generalized Cost Beta	-6.067e-04	-6.939e-04	3.702e-03
20	Generalized Cost Beta	9.594e-06	-1.248e-04	4.325e-03
50	Generalized Cost Beta	-2.721e-04	-4.322e-04	4.639e-03
10	Procedure Index Beta	3.321e-03	2.378e-03	0.033
20	Procedure Index Beta	4.432e-03	-2.780e-04	0.042
50	Procedure Index Beta	-6.737e-03	2.934e-04	0.038

Table B.21: Compton and Stanley Mosk (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	1.887e-06	-5.396e-07	2.760e-05
procedure_core	Distance-to-Court Beta	-1.597e-05	-1.479e-05	4.485e-05
procedure_extended	Distance-to-Court Beta	1.895e-06	3.067e-07	3.310e-05
procedure_all	Procedure Index Beta	-0.019	-0.029	0.037
procedure_core	Procedure Index Beta	-5.065e-03	-3.627e-03	9.381e-03
procedure_extended	Procedure Index Beta	-0.030	-0.041	0.125
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-7.112e-06	-1.145e-06	6.432e-05
imputed_unit_defaults	Distance-to-Court Beta	-1.621e-06	1.480e-07	1.815e-05
case_default_flag	Procedure Index Beta	-0.026	-0.029	0.036
imputed_unit_defaults	Procedure Index Beta	-0.011	9.094e-04	0.043
First-Stage Method				
local_linear	Distance-to-Court Beta	-7.335e-06	-1.915e-06	4.851e-05
means	Distance-to-Court Beta	-7.873e-07	3.610e-07	2.094e-05
local_linear	Procedure Index Beta	-0.027	-0.029	0.038
means	Procedure Index Beta	-8.962e-03	1.153e-03	0.040
Local Bandwidth				
1 km	Distance-to-Court Beta	-6.804e-06	-4.953e-07	5.170e-05
2 km	Distance-to-Court Beta	-2.690e-06	-2.350e-07	2.562e-05
1 km	Procedure Index Beta	5.367e-03	-1.545e-03	0.049
2 km	Procedure Index Beta	-0.030	-0.022	0.034
Boundary Points				
10	Distance-to-Court Beta	-2.356e-06	1.536e-06	3.311e-05
20	Distance-to-Court Beta	-2.787e-06	-8.832e-07	2.067e-05
50	Distance-to-Court Beta	-5.870e-06	-1.862e-06	3.644e-05
10	Procedure Index Beta	-0.041	-0.024	0.047
20	Procedure Index Beta	-8.339e-03	-3.988e-03	0.037
50	Procedure Index Beta	-0.014	-0.012	0.042

Table B.22: Inglewood and Long Beach (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	-2.353e-04	-7.733e-05	3.198e-03
procedure_core	Generalized Cost Beta	6.451e-04	-2.998e-04	4.408e-03
procedure_extended	Generalized Cost Beta	-2.640e-04	-5.330e-04	4.233e-03
procedure_all	Procedure Index Beta	-0.015	-0.025	0.036
procedure_core	Procedure Index Beta	-5.941e-04	-7.309e-04	7.768e-03
procedure_extended	Procedure Index Beta	-0.015	-0.020	0.093
Outcome Spec				
case_default_flag	Generalized Cost Beta	1.183e-03	7.281e-04	5.690e-03
imputed_unit_defaults	Generalized Cost Beta	-8.590e-04	-1.139e-03	1.733e-03
case_default_flag	Procedure Index Beta	-0.022	-0.023	0.046
imputed_unit_defaults	Procedure Index Beta	-3.281e-04	4.251e-04	0.040
First-Stage Method				
local_linear	Generalized Cost Beta	5.178e-04	-2.274e-04	5.727e-03
means	Generalized Cost Beta	-4.206e-04	-1.011e-03	1.476e-03
local_linear	Procedure Index Beta	-0.018	-0.023	0.054
means	Procedure Index Beta	-2.648e-03	4.313e-04	0.041
Local Bandwidth				
1 km	Generalized Cost Beta	-5.009e-04	-4.694e-04	4.417e-03
2 km	Generalized Cost Beta	3.233e-04	-7.733e-05	3.853e-03
1 km	Procedure Index Beta	0.012	-1.131e-03	0.056
2 km	Procedure Index Beta	-0.021	-0.011	0.044
Boundary Points				
10	Generalized Cost Beta	5.270e-04	-4.353e-05	7.240e-03
20	Generalized Cost Beta	-4.595e-05	-5.330e-04	2.744e-03
50	Generalized Cost Beta	-1.197e-04	-2.657e-04	3.853e-03
10	Procedure Index Beta	-0.024	-0.011	0.054
20	Procedure Index Beta	1.229e-04	-1.625e-04	0.036
50	Procedure Index Beta	-0.011	-3.488e-03	0.043

Table B.23: Inglewood and Long Beach (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	-6.473e-06	-4.865e-06	4.489e-05
procedure_core	Distance-to-Court Beta	-1.148e-05	-7.326e-06	4.468e-05
procedure_extended	Distance-to-Court Beta	-8.982e-06	-4.018e-06	5.327e-05
procedure_all	Procedure Index Beta	-0.013	-0.019	0.033
procedure_core	Procedure Index Beta	5.460e-03	3.322e-03	0.013
procedure_extended	Procedure Index Beta	-0.031	-0.039	0.087
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-2.402e-06	-3.086e-06	5.718e-05
imputed_unit_defaults	Distance-to-Court Beta	-1.424e-05	-7.942e-06	2.240e-05
case_default_flag	Procedure Index Beta	-0.025	-0.028	0.047
imputed_unit_defaults	Procedure Index Beta	-3.484e-03	2.393e-03	0.020
First-Stage Method				
local_linear	Distance-to-Court Beta	-1.352e-05	-9.524e-06	6.655e-05
means	Distance-to-Court Beta	-4.431e-06	-3.496e-06	1.799e-05
local_linear	Procedure Index Beta	-0.020	-0.012	0.047
means	Procedure Index Beta	-5.759e-03	4.227e-04	0.028
Local Bandwidth				
1 km	Distance-to-Court Beta	-1.359e-05	-9.208e-06	4.681e-05
2 km	Distance-to-Court Beta	-6.673e-06	-3.061e-06	4.749e-05
1 km	Procedure Index Beta	-0.022	-8.283e-05	0.032
2 km	Procedure Index Beta	-8.515e-03	-5.158e-03	0.042
Boundary Points				
10	Distance-to-Court Beta	-6.054e-07	-3.456e-06	5.786e-05
20	Distance-to-Court Beta	-1.720e-05	-5.310e-06	3.192e-05
50	Distance-to-Court Beta	-6.994e-06	-5.376e-06	4.659e-05
10	Procedure Index Beta	-6.161e-03	-5.158e-03	0.055
20	Procedure Index Beta	-0.020	4.837e-04	0.041
50	Procedure Index Beta	-0.011	-1.561e-04	0.031

Table B.24: Pasadena and Stanley Mosk (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	-2.102e-04	-4.725e-05	1.407e-03
procedure_core	Generalized Cost Beta	-4.182e-04	-8.877e-05	1.934e-03
procedure_extended	Generalized Cost Beta	-2.010e-04	-8.449e-05	1.695e-03
procedure_all	Procedure Index Beta	-0.010	-0.012	0.027
procedure_core	Procedure Index Beta	3.894e-03	4.617e-03	0.013
procedure_extended	Procedure Index Beta	-0.027	-0.028	0.059
Outcome Spec				
case_default_flag	Generalized Cost Beta	-1.354e-04	2.950e-05	2.384e-03
imputed_unit_defaults	Generalized Cost Beta	-3.893e-04	-1.264e-04	6.789e-04
case_default_flag	Procedure Index Beta	-0.021	-0.023	0.048
imputed_unit_defaults	Procedure Index Beta	-3.416e-03	3.272e-03	0.013
First-Stage Method				
local_linear	Generalized Cost Beta	-6.553e-04	-5.939e-04	2.212e-03
means	Generalized Cost Beta	1.024e-04	1.092e-04	5.021e-04
local_linear	Procedure Index Beta	-0.018	-0.011	0.040
means	Procedure Index Beta	-4.410e-03	2.995e-03	0.017
Local Bandwidth				
1 km	Generalized Cost Beta	-3.838e-04	7.620e-05	2.463e-03
2 km	Generalized Cost Beta	-2.228e-04	-9.999e-05	1.501e-03
1 km	Procedure Index Beta	-0.015	1.168e-03	0.029
2 km	Procedure Index Beta	-9.189e-03	3.145e-04	0.025
Boundary Points				
10	Generalized Cost Beta	-3.645e-04	-4.310e-04	2.446e-03
20	Generalized Cost Beta	-3.874e-04	-9.514e-05	1.486e-03
50	Generalized Cost Beta	-1.493e-04	5.111e-06	1.643e-03
10	Procedure Index Beta	-6.741e-03	7.240e-04	0.042
20	Procedure Index Beta	-0.017	8.816e-04	0.028
50	Procedure Index Beta	-8.709e-03	6.274e-04	0.025

Table B.25: Pasadena and Stanley Mosk (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	-1.762e-05	-1.408e-05	9.222e-05
procedure_core	Distance-to-Court Beta	-1.411e-05	-8.342e-06	9.927e-05
procedure_extended	Distance-to-Court Beta	-2.027e-05	-1.618e-05	1.210e-04
procedure_all	Procedure Index Beta	-0.036	-0.019	0.068
procedure_core	Procedure Index Beta	-2.084e-03	-3.366e-03	0.033
procedure_extended	Procedure Index Beta	-0.110	-0.041	0.161
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-1.849e-05	-1.927e-05	1.744e-04
imputed_unit_defaults	Distance-to-Court Beta	-1.694e-05	-3.397e-06	4.711e-05
case_default_flag	Procedure Index Beta	-0.023	-0.019	0.058
imputed_unit_defaults	Procedure Index Beta	-0.058	-8.495e-03	0.109
First-Stage Method				
local_linear	Distance-to-Court Beta	-3.537e-05	-3.254e-05	1.322e-04
means	Distance-to-Court Beta	7.028e-07	3.907e-06	2.245e-05
local_linear	Procedure Index Beta	-0.087	-0.018	0.292
means	Procedure Index Beta	-0.011	-0.011	0.060
Local Bandwidth				
1 km	Distance-to-Court Beta	-1.177e-05	-3.397e-06	4.711e-05
2 km	Distance-to-Court Beta	-1.919e-05	-1.912e-05	1.242e-04
1 km	Procedure Index Beta	0.144	-9.952e-04	0.104
2 km	Procedure Index Beta	-0.114	-0.019	0.064
Boundary Points				
20	Distance-to-Court Beta	-2.436e-05	-1.044e-05	7.053e-05
50	Distance-to-Court Beta	-1.499e-05	-8.342e-06	1.086e-04
20	Procedure Index Beta	-0.260	-0.043	0.101
50	Procedure Index Beta	0.021	-0.011	0.064

Table B.26: Pasadena and West Covina (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	4.062e-03	2.603e-03	6.178e-03
procedure_core	Generalized Cost Beta	6.076e-03	2.589e-03	6.674e-03
procedure_extended	Generalized Cost Beta	4.264e-03	2.149e-03	6.573e-03
procedure_all	Procedure Index Beta	-8.315e-03	-0.013	0.056
procedure_core	Procedure Index Beta	3.366e-04	-1.123e-03	0.029
procedure_extended	Procedure Index Beta	-0.019	-0.014	0.153
Outcome Spec				
case_default_flag	Generalized Cost Beta	-2.243e-04	-1.167e-04	4.687e-03
imputed_unit_defaults	Generalized Cost Beta	6.476e-03	3.001e-03	0.012
case_default_flag	Procedure Index Beta	-0.021	-0.020	0.057
imputed_unit_defaults	Procedure Index Beta	-4.994e-03	6.796e-04	0.093
First-Stage Method				
local_linear	Generalized Cost Beta	7.761e-03	8.287e-03	0.021
means	Generalized Cost Beta	1.840e-03	2.186e-03	1.887e-03
local_linear	Procedure Index Beta	-0.014	-5.212e-03	0.235
means	Procedure Index Beta	-4.066e-03	-1.123e-03	0.056
Local Bandwidth				
1 km	Generalized Cost Beta	4.262e-03	2.086e-03	0.030
2 km	Generalized Cost Beta	4.980e-03	2.589e-03	4.687e-03
1 km	Procedure Index Beta	0.073	5.968e-03	0.107
2 km	Procedure Index Beta	-0.036	-0.012	0.062
Boundary Points				
20	Generalized Cost Beta	8.747e-03	8.347e-03	8.870e-03
50	Generalized Cost Beta	3.485e-03	2.086e-03	6.655e-03
20	Procedure Index Beta	-0.111	-0.019	0.096
50	Procedure Index Beta	0.025	-1.123e-03	0.062

Table B.27: Pasadena and West Covina (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Distance-to-Court Beta	-1.360e-06	-4.747e-07	3.186e-05
procedure_core	Distance-to-Court Beta	-4.221e-06	-3.475e-06	5.590e-05
procedure_extended	Distance-to-Court Beta	-5.865e-06	-2.453e-06	3.705e-05
procedure_all	Procedure Index Beta	-0.016	-0.015	0.026
procedure_core	Procedure Index Beta	4.847e-03	7.360e-04	6.994e-03
procedure_extended	Procedure Index Beta	-0.025	-0.029	0.059
Outcome Spec				
case_default_flag	Distance-to-Court Beta	-1.402e-05	-7.455e-06	4.810e-05
imputed_unit_defaults	Distance-to-Court Beta	2.986e-06	1.665e-06	1.989e-05
case_default_flag	Procedure Index Beta	-0.025	-0.030	0.028
imputed_unit_defaults	Procedure Index Beta	-3.666e-03	-5.021e-03	0.019
First-Stage Method				
local_linear	Distance-to-Court Beta	-5.204e-06	-5.093e-06	7.085e-05
means	Distance-to-Court Beta	-2.426e-06	-2.049e-06	2.044e-05
local_linear	Procedure Index Beta	-0.011	-8.157e-03	0.034
means	Procedure Index Beta	-0.013	-7.862e-03	0.019
Local Bandwidth				
0.5 km	Distance-to-Court Beta	-3.740e-05	-3.373e-05	2.083e-04
1 km	Distance-to-Court Beta	1.162e-05	6.527e-06	4.982e-05
2 km	Distance-to-Court Beta	-5.938e-06	-3.928e-06	2.545e-05
0.5 km	Procedure Index Beta	0.039	0.040	0.113
1 km	Procedure Index Beta	-0.021	-7.604e-03	0.031
2 km	Procedure Index Beta	-0.016	-0.011	0.019
Boundary Points				
10	Distance-to-Court Beta	-8.319e-06	2.902e-06	4.191e-05
20	Distance-to-Court Beta	5.512e-06	-1.169e-06	2.071e-05
50	Distance-to-Court Beta	-7.610e-06	-3.928e-06	4.810e-05
10	Procedure Index Beta	-0.013	-8.592e-03	0.022
20	Procedure Index Beta	-0.015	-7.427e-03	0.021
50	Procedure Index Beta	-0.010	-8.485e-03	0.028

Table B.28: Santa Monica and Stanley Mosk (Distance Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is distance-to-court.

	Beta Coefficient	Average Estimate	Median Estimate	Median CI Width
Procedure Spec				
procedure_all	Generalized Cost Beta	3.030e-04	-3.009e-05	1.867e-03
procedure_core	Generalized Cost Beta	5.759e-04	-1.520e-04	3.654e-03
procedure_extended	Generalized Cost Beta	5.035e-04	1.219e-04	2.070e-03
procedure_all	Procedure Index Beta	-0.014	-8.875e-03	0.024
procedure_core	Procedure Index Beta	4.760e-03	-1.815e-04	8.011e-03
procedure_extended	Procedure Index Beta	-0.021	-0.028	0.054
Outcome Spec				
case_default_flag	Generalized Cost Beta	1.838e-03	1.314e-03	3.762e-03
imputed_unit_defaults	Generalized Cost Beta	-4.571e-04	-3.115e-04	1.743e-03
case_default_flag	Procedure Index Beta	-0.029	-0.032	0.029
imputed_unit_defaults	Procedure Index Beta	2.424e-03	-2.294e-03	0.018
First-Stage Method				
local_linear	Generalized Cost Beta	7.980e-04	8.727e-05	3.183e-03
means	Generalized Cost Beta	1.236e-04	-9.926e-05	1.975e-03
local_linear	Procedure Index Beta	-6.241e-03	-7.605e-03	0.034
means	Procedure Index Beta	-0.014	-7.513e-03	0.021
Local Bandwidth				
0.5 km	Generalized Cost Beta	7.946e-05	-6.046e-04	9.838e-03
1 km	Generalized Cost Beta	4.128e-04	-1.284e-04	3.612e-03
2 km	Generalized Cost Beta	5.483e-04	2.596e-04	1.777e-03
0.5 km	Procedure Index Beta	0.046	0.043	0.101
1 km	Procedure Index Beta	-0.010	-7.140e-03	0.036
2 km	Procedure Index Beta	-0.020	-0.015	0.018
Boundary Points				
10	Generalized Cost Beta	9.785e-04	4.920e-04	5.710e-03
20	Generalized Cost Beta	4.563e-05	-1.767e-04	2.333e-03
50	Generalized Cost Beta	5.028e-04	-1.069e-04	1.940e-03
10	Procedure Index Beta	-0.021	-9.925e-03	0.023
20	Procedure Index Beta	-9.632e-03	-2.691e-03	0.026
50	Procedure Index Beta	-5.930e-03	-8.548e-03	0.027

Table B.29: Santa Monica and Stanley Mosk (Cost Specifications)

Note: Average non-standardized beta coefficient estimates are given by different specification settings. The measure of tenant costs is these specifications is generalized cost.

Appendix C

SUPPLEMENTARY MATERIAL FOR CHAPTER 3

C.1 Sharp RDD Constant Treatment Effect

Appendix C.1 gives regression results for the constant treatment effect sharp RDD example.

	Dependent variable:	
	Y	
	(1)	(2)
X-3	3.268***	2.893***
	(2.981,3.555)	(2.552,3.234)
Treated	7.000***	7.000***
	(6.290,7.710)	(6.713,7.287)
(X-3):Treated	0.000	−0.000
	(-0.406,0.406)	(-0.482,0.482)
Constant	4.950***	2.232***
	(4.448,5.452)	(2.029,2.434)
Observations	200	68
R^2	0.875	0.978
Adjusted R^2	0.873	0.977
F Statistic	456.389*** (df = 3; 196)	948.619*** (df = 3; 64)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table C.1: Linear Regression Results

As seen above, the estimates on "Treated" recover the true treatment effect of 7 from the example plot. Appendix Table C.1 reports estimates of the local treatment effect. The first column uses all observations, whereas the second column uses only observations with $X_i \in [2, 4]$, which is the local linear regression plotted in thick

dark red and dark blue lines between the dashed vertical lines in Figure 3.1. In both columns, we regress Y_i on $1, (X_i - 3), W_i, W_i(X_i - 3)$ and recover the true $\tau_C = 7.0$; however, the local linear regression (Column (2)) gives smaller CIs on τ_C.

C.2 Extrapolation to ATE

Appendix C.2 considers extrapolation to the ATE. For simplicity, we focus on the case where the running variable X is discrete.

Extrapolation: Proof of the ATE Bounds

We first obtain the upper bound by applying the LIE, the global monotonicity assumption, and observing the following:

$$\begin{aligned}
ATE &= \mathbb{E}[Y_1 - Y_0] \\
&= \sum_{x \in \text{supp}(X)} Pr(X = x) \cdot \mathbb{E}[Y_1 - Y_0 | X = x] \\
&= \sum_{x<C} Pr(X = x) \cdot \mathbb{E}[Y_1 - Y_0 | X = x] + \sum_{x \geq C} Pr(X = x) \cdot \mathbb{E}[Y_1 - Y_0 | X = x] \\
&\leq \sum_{x<C} Pr(X = x) \cdot (\mathbb{E}[Y_1 | W = 1, X = C] - \mathbb{E}[Y_0 | W = 0, X = x]) \\
&\quad + \sum_{x \geq C} Pr(X = x) \cdot \left(\mathbb{E}[Y_1 | W = 1, X = x] - \lim_{x \to C^-} \mathbb{E}[Y_0 | W = 0, X = x] \right) \\
&\leq \underbrace{\sum_{x<C} Pr(X = x) \cdot \left(\mu_1^R(C) - \mu_0^L(x) \right) + \sum_{x \geq C} Pr(X = x) \cdot \left(\mu_1^R(x) - \lim_{x \to C^-} \mu_0^L(x) \right)}_{= \overline{ATE}}
\end{aligned}$$

where we used that $\mu_1(x), \mu_0(x)$ are increasing functions of x. Note as well that $\mu_1^R(x)$ for $x \geq C$ is observable by the sharp assignment rule; and similarly for $\mu_0^L(x)$ for $x < C$.

The lower bound on ATE uses boundedness and monotonicity. Again, we apply the LIE and the assumptions:

$$\begin{aligned}
ATE &= \mathbb{E}[Y_1 - Y_0] \\
&= \sum_{x<C} Pr(X = x) \cdot \mathbb{E}[Y_1 - Y_0 | X = x] + \sum_{x \geq C} Pr(X = x) \cdot \mathbb{E}[Y_1 - Y_0 | X = x] \\
&\geq \underbrace{\sum_{x<C} Pr(X = x) \cdot (M_0 - \mu_0^L(x)) + \sum_{x \geq C} Pr(X = x) \cdot (\mu_1^R(x) - M_1)}_{= \underline{ATE}}
\end{aligned}$$

where we use that $\mu_1(x)$ is bounded below by M_0 for $x < C$ and that $\mu_0(x)$ is bounded above by M_1 for $x > C$.

Thus, the upper and lower bounds for the ATE are given by the quantities:

$$\underline{ATE} \leq ATE \leq \overline{ATE}$$

where

$$\overline{ATE} = \sum_{x<C} Pr(X = x) \cdot \left(\mu_1^R(C) - \mu_0^L(x)\right) + \sum_{x \geq c} Pr(X = x) \cdot \left(\mu_1^R(x) - \lim_{x \to C^-} \mu_0^L(x)\right)$$
$$\underline{ATE} = \sum_{x<C} Pr(X = x) \cdot (M_0 - \mu_0^L(x) + \sum_{x \geq C} Pr(X = x) \cdot (\mu_1^R(x) - M_1)$$

ATE: Bounding Example

To fix ideas, I also include an example of bounding the ATE. In the example, the forcing variable X is discrete and the CEFs are globally monotone increasing. I consider Figure C.1 below, which demonstrates the idea for the upper bound on ATE. First, in the top, left-most plot, we display the same CEFs for the treated (blue) and control (maroon) groups as above in Figure 3.2. The specific CEFs defined on $\text{supp}(X) = [0, 6]$ are recorded below again:

$$\mu_1(x) = \mathbf{1}\{x < 3\}(0.2x^2 + 2) + \mathbf{1}\{x \geq 3\}(0.7x^2 + 2)$$
$$\mu_0(x) = \mathbf{1}\{x < 3\}(0.2x^2 - 2) + \mathbf{1}\{x \geq 3\}(0.7x^2 - 5)$$

The treatment assignment cutoff occurs at $C = 3$. The solid lines are observed: because of the assignment rule $W_i = \mathbf{1}\{X_i \geq 3\}$, $\mu_1(x) = \mathbb{E}[Y_1|X = x]$ is observable for $x \geq 3$, while $\mu_0(x) = \mathbb{E}[Y_0|X = x]$ is observable for $x < 3$.[1] The dashed lines are the true values for the CEFs, but they are unobservable.

As before, note that both CEFs are discontinuous at the regression cutoff value $C = 3$. The typical RDD assumption that both conditional means $\mu_1(x), \mu_0(x)$ are continuous in x **does not hold.** However, under (global) monotonicity we can obtain an upper bound on the ATE. Consider the top, right-most panel in Figure C.1: the dotted lines for each potential outcome CEF are those that are compatible with monotonicity for the purposes of calculating an upper bound on the ATE.

Since we know the DGP, we can calculate the upper bound on the ATE as follows. First, we assume the X are uniformly distributed on $[0, 6]$ and we calculate the highest values of $\mu_1(x)$ for each x and the lowest values of $\mu_0(x)$ that are compatible with monotonicity and the observed data. The ATE is the difference between these values integrated over the support of the forcing variable. In our Figure C.1 example, this means, for $x \in [0, 3)$, taking the difference between the dotted blue line and the

[1] Again, we use "observable" to mean "can be estimated consistently from the data."

solid maroon line; whereas, for $x \in [3, 6]$, we take the difference between the solid blue line and the dotted maroon line. This results in an upper bound for the ATE of:

$$\begin{aligned}\overline{ATE} &= \frac{1}{6}\int_{x=0}^{3}(8.3 - (0.2x^2 - 2))dx + \frac{1}{6}\int_{x=3}^{6}(0.7x^2 + 2 - (-0.2))dx \\ &= \frac{1}{6}(29.1 + 50.7) = \boxed{13.3}\end{aligned}$$

Note that this can also be done for RDD cases where there are only a few discrete values that the forcing variable X can take. For example, in the bottom panel of Figure C.1, we keep the same CEFs but suppose we have only 6 support points on each side of the cutoff C. The values for $\mathbb{E}[Y|X = x]$ are displayed above each point in the panel. Supposing that X is uniformly distributed, this gives an estimated upper bound on the ATE of 12.98, which is close to the 13.3 above in the continuously distributed X scenario.

Compare this to the true ATE, which is, for $x \in [0, 3)$, given by taking the difference between the dashed blue line and the solid maroon line; whereas, for $x \in [3, 6]$, we take the difference between the solid blue line and the dashed maroon line, and calculated as follows:

$$\begin{aligned}ATE &= \frac{1}{6}\left(\int_{x=0}^{3}\Big((0.2x^2 + 2) - (0.2x^2 - 2)\Big)\,dx + \int_{x=3}^{6}\Big((0.7x^2 + 2) - (0.7x^2 - 5)\Big)\,dx\right) \\ &= \frac{1}{6}(3 * 4 + 3 * 7) = \boxed{5.5}\end{aligned}$$

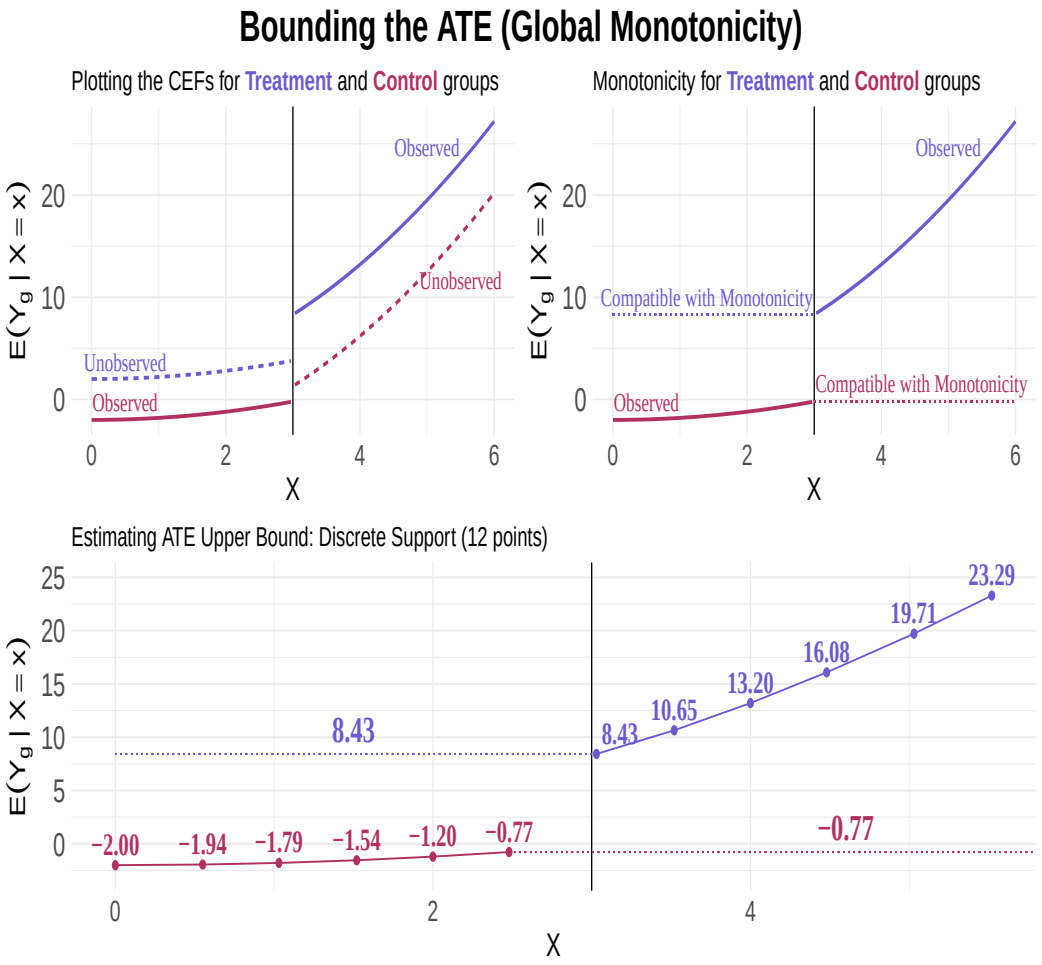

Figure C.1: Bounding the ATE Under Global Monotonicity

Note: The top left panel shows the DGP. The top right panel shows the imputed values for each CEF compatible with monotonicity for the purposes of computing the upper bound on ATE. The bottom panel is included for a computational example: using these values there, and if the forcing variable X is uniformly distributed, the ATE upper bound is approximately 12.98.

C.3 Data Generating Processes (DGPs)

Appendix C.3 gives more information on the data generating processes used in the simulation study. Each plot belows shows the true potential outcome CEFs, the population data with added noise, and the observable population CEF.

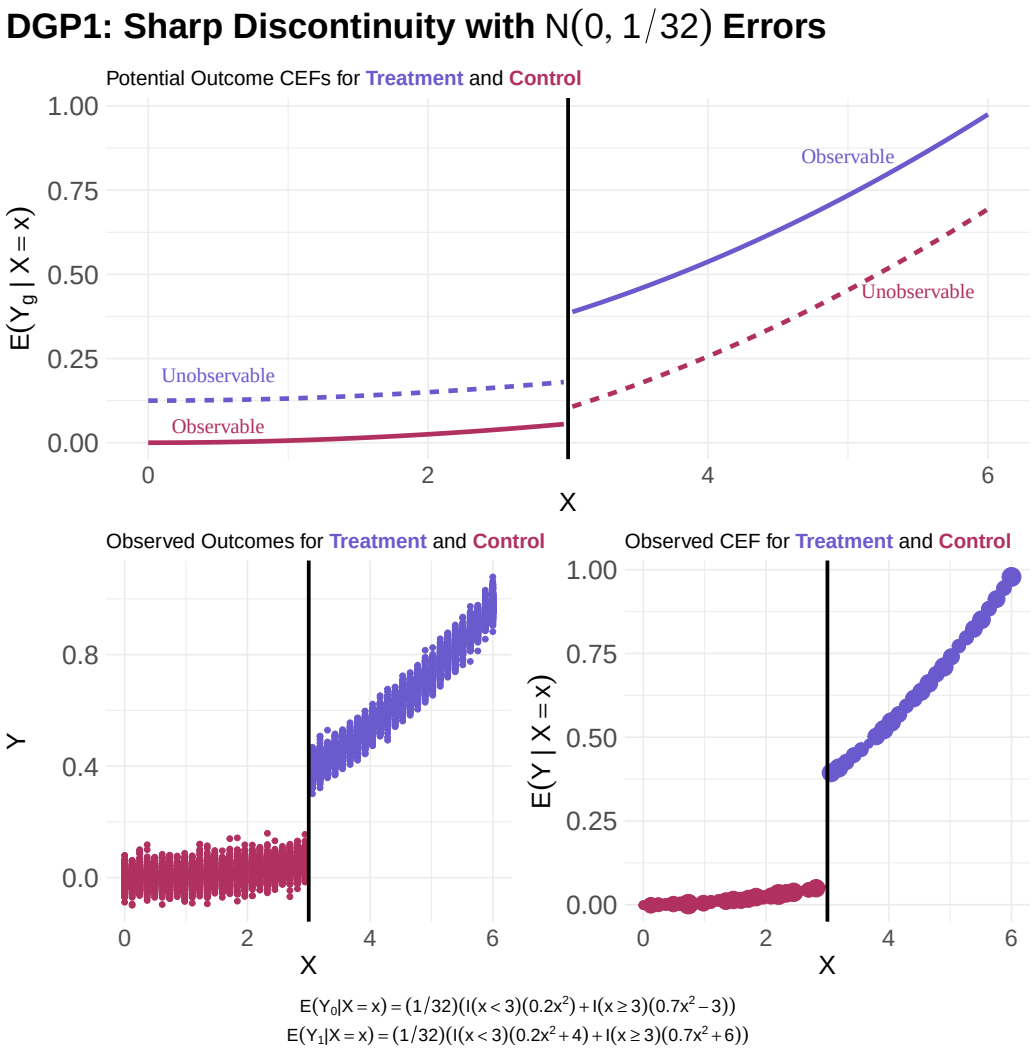

Figure C.2: Data Generating Process 1 (DGP1)

Note: By the sharp design, the solid lines of the potential outcome CEFs are observable for treated units (control units) to the right (left) of the cutoff $C = 3$. The dashed lines are the unobservable part of the potential outcome CEFs because of the sharp design. The equations for the potential outcome CEFs are given below panels, where I() is the indicator function. The true local treatment effect at $C = 3$ is approximately 0.281. The bottom-left panel shows that there are 50 unique X values in the simulated dataset. In the bottom-right panel, the volume of the points is proportional to the share of observations with $X = x$.

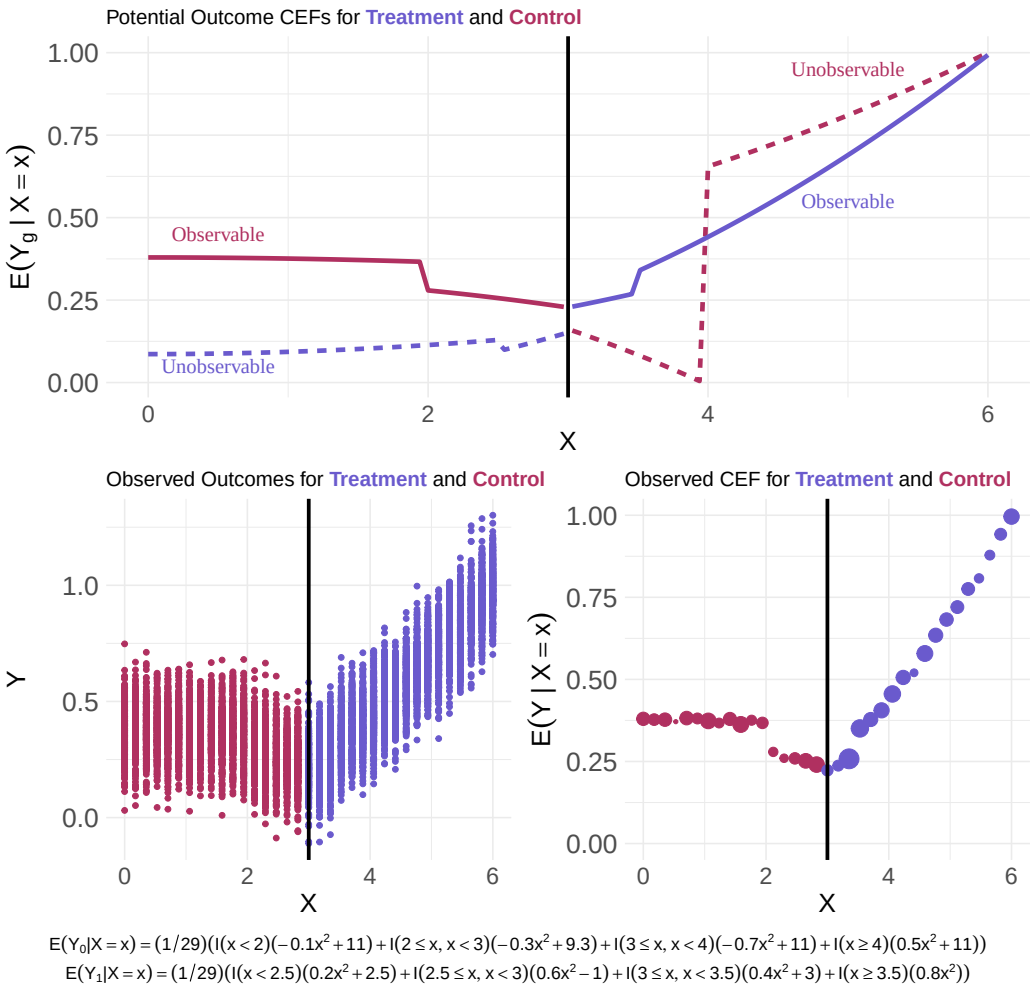

Figure C.3: Data Generating Process 2 (DGP2)

Note: By the sharp design, the solid lines of the potential outcome CEFs are observable for treated units (control units) to the right (left) of the cutoff $C = 3$. The dashed lines are the unobservable part of the potential outcome CEFs because of the sharp design. The equations for the potential outcome CEFs are given below panels, where I() is the indicator function. The true local treatment effect at $C = 3$ is approximately 0.066. The bottom-left panel shows that there are 35 unique X values in the simulated dataset. In the bottom-right panel, the volume of the points is proportional to the share of observations with $X = x$.

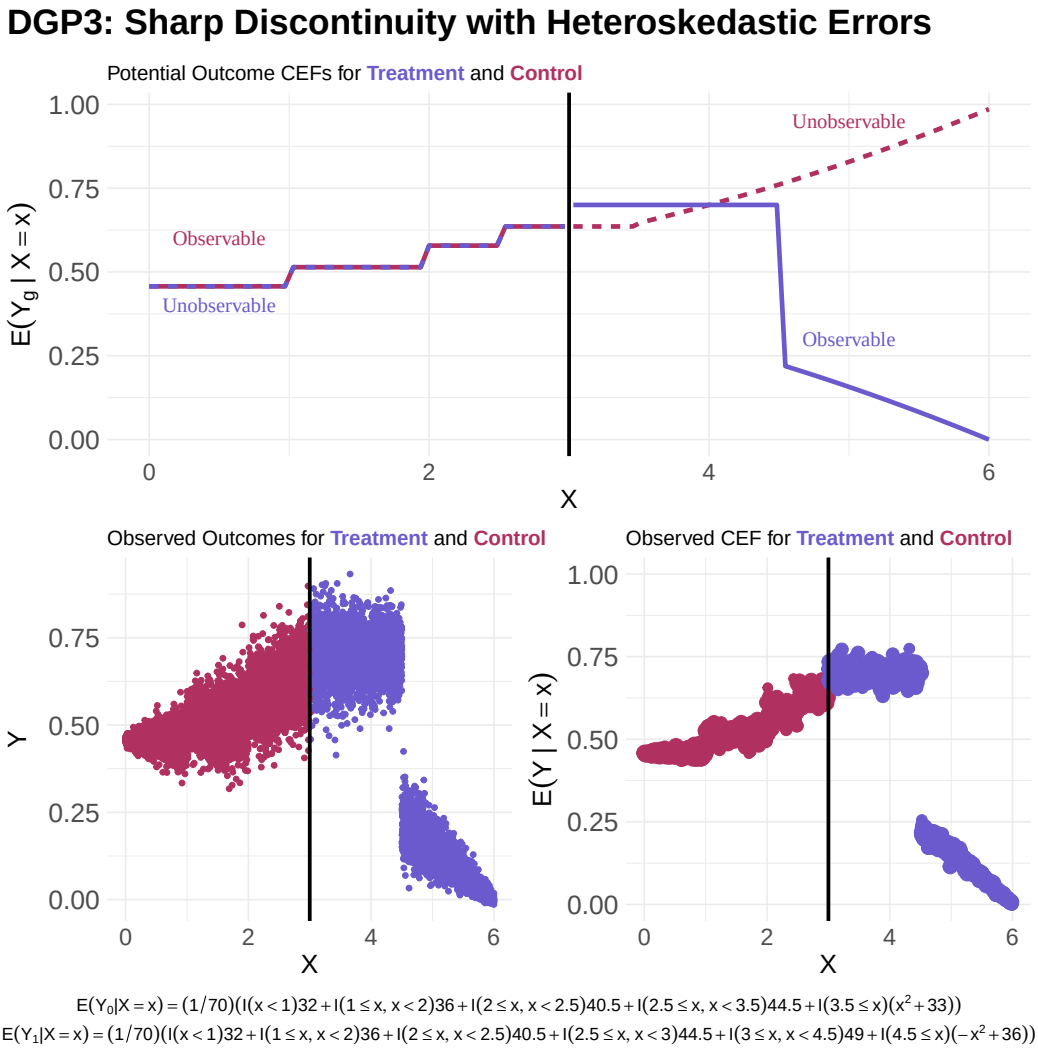

Figure C.4: Data Generating Process 3 (DGP3)

Note: By the sharp design, the solid lines of the potential outcome CEFs are observable for treated units (control units) to the right (left) of the cutoff $C = 3$. The dashed lines are the unobservable part of the potential outcome CEFs because of the sharp design. The equations for the potential outcome CEFs are given below panels, where I() is the indicator function. The true local treatment effect at $C = 3$ is approximately 0.064. The bottom-left panel shows that there are 1000 unique X values in the simulated dataset. In the bottom-right panel, the volume of the points is proportional to the share of observations with $X = x$.

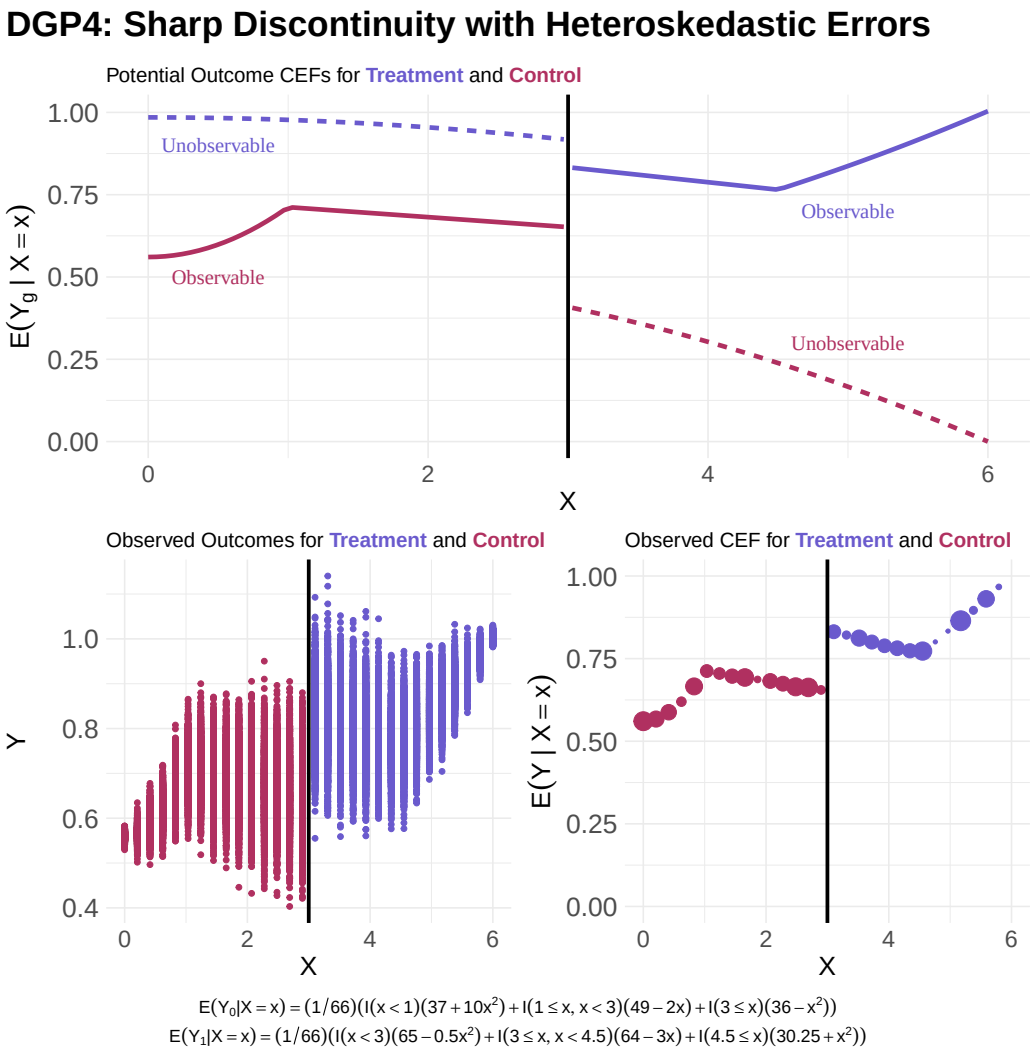

Figure C.5: Data Generating Process 4 (DGP4)

Note: By the sharp design, the solid lines of the potential outcome CEFs are observable for treated units (control units) to the right (left) of the cutoff $C = 3$. The dashed lines are the unobservable part of the potential outcome CEFs because of the sharp design. The equations for the potential outcome CEFs are given below panels, where I() is the indicator function. The true local treatment effect at $C = 3$ is approximately 0.424. The bottom-left panel shows that there are 30 unique X values in the simulated dataset. In the bottom-right panel, the volume of the points is proportional to the share of observations with $X = x$.

C.4 Reverse Kernel Examples

Appendix C.4 shows further examples of reverse kernels. The top plots are the kernels used (weights). The middle plots are two unknown functions with estimates based on the above kernels: the solid red (reverse triangular) and solid blue (reverse parabolic) lines are the reverse kernel estimates, whereas the dashed yellow (triangular) and dashed green (parabolic) lines are the standard kernel estimates. For viewability, we use the triangular kernels in the left-side plots and the parabolic kernels in the right-side plots. The bottom two plots are kernel estimates by bandwidth for the increasing (bottom-left) and decreasing (bottom-right) DGPs, and for viewability we only include either the parabolic and reverse parabolic (resp. triangular and reverse triangular) estimates in the bottom-left (resp. bottom-right) plot. The bandwidths (left h=0.25 and right h=0.50) are indicated by dashed vertical lines in the middle and bottom panels.

Reverse Kernel Estimators

Bounding the Boundaries of Monotone Functions

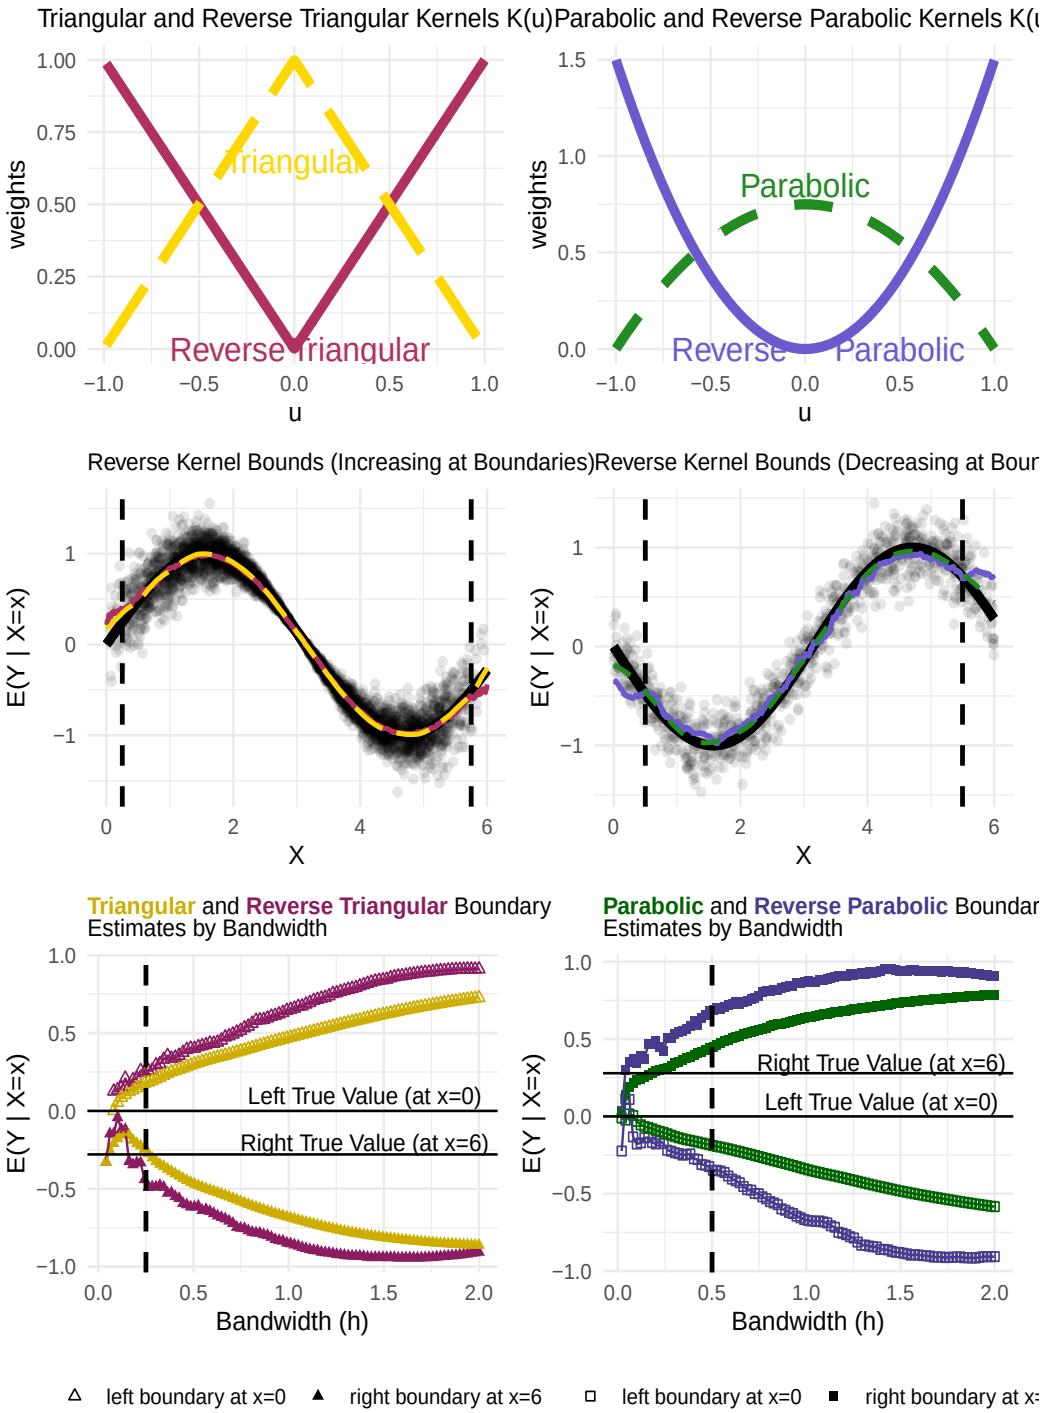

Figure C.6: Reverse Kernel Estimation

C.5 Local Simulation Study Results

Appendix C.5 shows the local simulation results as density plots for each DGP and both the upper and lower bounds.

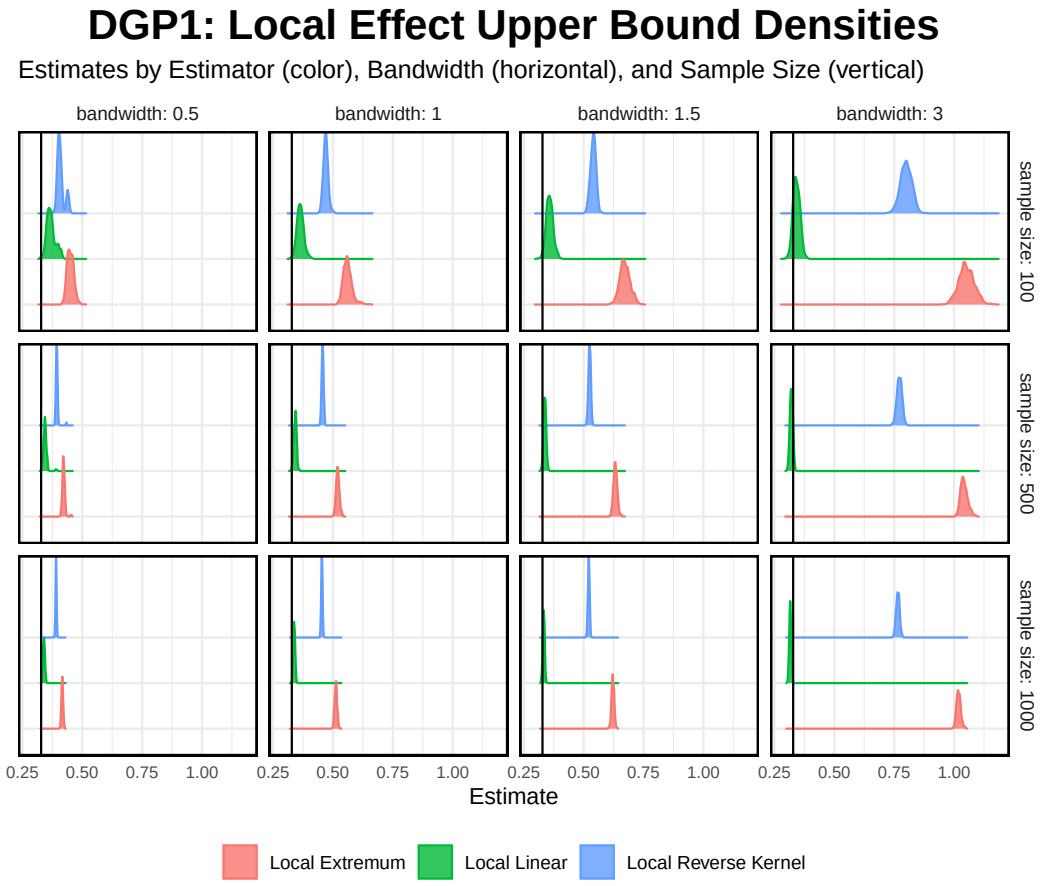

Figure C.7: DGP1: Local Effect Upper Bound Simulation Results

Note: The true local upper bound is given by the black vertical line in each panel.

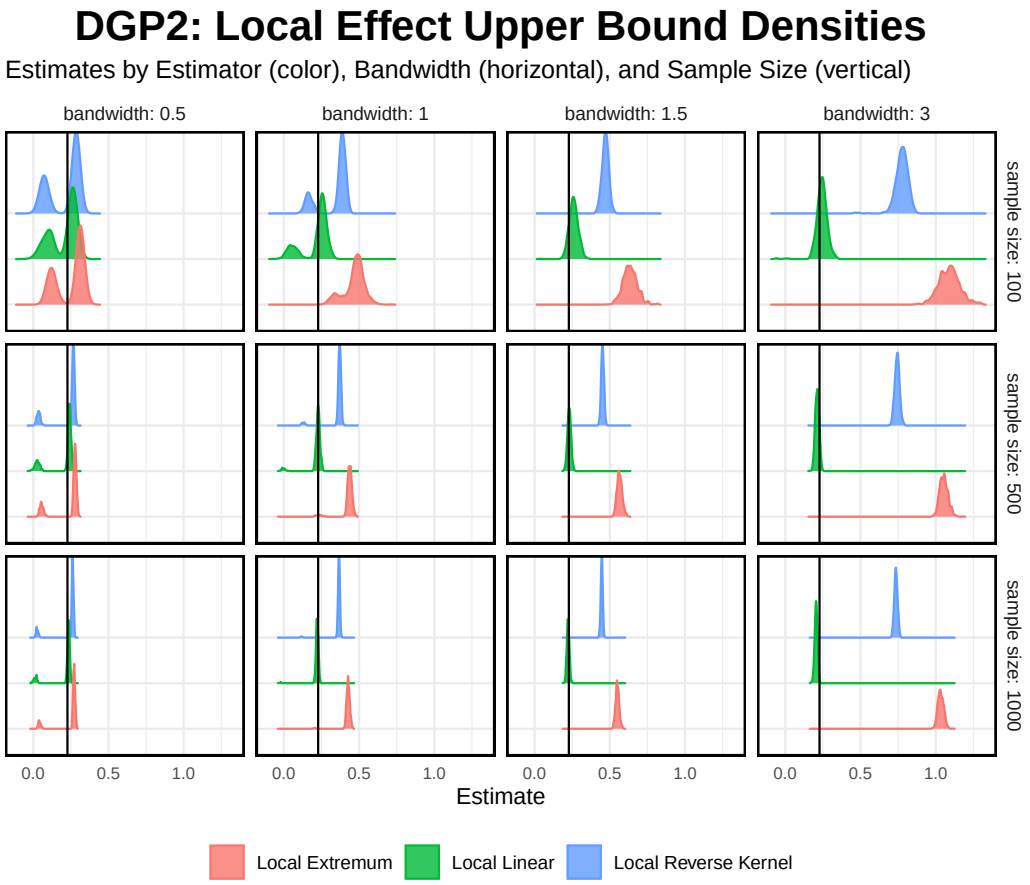

Figure C.8: DGP2: Local Effect Upper Bound Simulation Results

Note: The true local upper bound is given by the black vertical line in each panel.

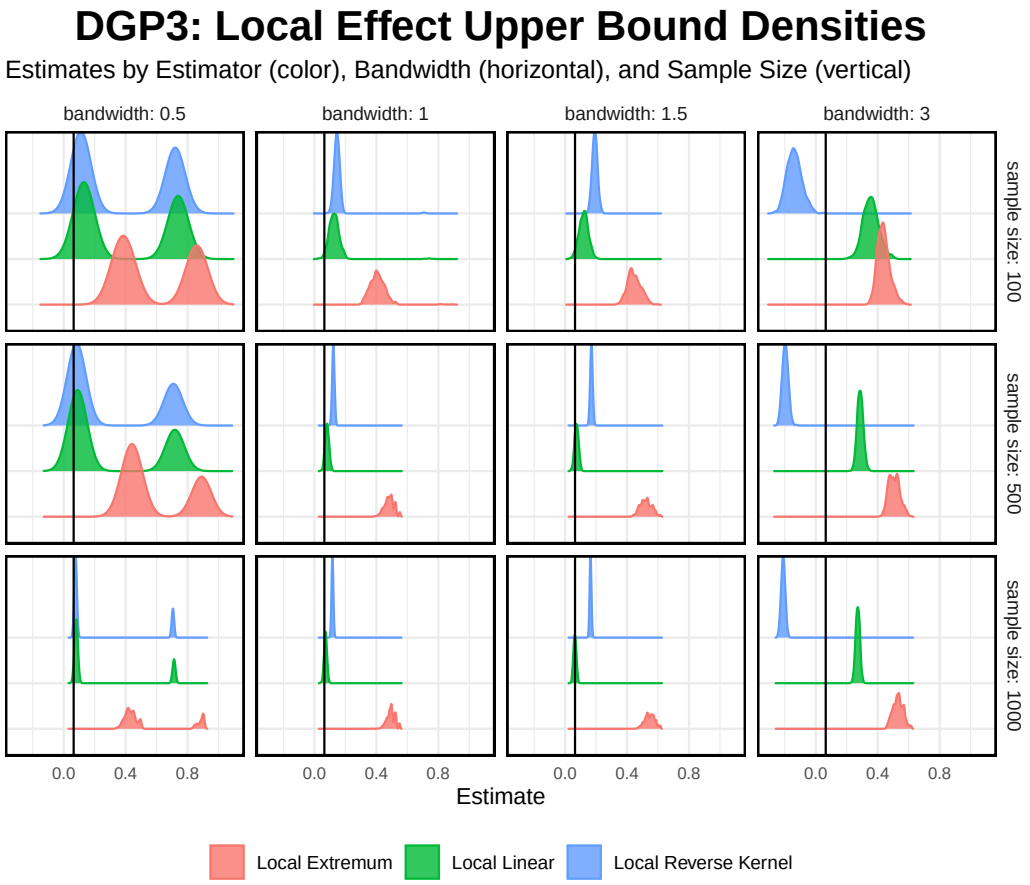

Figure C.9: DGP3: Local Effect Upper Bound Simulation Results

Note: The true local upper bound is given by the black vertical line in each panel.

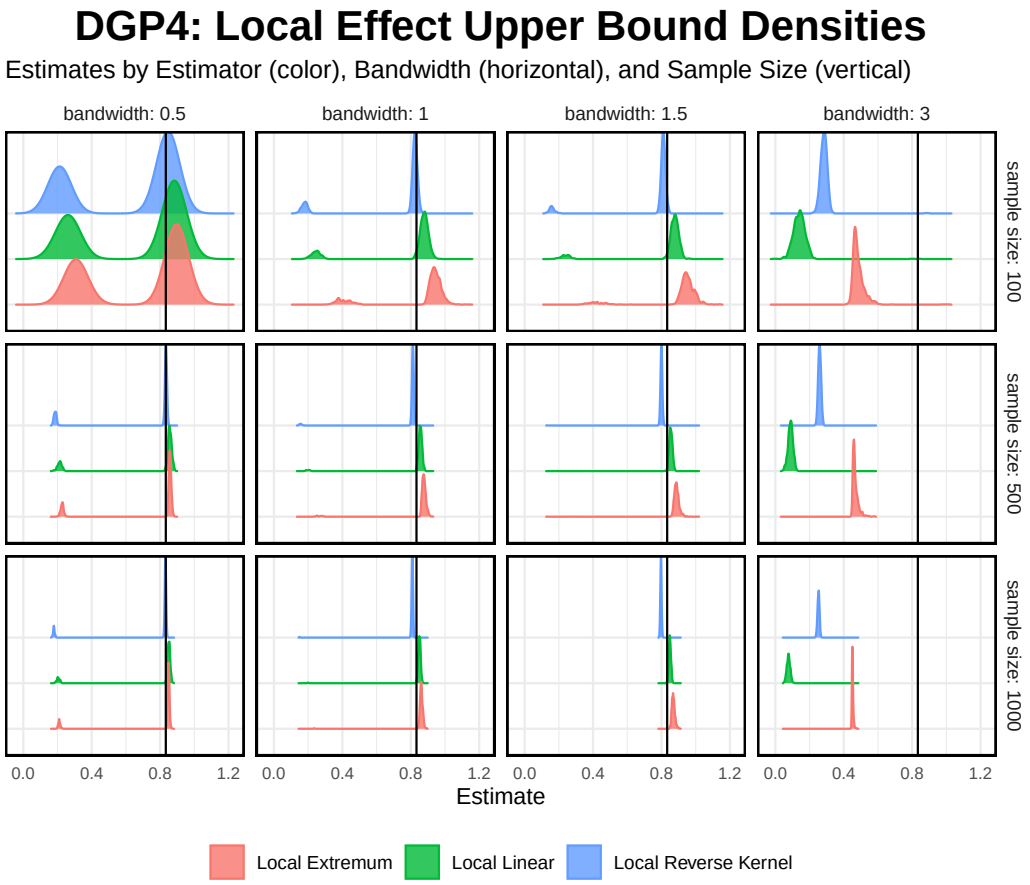

Figure C.10: DGP4: Local Effect Upper Bound Simulation Results

Note: The true local upper bound is given by the black vertical line in each panel.

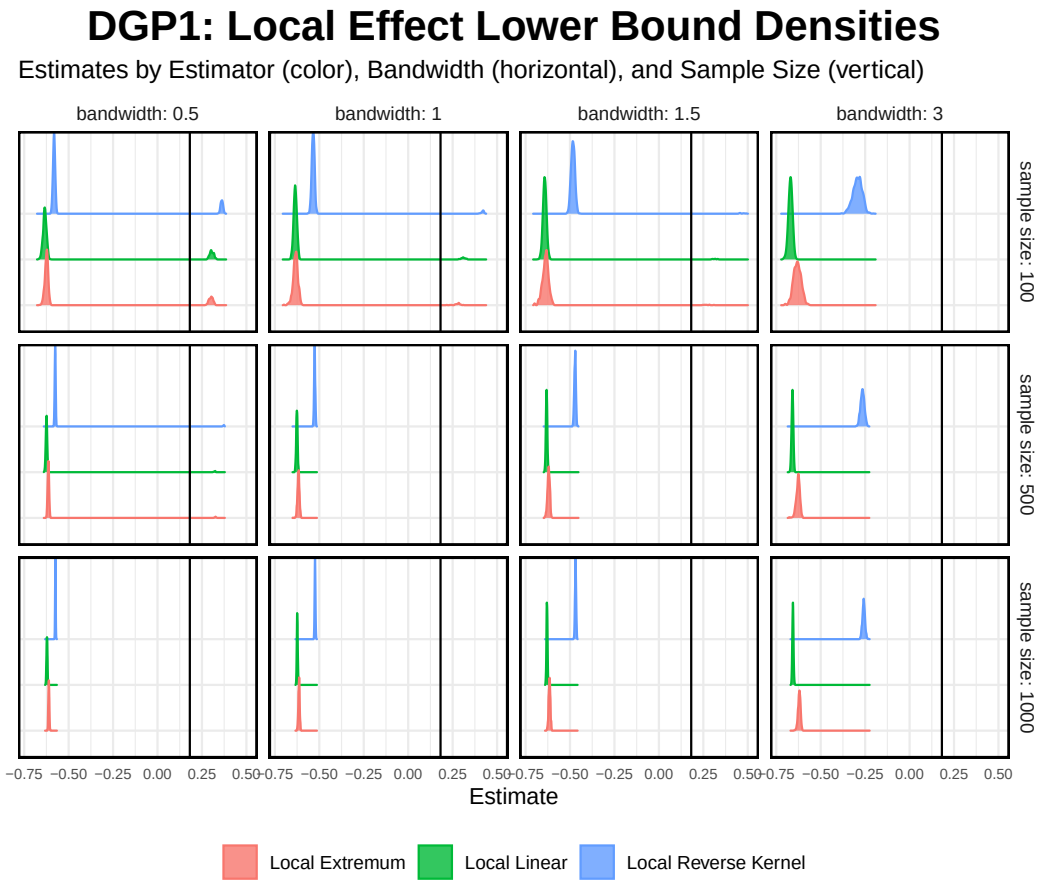

Figure C.11: DGP1: Local Effect Lower Bound Simulation Results

Note: The true local lower bound is given by the black vertical line in each panel.

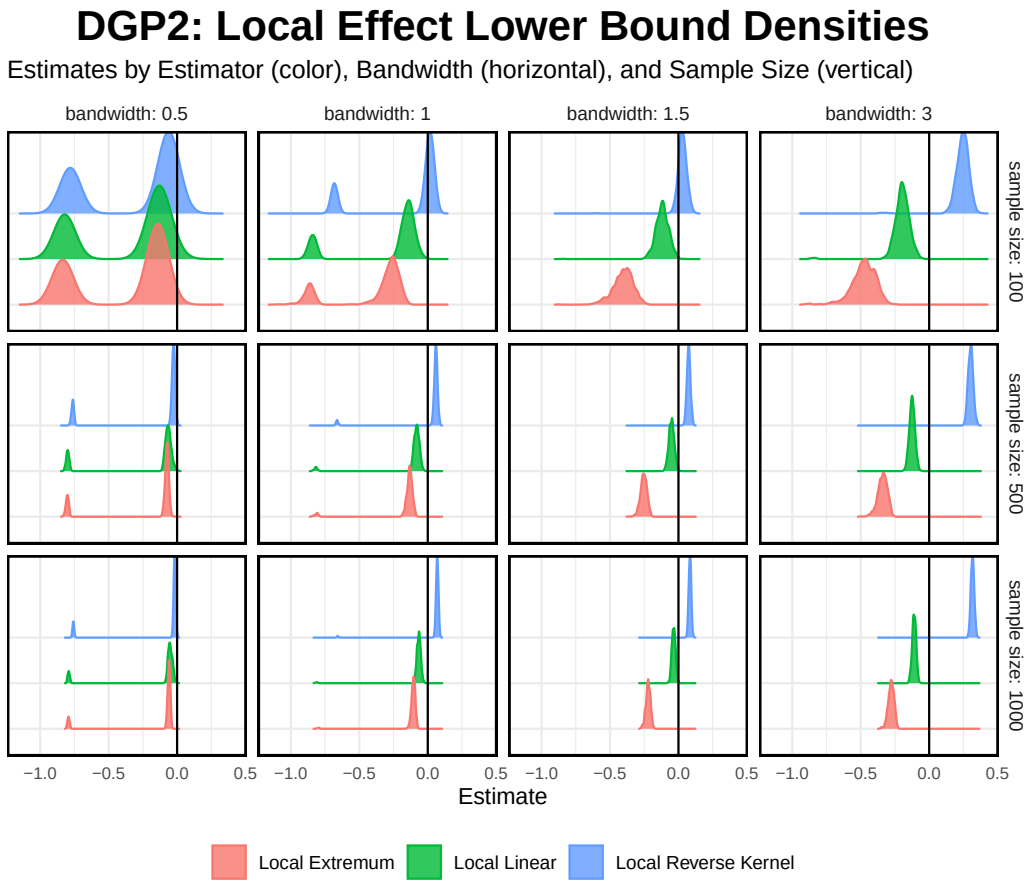

Figure C.12: DGP2: Local Effect Lower Bound Simulation Results

Note: The true local lower bound is given by the black vertical line in each panel.

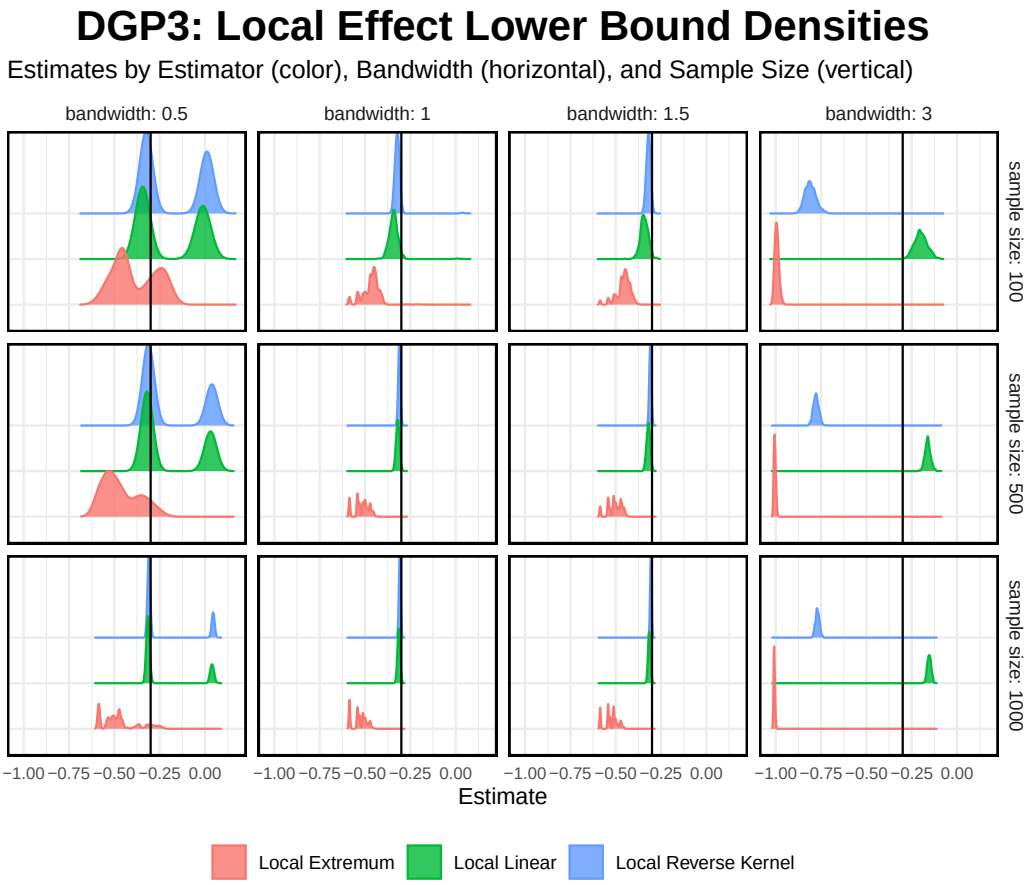

Figure C.13: DGP3: Local Effect Lower Bound Simulation Results

Note: The true local lower bound is given by the black vertical line in each panel.

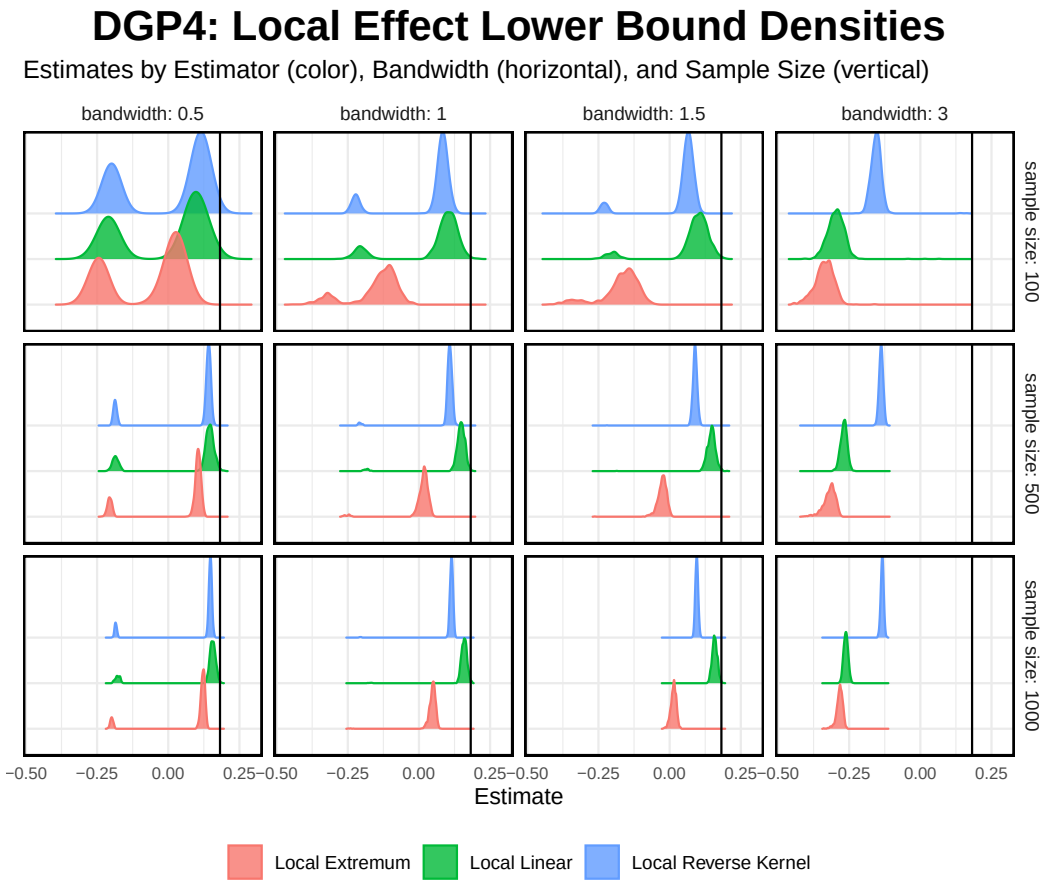

Figure C.14: DGP4: Local Effect Lower Bound Simulation Results

Note: The true local lower bound is given by the black vertical line in each panel.

C.6 Local Simulation Study Coverage Results

Appendix C.6 shows coverage results for the local simulation study in plots rather than a table.

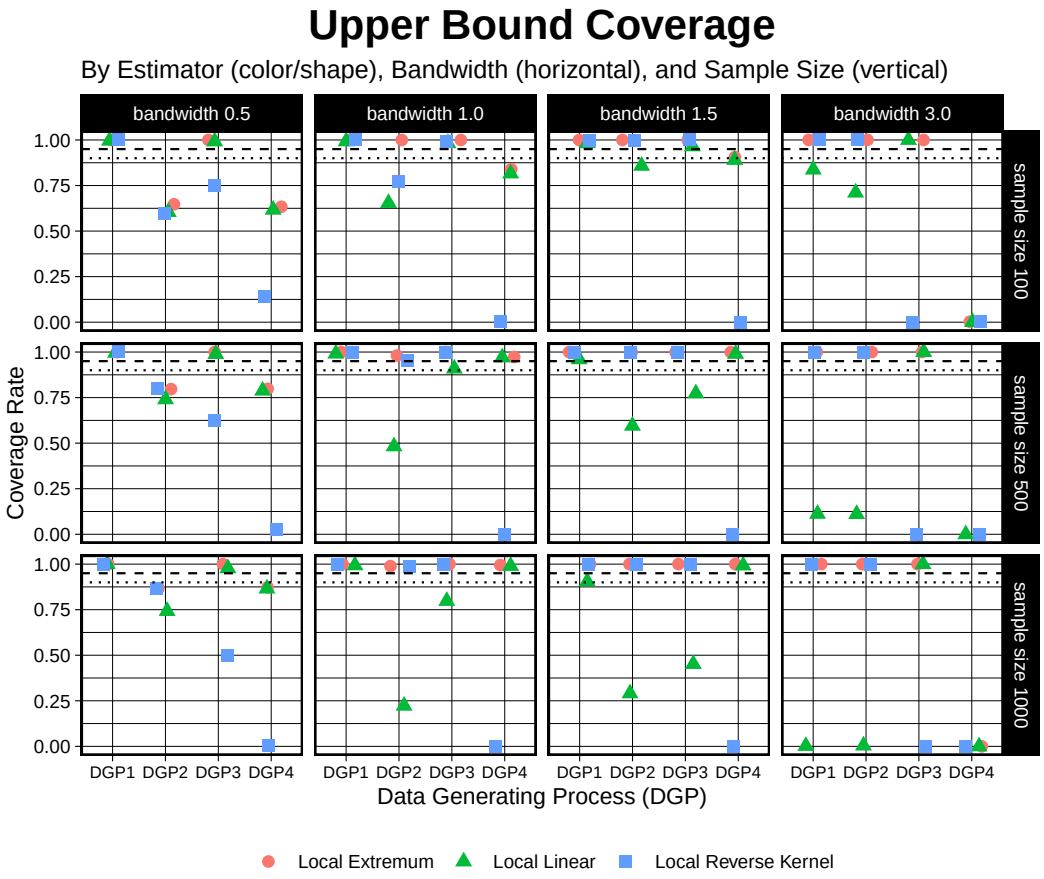

Figure C.15: Upper Bound Coverage

Note: The coverage rate for the upper bound is the number of simulation runs where estimated upper bound CIs for a given estimator include the true upper bound. For each bandwidth and sample size pair, there were 1000 simulation runs.

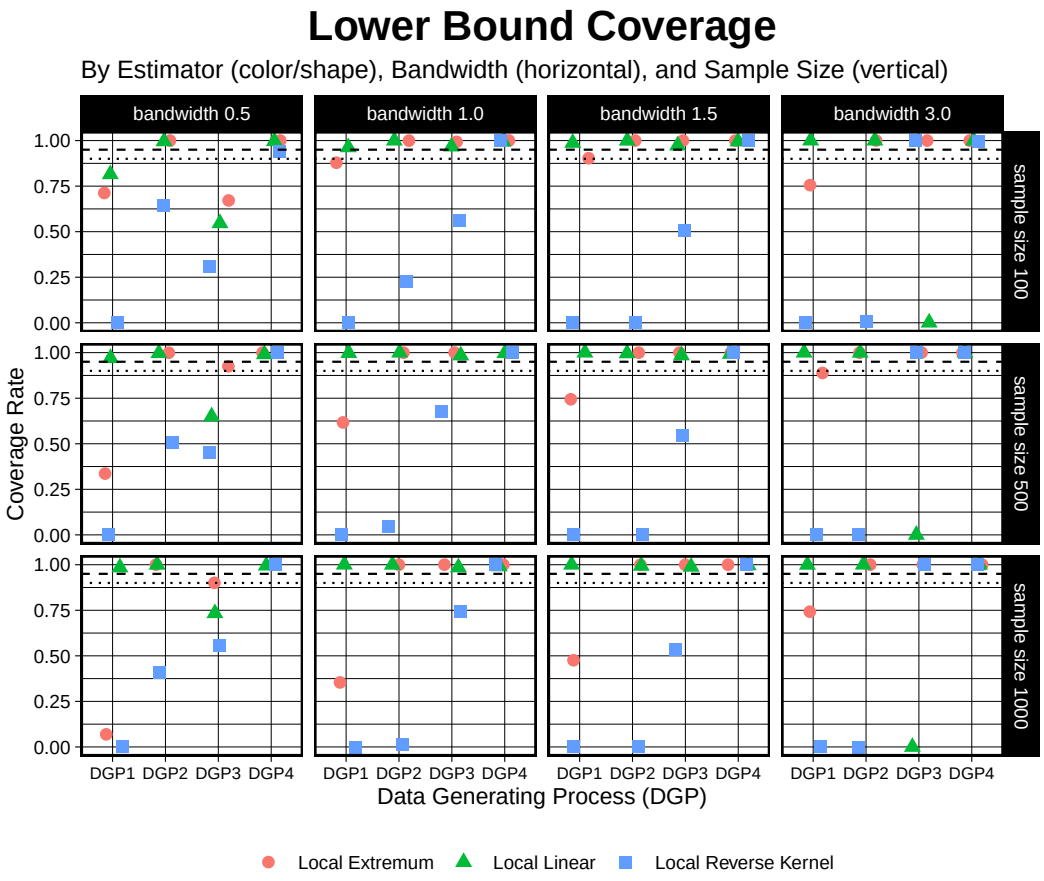

Figure C.16: Lower Bound Coverage

Note: The coverage rate for the lower bound is the number of simulation runs where estimated lower bound CIs for a given estimator include the true lower bound. For each bandwidth and sample size pair, there were 1000 simulation runs.

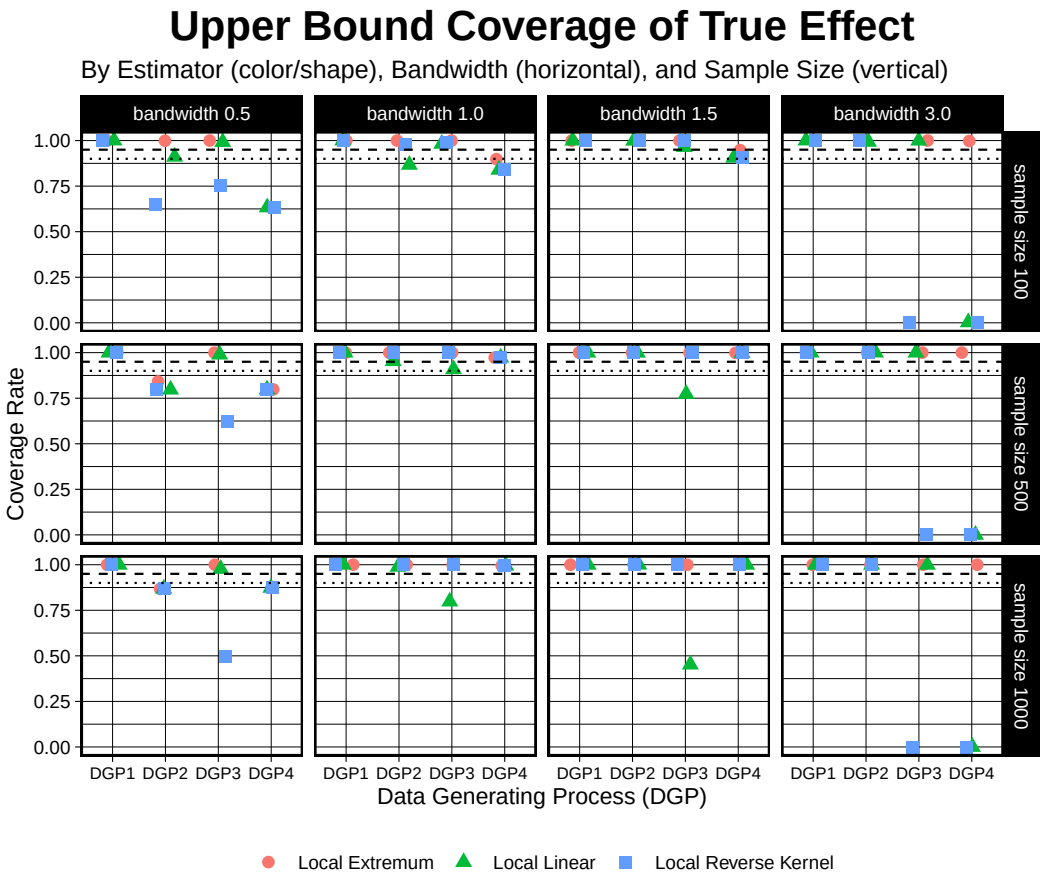

Figure C.17: Upper Bound Coverage of True Effect

Note: The upper bound coverage of the true effect rate is the number of simulation runs where the estimated upper bound CI for a given estimator is greater than or equal to the true treatment effect. For each bandwidth and sample size pair, there were 1000 simulation runs.

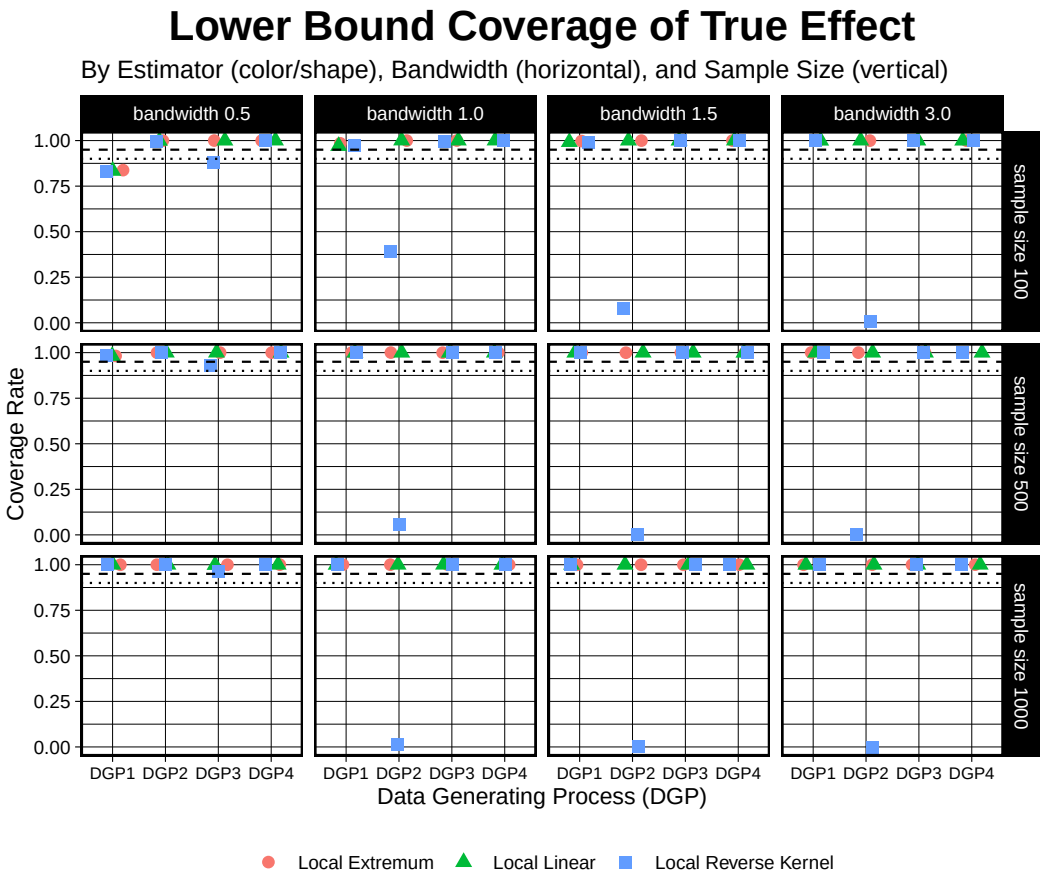

Figure C.18: Lower Bound Coverage of True Effect

Note: The lower bound coverage of the true effect rate is the number of simulation runs where the estimated lower bound CI for a given estimator is less than or equal to the true treatment effect. For each bandwidth and sample size pair, there were 1000 simulation runs.

C.7 Additional Simulation Results

Appendix C.7 shows additional simulation results, including: baseline simulations with more bandwidths, missing data simulations, and measurement error simulations.

More Bandwidths Simulations

The more bandwidths simulation use sample sizes of $n = 1000$ and 1000 simulation runs for each DGP, as specified above. Then I re-run the local effect simulations as before with more bandwidths in 0.10 increments from, e.g., $h = 0.20, 0.30, \ldots, 3.00$ for DGP1. Coverage rates (for upper bounds, lower bounds, and both bounds) and partial ID region widths are shown by estimator, bandwidth, and DGP.

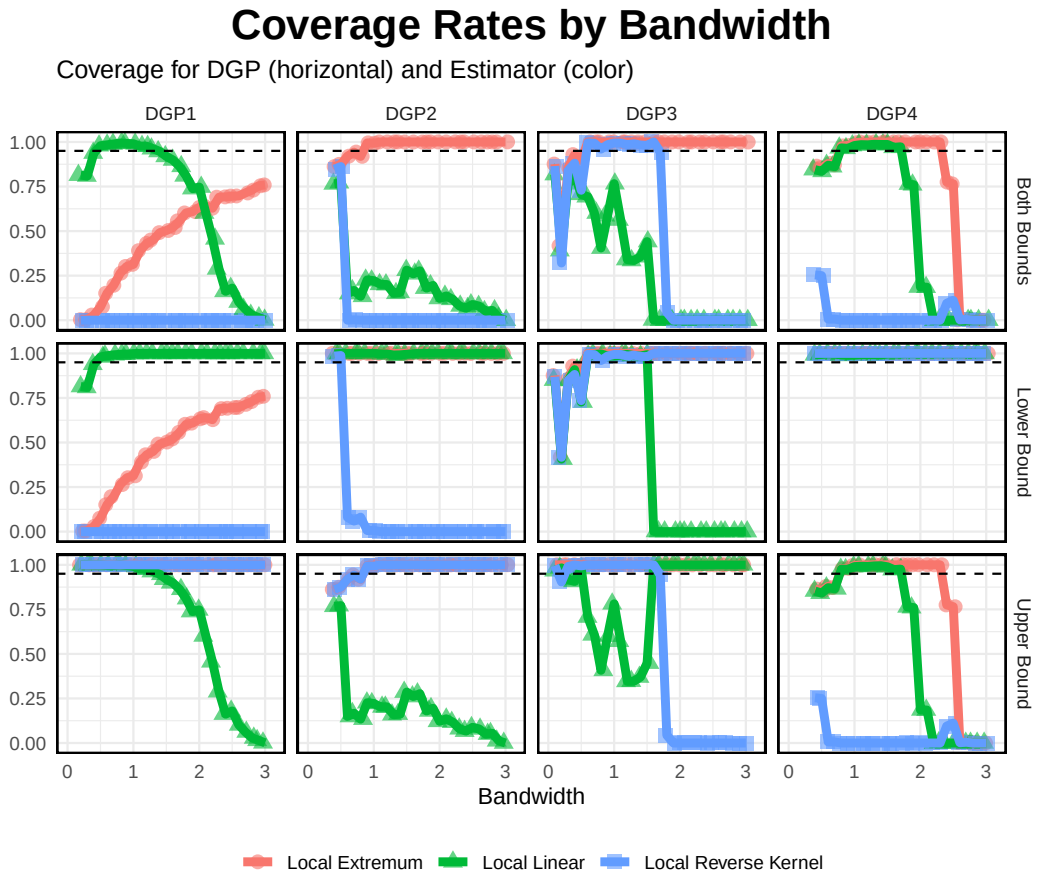

Figure C.19: More Bandwidths: Coverage Rates

Note: The 95% nominal coverage rate is plotted as the dashed black horizontal line in each panel.

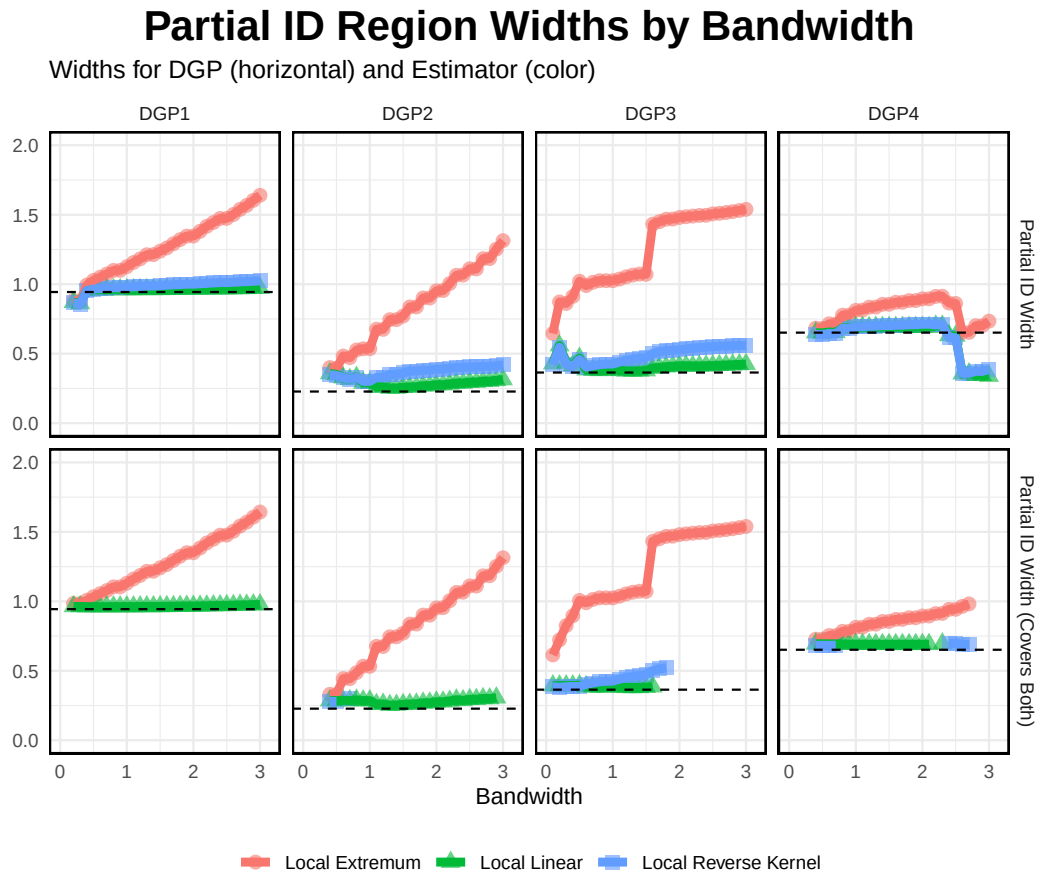

Figure C.20: More Bandwidths: Partial ID Region Widths

Note: The true partially identified region width is plotted as the dashed black horizontal line in each panel.

Missing Data Simulations

The missing data simulations use sample sizes of $n = 1000$ and 1000 simulation runs for each DGP, as specified above. Then I re-run the local effect simulations removing observations with running variable values near the cutoff: those with $|X - 3| \leq 0.5$ are removed from the dataset. Such missingness may occur either due to true missingness or as a robustness check as in donut RDD designs (Noack & Rothe, 2023). Coverage rates (for upper bounds, lower bounds, and both bounds) and partial ID region widths are shown by estimator, bandwidth, and DGP.

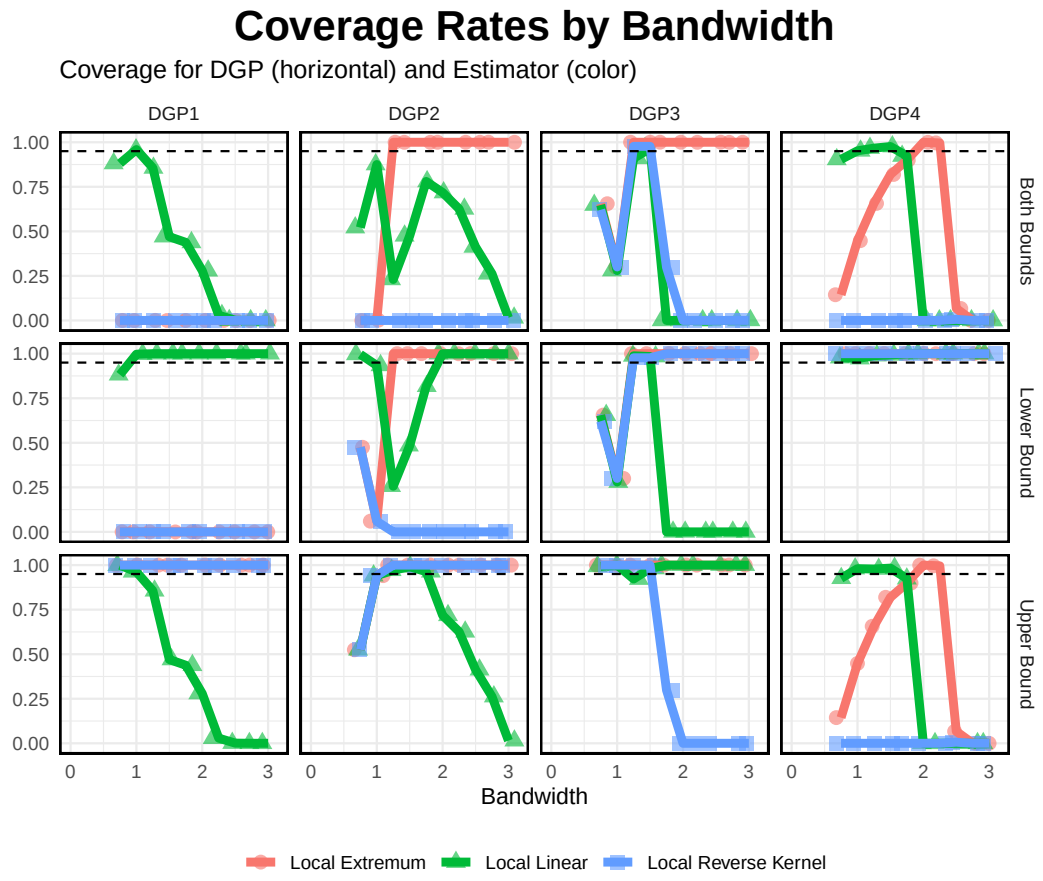

Figure C.21: Missing Data: Coverage Rates

Note: The 95% nominal coverage rate is plotted as the dashed black horizontal line in each panel.

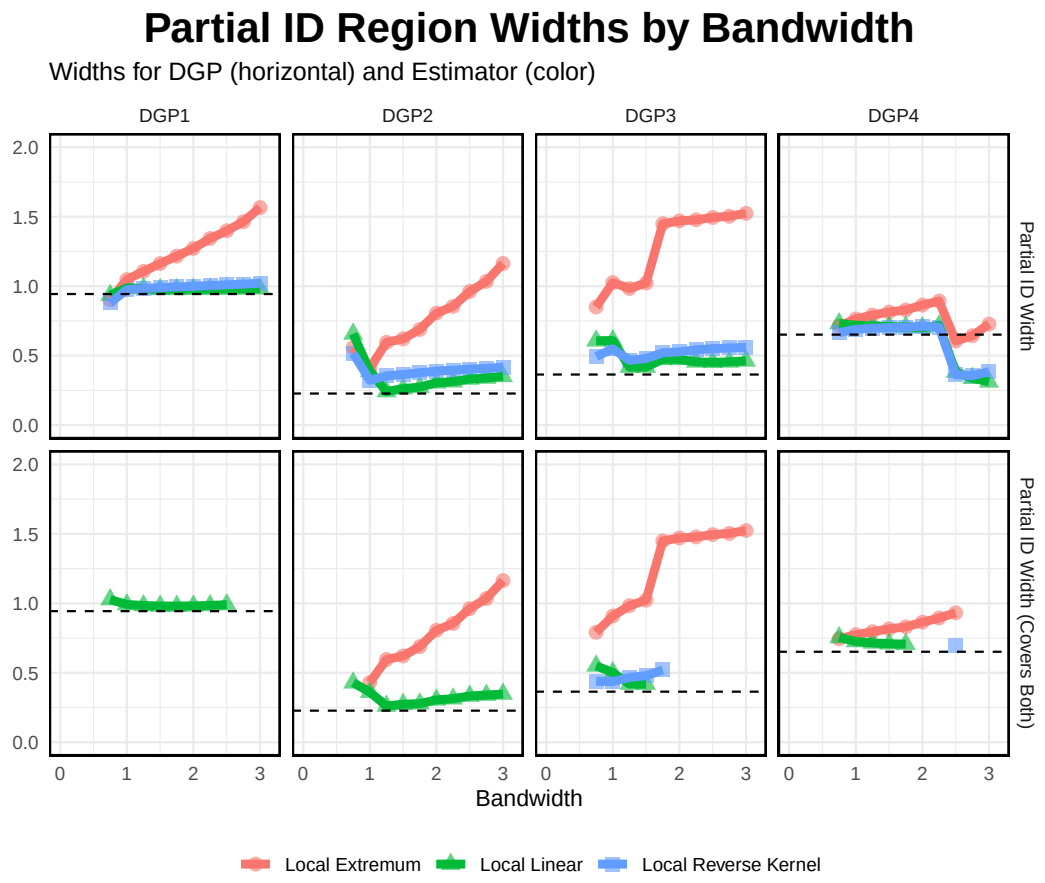

Figure C.22: Missing Data: Partial ID Region Widths

Note: The true partially identified region width is plotted as the dashed black horizontal line in each panel.

Measurement Error Simulations

The measurement error simulations use sample sizes of $n = 1000$ and 1000 simulation runs for each DGP, as specified above. Then I re-run the local effect simulations adding noise (standard normal) and rounding (to nearest two decimals) to each running variable value, while ensuring running variable values lie in $[0, 6]$ and do not shift observations to the other side of the cutoff. This is distortion of the running variable rather than treatment status manipulation. Rounding and classical measurement error may occur in a variety of settings where the running variable is imperfectly reported (by agents or agencies) or estimated by the econometrician. Coverage rates (for upper bounds, lower bounds, and both bounds) and partial ID region widths are shown by estimator, bandwidth, and DGP.

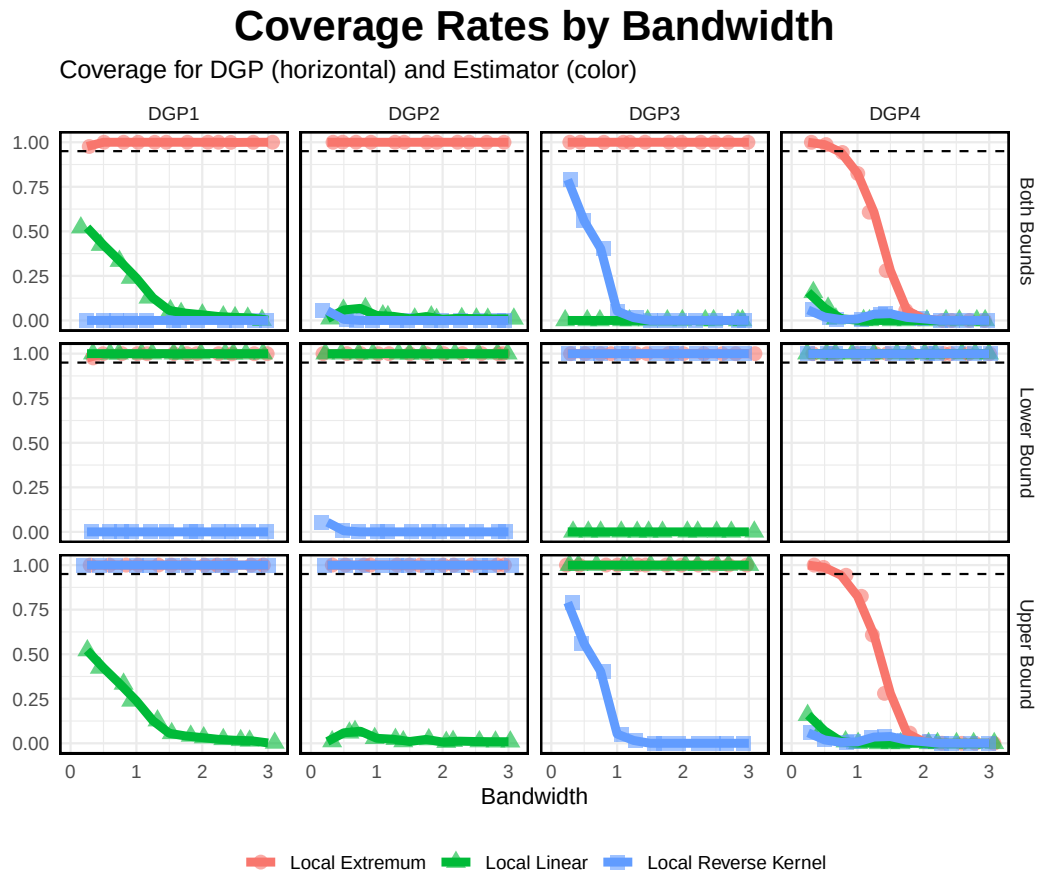

Figure C.23: Measurement Error: Coverage Rates

Note: The 95% nominal coverage rate is plotted as the dashed black horizontal line in each panel.

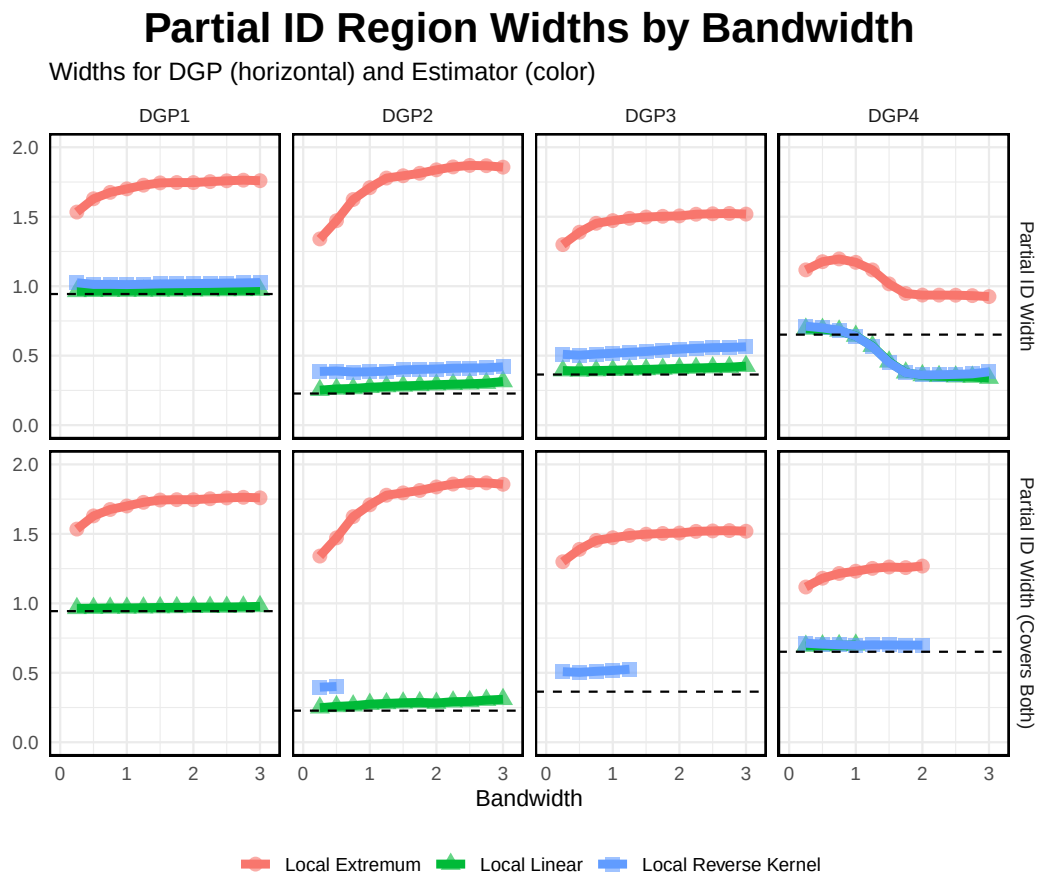

Figure C.24: Measurement Error: Partial ID Region Widths

Note: The true partially identified region width is plotted as the dashed black horizontal line in each panel.

C.8 Additional Application Results

Appendix C.8 contains additional application figures and results.

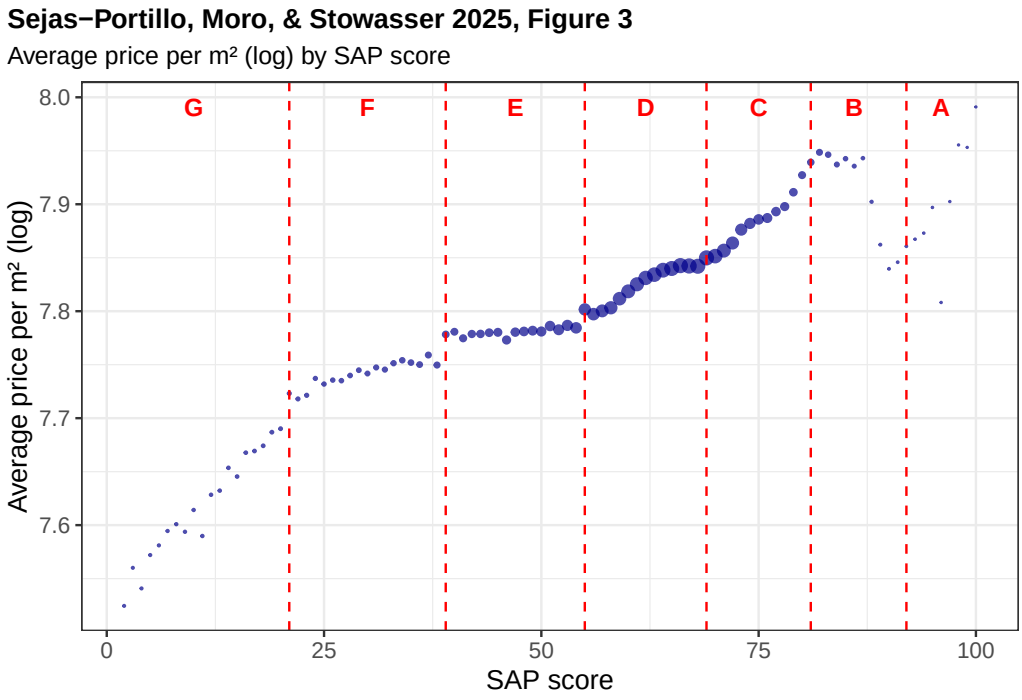

Figure C.25: Sejas-Portillo, Moro, and Stowasser 2025, Figure 3 Reproduction

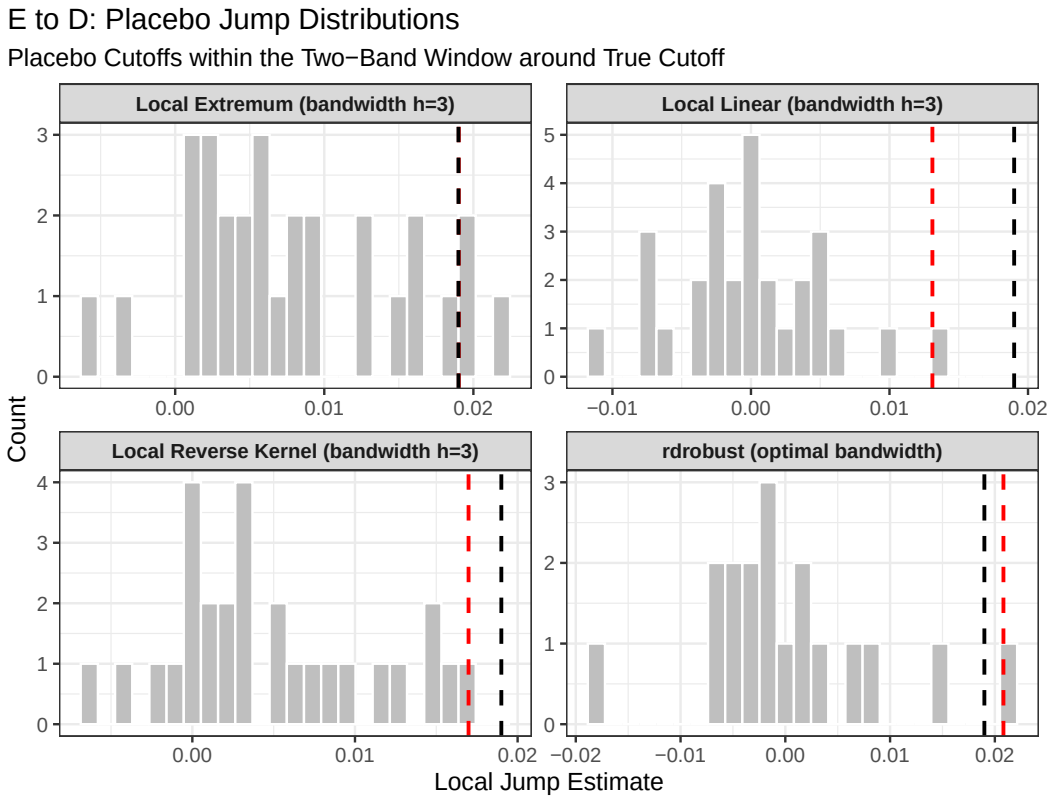

Figure C.26: Placebo Jump Distributions near E to D Cutoff

Note: Estimated placebo effects at each running variable value near the cutoff (but not beyond the next cutoff) are shown for each estimator. The dashed red line shows the estimated cutoff effect at the true cutoff value, and the dashed black line shows the baseline estimate from SpMS 2025.

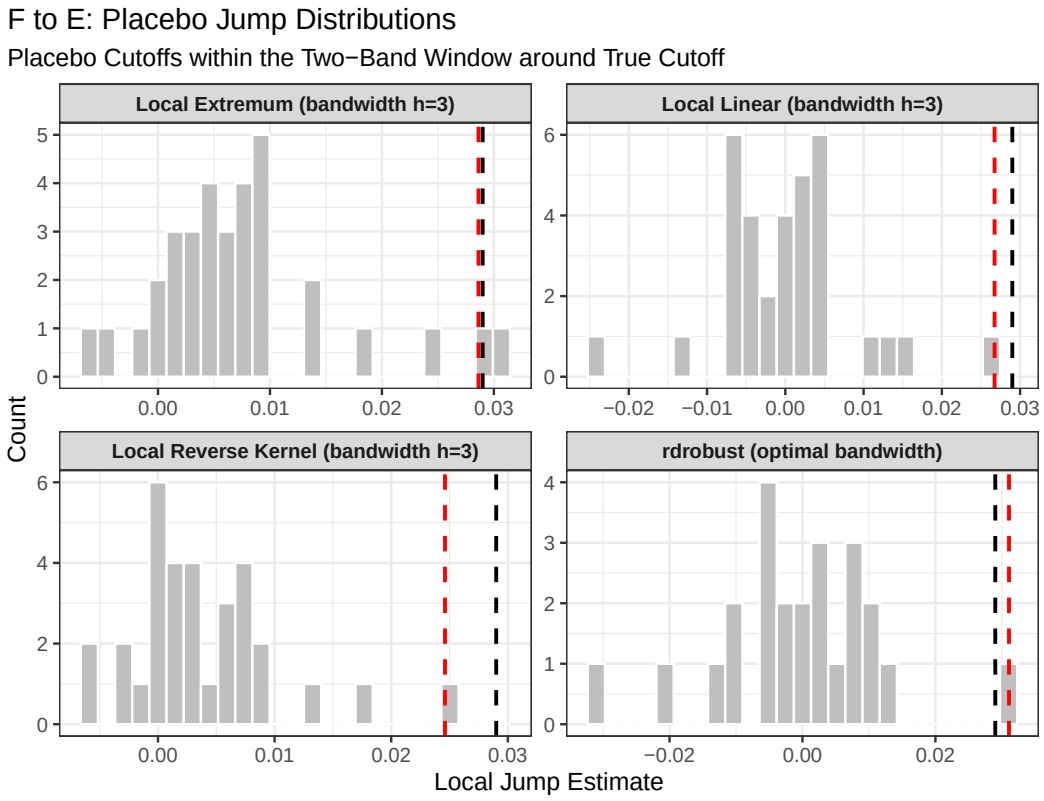

Figure C.27: Placebo Jump Distributions near F to E Cutoff

Note: Estimated placebo effects at each running variable value near the cutoff (but not beyond the next cutoff) are shown for each estimator. The dashed red line shows the estimated cutoff effect at the true cutoff value, and the dashed black line shows the baseline estimate from SpMS 2025.

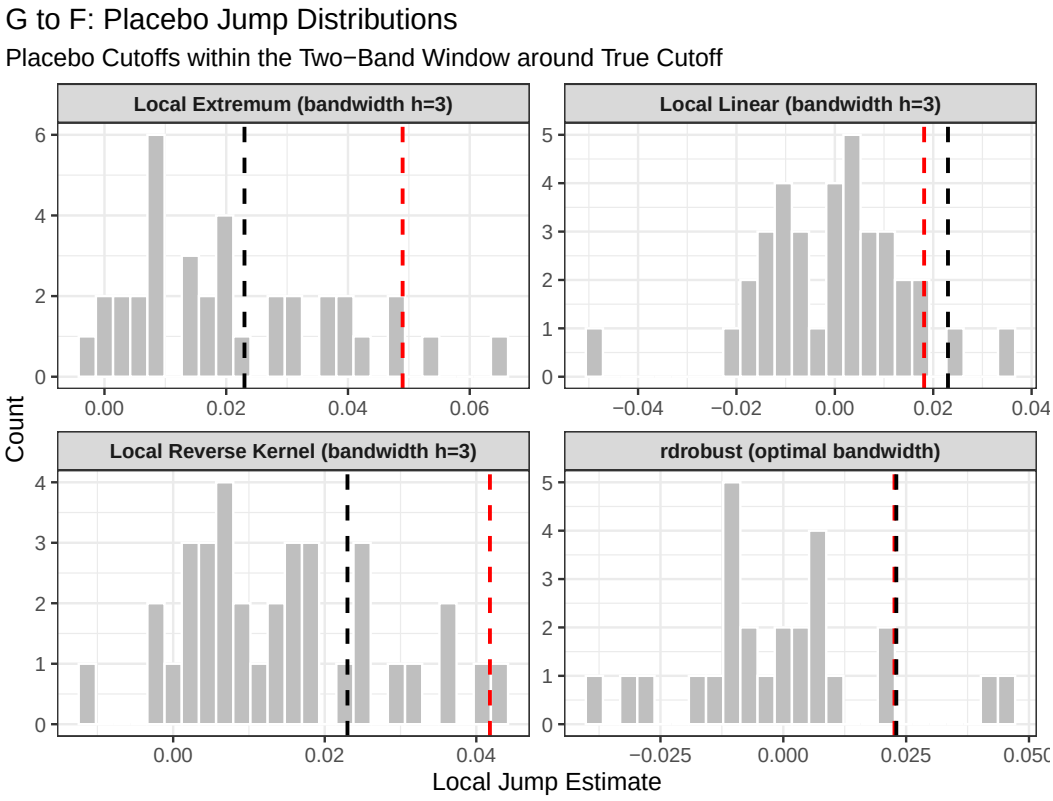

Figure C.28: Placebo Jump Distributions near G to F Cutoff

Note: Estimated placebo effects at each running variable value near the cutoff (but not beyond the next cutoff) are shown for each estimator. The dashed red line shows the estimated cutoff effect at the true cutoff value, and the dashed black line shows the baseline estimate from SpMS 2025.

C.9 Fuzzy RDD

Appendix C.9 considers fuzzy regression discontinuity designs. This section illustrates how causal bounds may be obtained under certain counterfactual assumptions. An example visualization and data generating process are also included.

Fuzzy Monotone RDD: Partial Identification of ATE under Monotone Treatment Selection (MTS)

In the fuzzy RDD, recall that there are treated and control observations on both sides of a cutoff C. Because treatment W is not a deterministic function of X, there are therefore four CEFs of interest:

$$\begin{aligned}
\mu_{00}(x) &= \mathbb{E}[Y_0|W = 0, X = x] \\
\mu_{01}(x) &= \mathbb{E}[Y_0|W = 1, X = x] \\
\mu_{10}(x) &= \mathbb{E}[Y_1|W = 0, X = x] \\
\mu_{11}(x) &= \mathbb{E}[Y_1|W = 1, X = x]
\end{aligned}$$

We consider an assumption to extrapolate to the ATE (or $ATE(\Omega)$). Bounds on the local or cutoff average treatment effect may also be obtained in the obvious way.[2] The assumption does not exploit the jump in treatment probability at the cutoff directly, but rather relies on the relationship between the different $\mu_{gg'}$ functions.

Under the following three assumptions, we can similarly bound the ATE by marginalizing and exploiting monotonicity. The first assumption is needed to deal with the fuzziness of the design:

Assumption 6 (Conditional Monotone Treatment Selection)**.** *For $g = 0, 1$ and $w = 0, 1$ and any $x \in supp(X)$, if $w \geq w'$, then*

$$\mathbb{E}[Y_g|W = w, X = x] \geq \mathbb{E}[Y_g|W = w', X = x]$$

The monotone treatment selection ("MTS") assumption is well-known in the partial identification literature.[3] Here, we only need it to hold conditional on $X = x$, so we call this Conditional Monotone Treatment Selection ("CMTS"). This assumption implies that $\mu_{11}(x) \geq \mu_{10}(x)$ and $\mu_{01}(x) \geq \mu_{00}(x)$ for all x-values.

[2] In particular, if three assumptions below hold, the bound is the same as that from the sharp RDD case.

[3] See, e.g., C. F. Manski (2007).

We further assume monotonicity and boundedness assumptions as in the sharp RDD case:

Assumption 7 (Global Monotonicity of μ_{11}, μ_{00})**.** *The conditional mean functions* $\mu_{11}(x), \mu_{00}(x)$ *are weakly monotone increasing in* x *for all* $x \in supp(X)$.

Assumption 8 (Boundedness)**.** *The conditional mean functions* $\mu_{gg'}$ *for any* $g, g' \in \{0, 1\}$ *are uniformly bounded above by* M_1 *and below by* M_0.

Finally, before proceeding to obtain the sharp upper bound on ATE under these assumptions, define the following sets: let $V_1 = \{v_1, \ldots, v_n\}$ be the set of x-values less than the cutoff C for which $Pr(W = 1|X \in V_1) \neq 0$; similarly, let $U_0 = \{u_0, \ldots, u_k\}$ be the set of x-values greater than or equal to C for which $Pr(W = 0|X \in U_0) \neq 0$. These are the sets of x-values to the left and right of the cutoff C, respectively, for which treated and control units are both observed (i.e., overlap is satisfied).[4]

[4]Technically, when X is continuously distributed there may still be a failure of overlap if, for example, we have a point $v_1 < C$ and $Pr(W = 1|X = v_1) = 1$. If there are treated and control units with $X_i = v_1$, then there is overlap (i.e., the probability of treatment assignment is not zero or one.) We can handle this case in the obvious way, but we simplify and assume overlap at all $v \in V_1$ and $u \in U_0$ for notational simplicity.

Consider then the following upper bound on ATE under these assumptions:

$$
\begin{aligned}
ATE &= \mathbb{E}[Y_1 - Y_0] \\
&= \sum_{w\in\{0,1\}} \sum_{x\in \text{supp}(X)} \mathbb{E}[Y_1 - Y_0 | W = w, X = x] Pr(W = w | X = x) Pr(X = x) \\
&= \sum_{x\in \text{supp}(X)} \mathbb{E}[Y_1 - Y_0 | W = 0, X = x] Pr(W = 0 | X = x) Pr(X = x) \\
&\quad + \sum_{x\in \text{supp}(X)} \mathbb{E}[Y_1 - Y_0 | W = 1, X = x] Pr(W = 1 | X = x) Pr(X = x) \\
&\leq \sum_{x\in \text{supp}(X)} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(W = 0 | X = x) Pr(X = x) \\
&\quad + \sum_{x\in \text{supp}(X)} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(W = 1 | X = x) Pr(X = x) \\
&\leq \sum_{x\in \text{supp}(X)} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) \\
&\leq \sum_{x\leq v_1} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) \\
&\quad + \sum_{v_1 < x \leq v_2} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) + \cdots \\
&\quad + \sum_{v_n < x < C} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) \\
&\quad + \sum_{C \leq x \leq u_0} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) \\
&\quad + \sum_{u_0 < x \leq u_1} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) + \cdots \\
&\quad + \sum_{u_k \leq x} \left(\mathbb{E}[Y_1 | W = 1, X = x] - \mathbb{E}[Y_0 | W = 0, X = x]\right) Pr(X = x) \\
&\leq \sum_{x\leq v_1} \left(\mu_{11}(v_1) - \mu_{00}(x)\right) Pr(X = x) + \cdots + \sum_{v_n < x < C} \left(\mu_{11}(C) - \mu_{00}(x)\right) Pr(X = x) \\
&\quad + \sum_{C \leq x \leq u_0} \left(\mu_{11}(x) - \lim_{x \to C^-} \mu_{00}(x)\right) Pr(X = x) + \cdots + \sum_{u_k \leq x} \left(\mu_{11}(x) - \mu_{00}(u_k)\right) Pr(X = x)
\end{aligned}
$$

In the first inequality, we use Assumption 6; in the second inequality, we grouped like terms and used $Pr(W = 0|X = x) + Pr(W = 1|X = x) = 1$; in the third inequality, we split the support into x-values split by the values in V_1, U_0; finally, in the fourth inequality we use the observed conditional means for x-values in sets V_1, U_0 and Assumption 7 to get bounds on the conditional means μ_{11}, μ_{00} for all $x \in \text{supp}(X)$. What remains in the final line is an estimable upper bound for ATE

under global monotonicity of μ_{11} and μ_{00} using observable data alone. Obtaining the lower bound is similar but requires also using Assumption 8.

Fuzzy RDD Example

To reinforce ideas, we provide an example for the fuzzy RDD upper bound estimation on the following page. We assume that $Pr(X = x)$ is uniformly distributed. In our simulated dataset, X is a discrete distribution with 100 different x-values, but for computing the exact ATE we will assume X is continuously uniformly distributed over $[0, 6]$. The treatment assignment is as follows:

$$Pr(W = 0|X = x) = \begin{cases} 1 & \text{if } 0 \leq x < 2 \\ \frac{1}{2} & \text{if } x = 2 \\ 1 & \text{if } 2 < x < 3 \\ 0 & \text{if } 3 \leq x < 4 \\ \frac{1}{2} & \text{if } x = 4 \\ 0 & \text{if } 4 \leq x \leq 6 \end{cases}$$

This means that there is overlap at $x = 2$ and $x = 4$, but otherwise units with $X_i < 3$ are control units and units with $X_i \geq 3$ are treated units.

The four functions that define the CEFs are as follows:

$$\mu_{00}(x) = \begin{cases} 0.2x^2 - 2 & \text{if } 0 \leq x < 3 \\ 0.7x^2 - 5 & \text{if } 3 \leq x \leq 6 \end{cases} \quad \mu_{01}(x) = \begin{cases} x + 1 & \text{if } 0 \leq x < 3 \\ 5.5x - 10 & \text{if } 3 \leq x \leq 6 \end{cases}$$

$$\mu_{10}(x) = \begin{cases} 0.5x & \text{if } 0 \leq x < 3 \\ -0.1x^2 + 5 & \text{if } 3 \leq x \leq 6 \end{cases} \quad \mu_{11}(x) = \begin{cases} 0.2x^2 + 7 & \text{if } 0 \leq x < 3 \\ 0.6x^2 + 5 & \text{if } 3 \leq x \leq 6 \end{cases}$$

Observe that all four CEFs are discontinuous at the cutoff $C = 3$. This is plotted in the figure below:

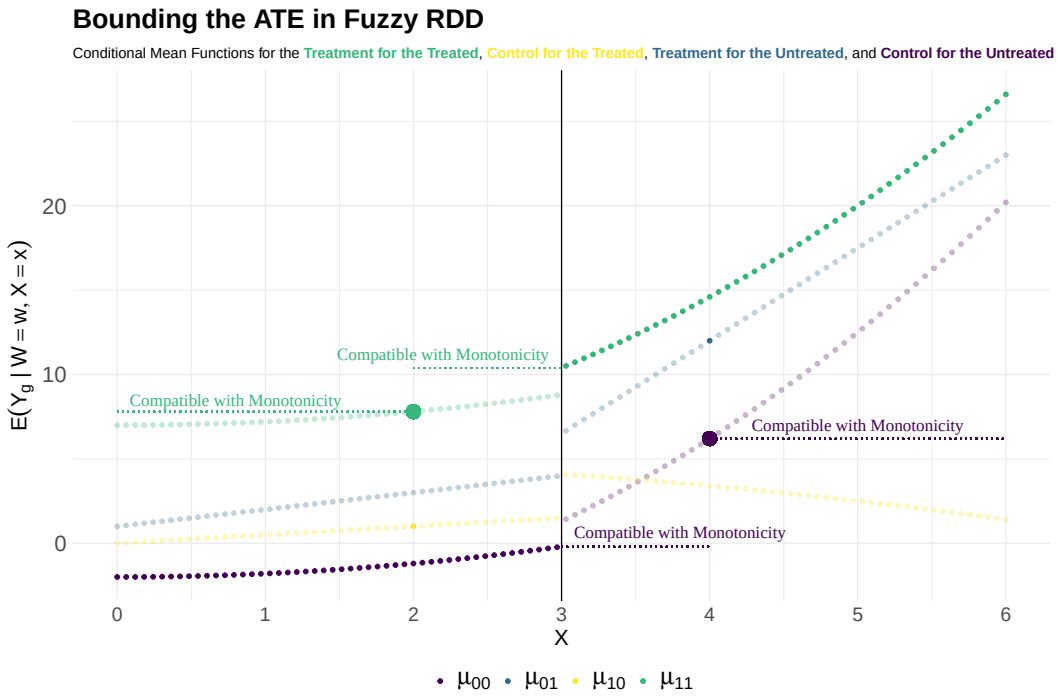

Figure C.29: Bounding ATE in Fuzzy RDD Under Global Monotonicity

Note: This plot shows the four conditional mean functions in the fuzzy discontinuity design, with the lower transparency points unobserved. The two large points are for observed treated units to the left of the cutoff and observed control units to the right of the cutoff, at 2 and 4, respectively. The true treatment effect at the boundary is 3.9, whereas the upper bound is 10.75. The true ATE is 2.49, whereas the upper bound is approximately 11.81.

In the plot, the CMTS assumption is satisfied because the green points lie above the yellow points, and the purple points lie below the blue points. The monotonicity assumption holds for $\mu_{11}(x)$ and $\mu_{00}(x)$ on the entire support of X, but observe that in this example $\mu_{10}(x)$ is not monotone increasing on the entire domain. Recall that we do not require that the unobservable functions $\mu_{10}(x)$ and $\mu_{01}(x)$ be monotone.

In this particular example, the true treatment effect at the cutoff is 3.9, which can be

found analytically as follows for the cutoff $C = 3$:

$$\begin{aligned}
\tau_C &= \mathbb{E}[Y_1 - Y_0 | X = C] \\
&= \mathbb{E}[Y_1 - Y_0 | W = 0, X = C] \underbrace{Pr(W = 0 | X = C)}_{=0} + \mathbb{E}[Y_1 - Y_0 | W = 1, X = C] \underbrace{Pr(W = 1 | X = C)}_{=1} \\
&= \mathbb{E}[Y_1 | W = 1, X = C] - \mathbb{E}[Y_0 | W = 1, X = C] \\
&= \mu_{11}(C) - \mu_{01}(C) \\
&= (0.6 * (3)^2 + 5) - (5.5 * 3 - 10) \\
&= \boxed{3.9}
\end{aligned}$$

The upper bound at the cutoff is found by using that $\mu_{01}(C) \geq \mu_{00}(C)$ and then replacing C with the last observable value of x for control units to the left of the cutoff. Similarly, since we do not observe $\mu_{11}(C)$ but we do observe $\mu_{11}(3.03)$ we will use that for bounding. In this case, this results in the following upper bound:

$$\begin{aligned}
\overline{\tau}_C &= \mu_{11}(C) - \mu_{01}(C) \\
&\leq \mu_{11}(3.03) - \mu_{00}(2.97) = 10.51 - (-0.24) = \boxed{10.75}
\end{aligned}$$

The ATE is given by the following computation (for X continuously uniformly distributed on $[0, 6]$ so that $Pr(X = x) = \frac{1}{6}$ for all x):

$$\begin{aligned}
ATE &= \frac{1}{6} \left(\int_{x=0}^{3} (\mu_{10}(x) - \mu_{00}(x)) dx + \int_{x=3}^{6} \mu_{11}(x) - \mu_{01}(x) dx \right) \\
&= \frac{1}{6} (6.45 + 8.55) \\
&= \boxed{2.5}
\end{aligned}$$

We are using that $Pr(W = 0 | X = x)$ is zero or one except at two points, which does not matter for the integration. In the caption below the plot, we note that ATE is 2.49 for the dataset generated from the $\mu_{gg'}(x)$ functions with 100 x-values evenly spaced. Finally, we also compute the upper bound on ATE from the generated dataset and get an upper bound of

$$\overline{ATE} = \boxed{11.81}$$

Appendix D

SUPPLEMENTARY MATERIAL FOR CHAPTER 4

D.1 Additional Descriptive Plots

Appendix D.1 includes additional descriptive plots. We include the following plots: two additional academic citation series for the legal realist works (cumulative vs yearly citations); court mentions for realist works yearly; and court mentions for each author aggregated across works and for each decade.

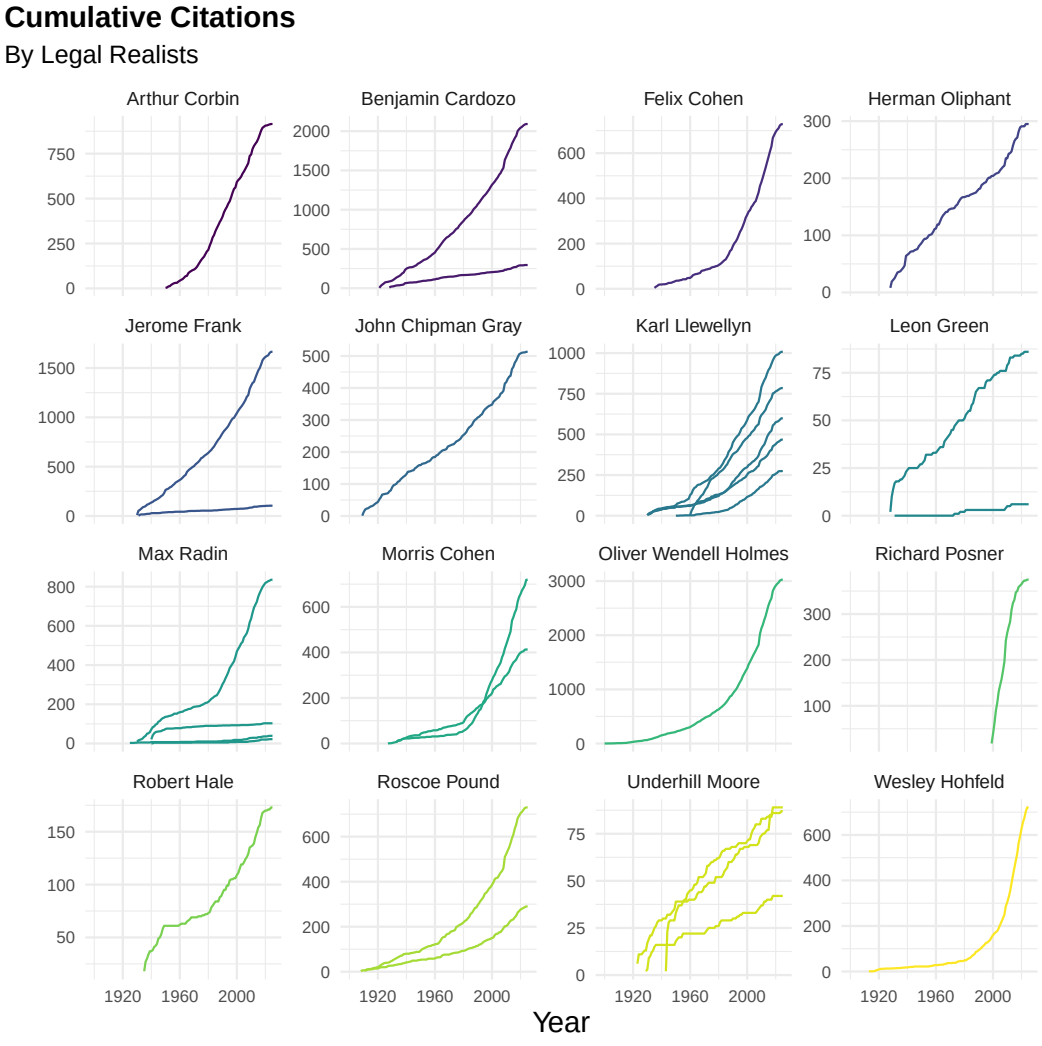

Figure D.1: Cumulative Citations

Note: The y-axis differs for each legal realist. The different lines for each legal realist are for their different texts.

Figure D.2: Yearly Citations

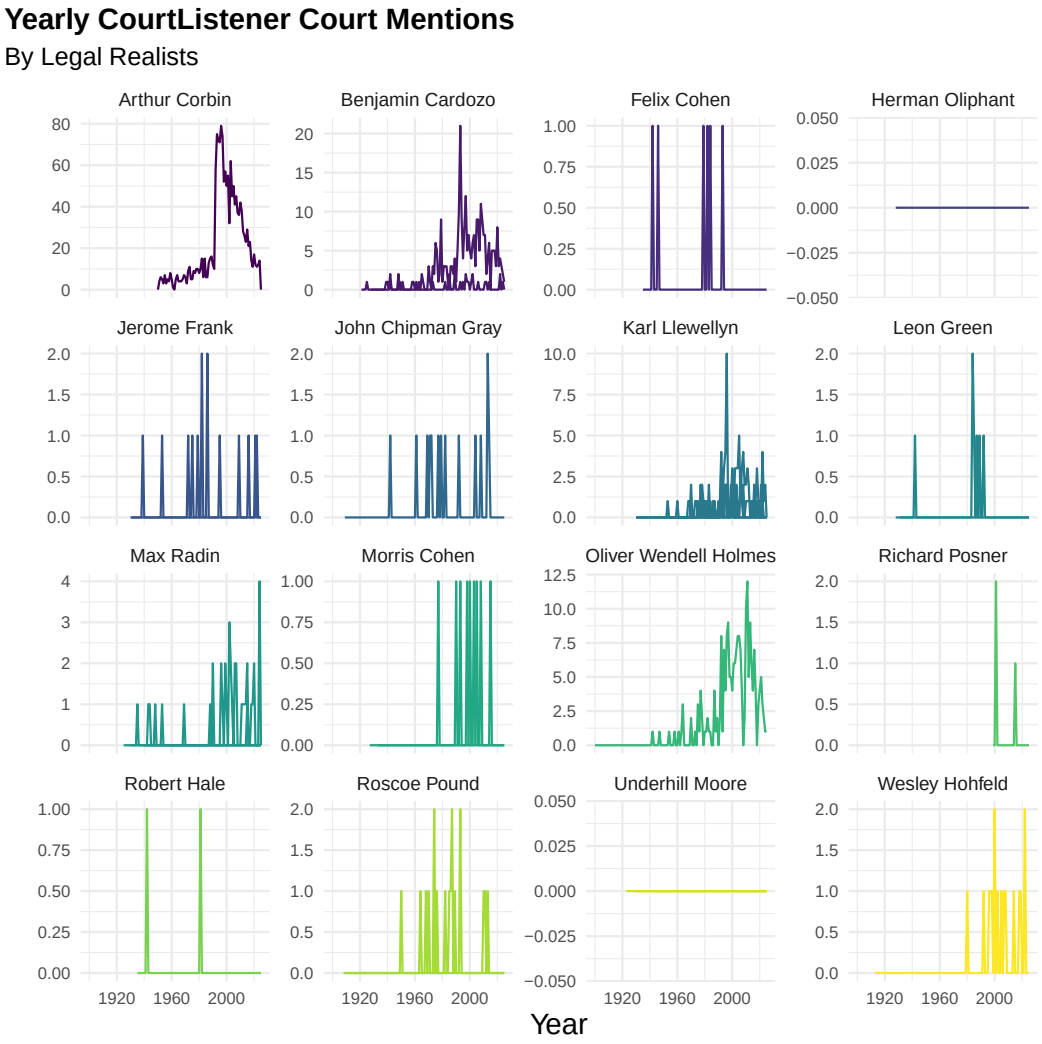

Figure D.3: Yearly CourtListener Court Mentions

Note: The y-axis differs for each legal realist.

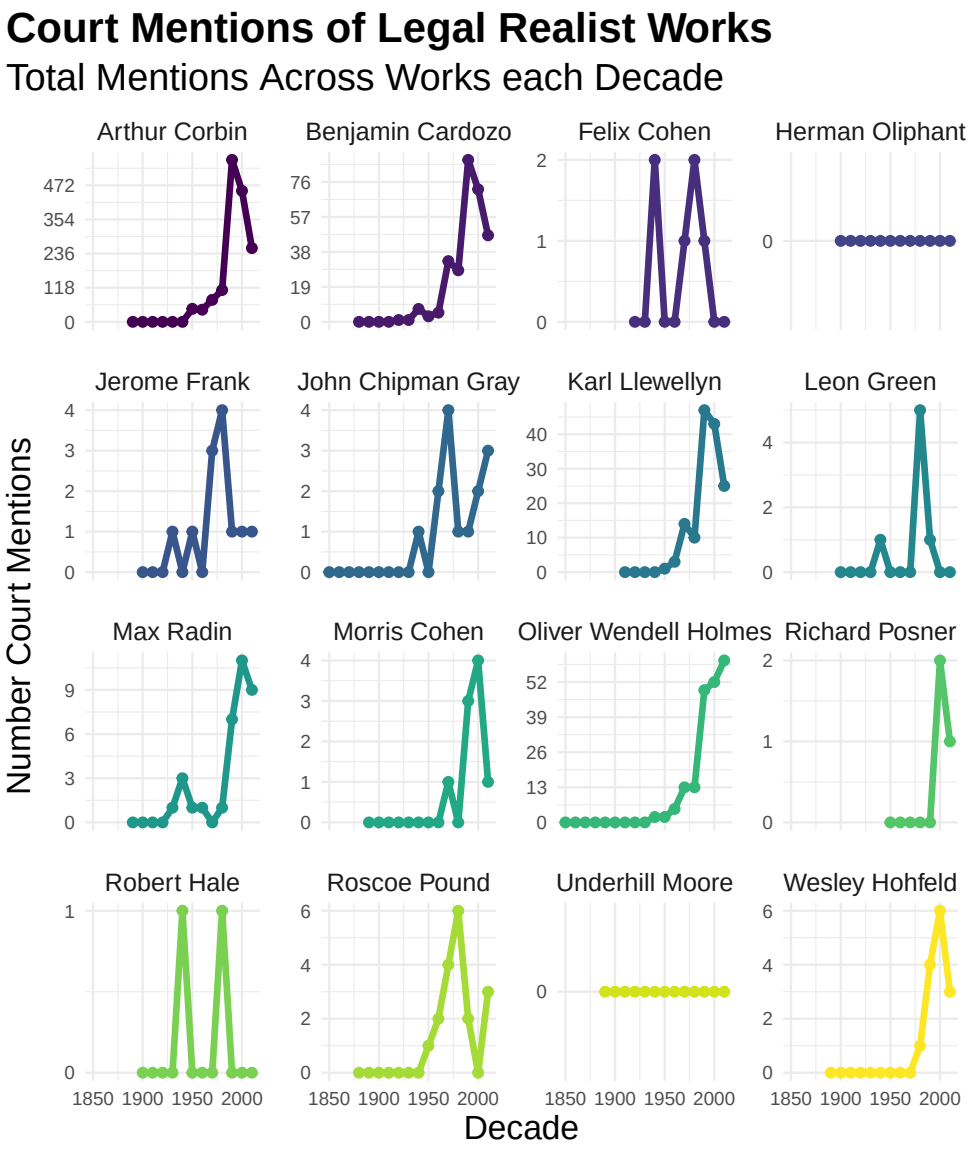

Figure D.4: Legal Realist Court Mentions

Word Clouds by Legal Movement

Legal Realists: Top 100 Terms

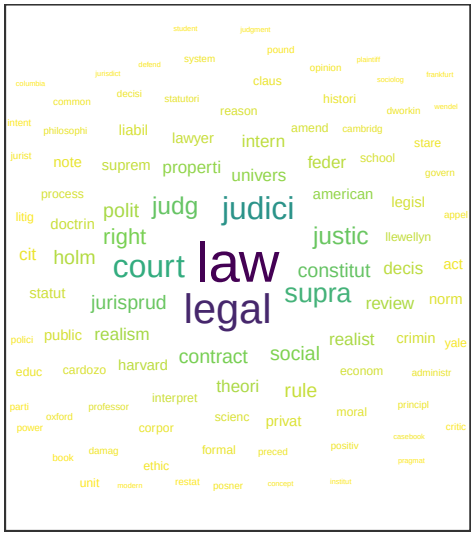

Originalists: Top 100 Terms

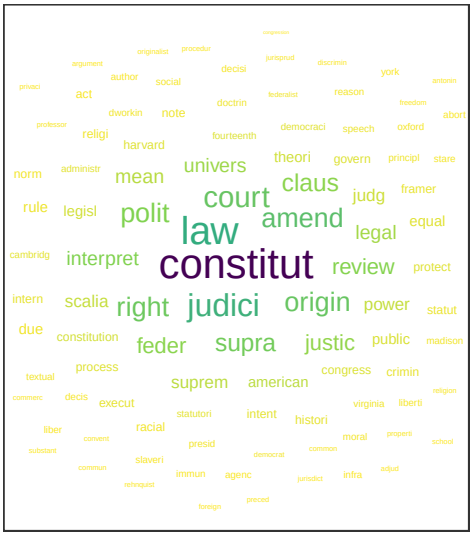

Figure D.5: Word Clouds by Legal Movement

D.2 Counterfactual Complexity & The Curse of Dimensionality

Appendix D.2 includes discussion and a plot of how potential outcomes grow as a function of the complexity of treatment paths. Researchers typically rely on simplifying assumptions rather than allowing for full path heterogeneity. Here, we discuss two types of assumptions that are (implicitly or explicitly) used in practice: memory and anticipation assumptions.

The first set of assumptions we term memory assumptions, which restrict how the potential outcome in period t can depend on treatments received prior to time t. Formally, we define *memory of degree m* to be the assumption that, for all i and in each time t, we have:

$$Y_{it}(\vec{k}) = Y_{it}(k_{t-m}, \ldots, k_t, \ldots k_T).$$

This restricts the potential outcomes in period t to depend only on the treatment path up to m periods prior to t. For example, memory of degree zero (*no memory*) means that potential outcomes in time t depend only on treatments received during or after time t.

The second set of assumptions are anticipation assumptions. Commonly used the DID literature, this assumption restricts how potential outcomes in time t can depend on treatments received after time t. Formally, we define *anticipation of degree a* to be the assumption that, for all i and in each time t, we have:

$$Y_{it}(\vec{k}) = Y_{it}(k_1, \ldots, k_t, \ldots k_{t+a})$$

This restricts the potential outcomes in period t to depend only on the treatment path up to a periods after t. For example, anticipation of degree zero (*no anticipation*) means that potential outcomes in time t depend only on treatments received during or before time t.

Appendix Figure D.6 illustrates the dimensionality problem. We plot the number of potential outcomes under different memory assumptions and no anticipation. We consider settings with ten time periods and between 1-10 treatment dosages other than the baseline "no treatment" dosage. As seen in Appendix Figure D.6, stronger memory assumptions greatly reduce the dimensionality of the problem. However, as the cardinality of the treatment set $\mathcal{K}$ grows, the total number of potential outcomes remains large: under no memory and no anticipation, there are 100 total potential outcomes for each unit.

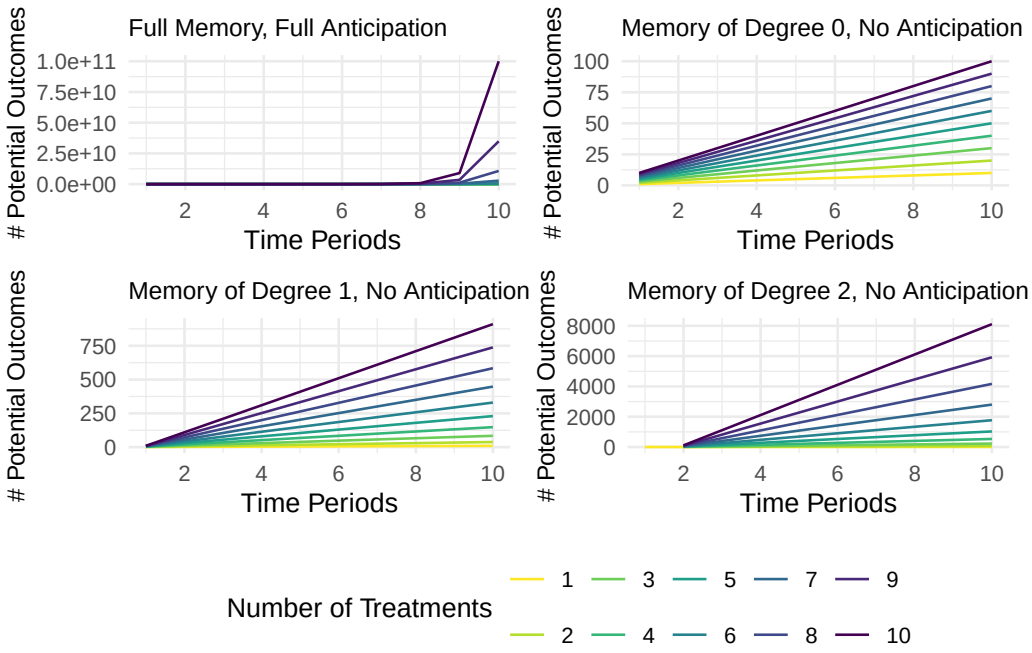

Figure D.6: Potential Outcomes: Combinatorial Explosion

Note: This plot shows the total number of potential outcomes as a function of time periods (x-axis) and the number of treatments (line color). The number of treatments is the number of treatments in addition to a baseline "no treatment" dosage. The yellow line (1 treatment) is, e.g., a binary treatment in each period.

Even with these restrictions on the nature of the treatment path, counterfactual imputation is not as straightforward as the simpler setting. For one, there are now more counterfactual paths. Researchers might ask what the unit's outcome would be if it were not treated at all, if it were treated earlier, or if it were treated later. The standard approach comparing treated units to not-yet-treated units to obtain treatment effects mimics the simpler settings described above.

These complications in dynamic settings are exacerbated as the complexity of the treatment path increases. Imputing counterfactuals with complex treatment paths is challenging for two reasons. The first is sparsity in the path space itself. Indeed, as shown in Figure 4.2, there is limited overlap in the treatment histories for units that are mentioned at least once in judicial opinions. The second is that even if many paths are observed, there will be too few observed units on each path to feasibly point identify a causal estimate. Both these make the usual machinery employed in causal inference difficult to work in these settings. Under such complex paths, there is no clear demarcation between treatment and control units, particularly if all units

get some treatment at some point.

Because it works to impute unit-level counterfactuals, the Bounded Deviations method we use is one possible solution. Our identification strategy partially identifies counterfactual paths for units receiving different treatment paths. If the length of a treatment path grows, the number of such counterfactual paths can be unmanageable quickly. To make this manageable, we could only limit to paths that we observe, or make parametric assumptions. For example, we could define a distance metric on treatment paths and assume that counterfactual outcomes are a function of the distance between paths. Counterfactual paths that are closer to the observed treatment path, e.g., may only have small, predictable deviations from the observed outcomes.

Possible Problems with Dimension-Reduction Assumptions

Memory-anticipation assumptions clearly reduce the dimensionality of the identification problem. Still, computational and interpretational problems may remain, especially in settings with numerous treatment levels (i.e., doses). But the main problem is that such assumptions can introduce bias and are difficult to motivate or justify in settings without well-founded theory. That is, if there are meaningful differences in potential outcomes along different treatment paths, comparing units along different paths can bias causal estimates.

Consider, for example, the simple two-period model with binary treatments available both time periods. If there is genuine anticipation, then the counterfactual averages along paths $(0,0)$ and $(0,1)$, say, may differ. That is, units that are untreated at time $t = 1$ but treated in time $t = 2$ may anticipate receiving treatment, act on it, and thereby experience different outcomes than units that are never treated (i.e., whose treatment paths are $(0,0)$). If we then assume there is no anticipation, or that $Y(0,0) = Y(0,1) \equiv Y(0)$, we improperly "mix" these units together. This can result in estimates of $\mathbb{E}[Y(0)]$ that differ substantially from the true population counterfactual averages $\mathbb{E}[Y(0,0)]$ and $\mathbb{E}[Y(0,1)]$. The direction and magnitude of the bias from assuming no anticipation can be difficult to predict or correct *ex post*.[1]

In our setting, modeling anticipation may be especially difficult. Although researchers studying, e.g., the effects of particular pre-announced, well-discussed congressional or executive legal actions may be able to model anticipation, it is not

[1]This problem is analogous to the problem of "forbidden comparisons" in staggered adoption settings. Goodman-Bacon (2021), Callaway and Sant'Anna (2021b).

clear what legal academics anticipate in terms of judicial citations. It may be very difficult to anticipate which academic articles judges will cite, yet some law review pieces are prepared with particular court cases (esp. Supreme Court cases) in mind. Indeed, some legal academics may expect their writings to be successful because of pending judicial decisions and even file *amicus* briefs encouraging judges to engage with their scholarship. Particularly for a legal movement like the originalism movement in constitutional law, judges seem to be openly inviting academic discussion of particular legal issues.[2]

Simplifying the treatment path space is, therefore, not without possible downsides. However, in spaces with high-dimensional (time-varying) treatments, simplifying from full-path treatment heterogeneity to a few treatment cohorts is necessary to obtain overlap. As stated above, the flexible partial identification assumptions we propose above can work to impute unit-level counterfactuals (as in, e.g., synthetic control). This is because the unobserved potential outcomes depend on the observable potential outcomes but differences are "bounded."

Individual Treatment Effects & Identifying Assumptions

Finally, we briefly sketch how one could consider individual treatment effect estimands. As in the synthetic control literature,[3] we consider individual treatment effect estimands of the following form:

$$\Delta_i(t, \vec{k}, \vec{k}') = Y_{it}(\vec{k}) - Y_{it}(\vec{k}').$$

In what follows, we can obtain similar bounds on these individual treatment effects by the modified assumption below:

Assumption 9 (Individual Bounded Deviation)**.** *Let $T_{div} = argmin_t(k_t \neq k'_t)$ for two units i and j on distinct treatment paths $\vec{k}$ and $\vec{k}'$, respectively. The following holds for each $t \geq T_{div}$ and some $C \in \mathbb{R}_+$:*

$$\left|Y_{it}(\vec{k}') - Y_{jt}(\vec{k}')\right| \leq C \cdot \max_{t<T_{div}} \left|Y_{it} - Y_{jt}\right|.$$

[2] See, e.g., *Franklin v. New York*, 604 U.S. (2025) (Alito, J., concurring) (discussing the possibility of overturning Confrontation Clause precedent based on "[h]istorical research").

[3] See, e.g., Abadie, Diamond, and Hainmueller (2010) and Abadie, Diamond, and Hainmueller (2015).

D.3 Full Treatment Histories

Appendix D.3 contains the full treatment histories for each of the realist and originalist works studied here. For each work, we plot the number of judicial mentions the work receives each year following publication. The levels are colored according to the legend at the top of each figure.

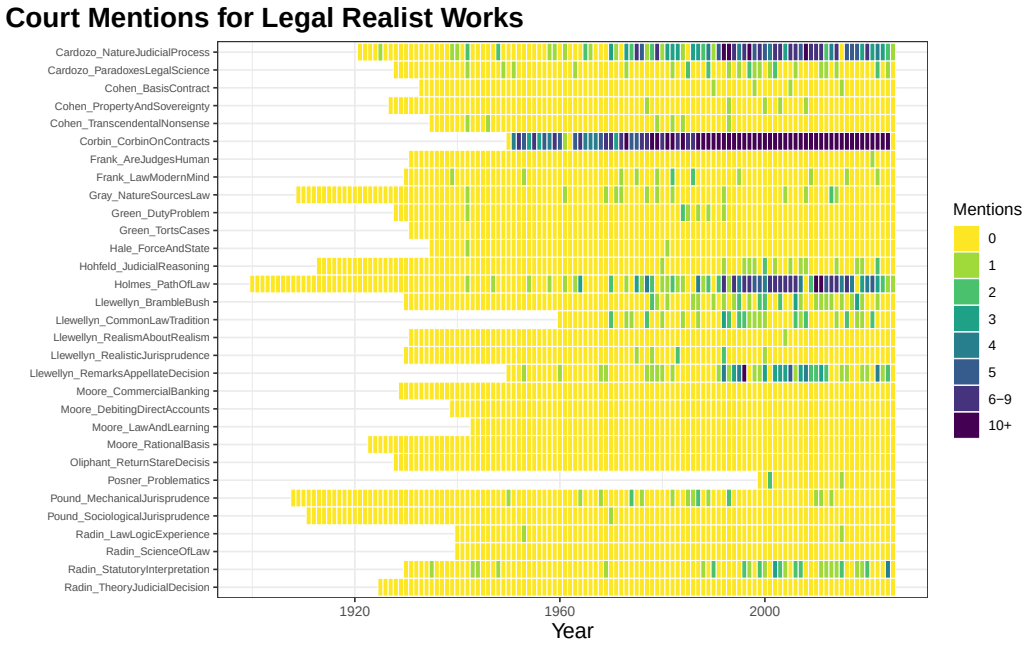

Figure D.7: Legal Realists: Court Mentions

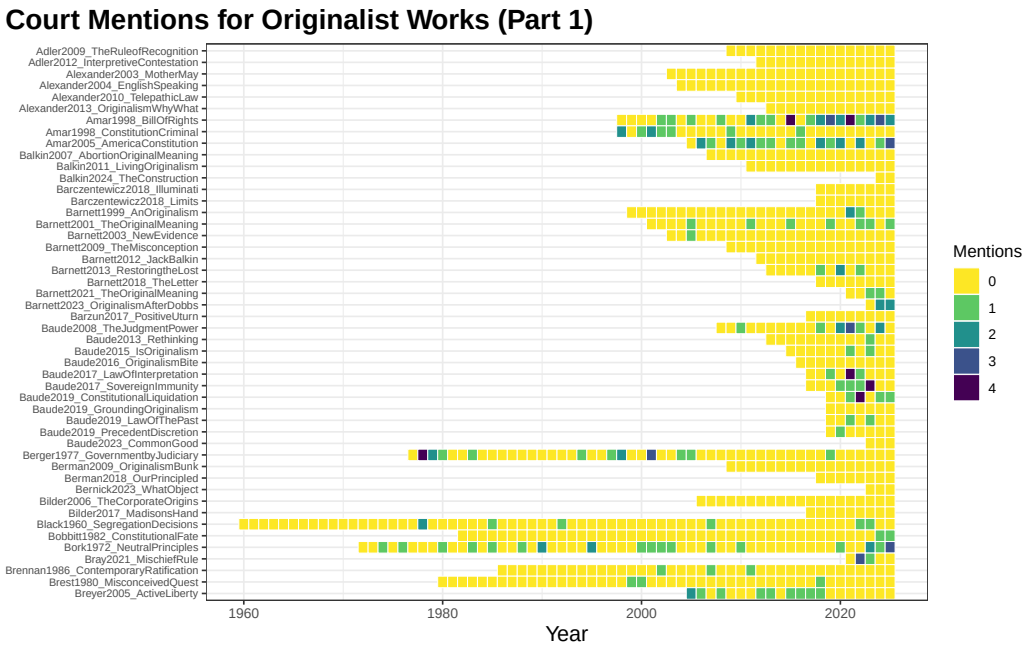

Figure D.8: Originalists: Court Mentions (Part 1)

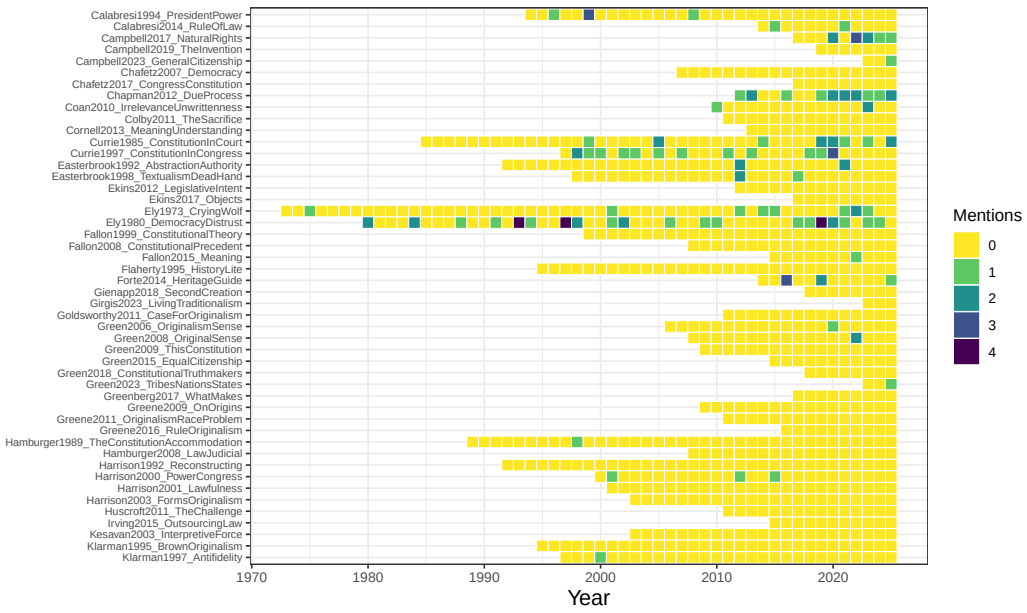

Figure D.9: Originalists: Court Mentions (Part 2)

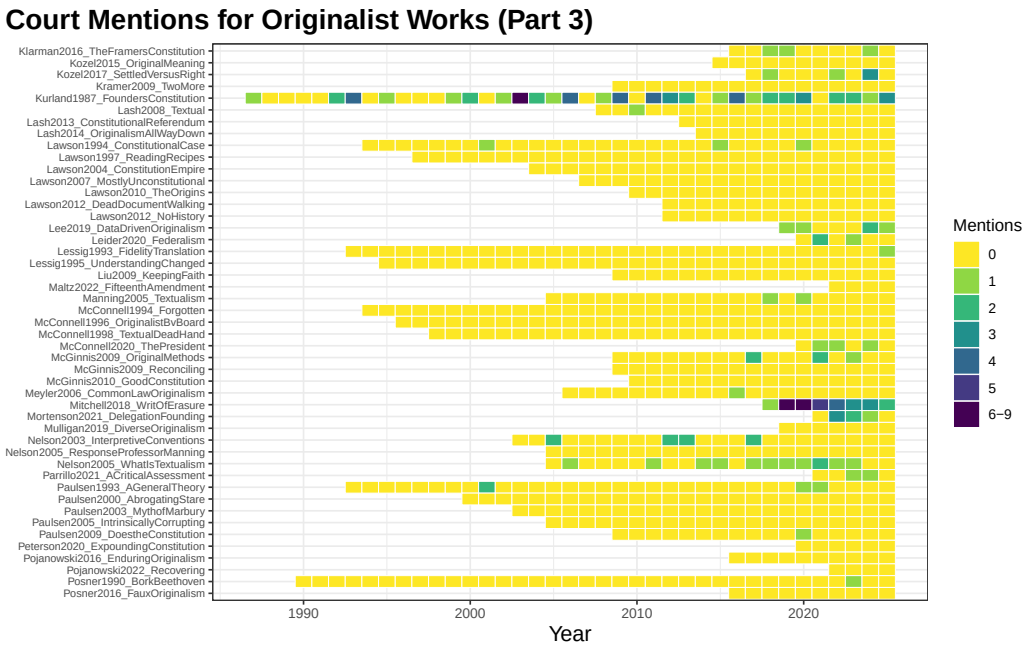

Figure D.10: Originalists: Court Mentions (Part 3)

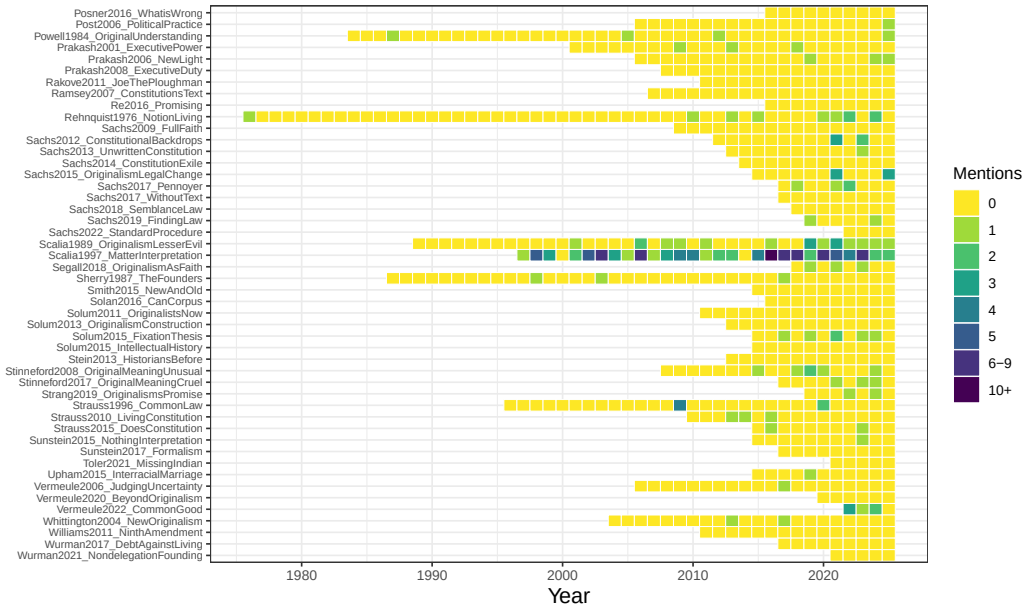

Figure D.11: Originalists: Court Mentions (Part 4)

D.4 Counterfactual Framework and Identification

Appendix D.4 further describes our counterfactual framework and identification strategy.

Counterfactual Framework

We first introduce notation to describe the potential outcomes as a function of the treatment histories or treatment paths for each individual i in each time period t. Consider a setting with time periods $t \in \mathcal{T} = \{1, \ldots, T\}$ and treatment histories $\vec{K}_i = (k_{i1}, \ldots, k_{iT})$, where k_{it} is the observed treatment level for individual i at time period t. Note that, in principle, each k_{it} can be multidimensional if there are multiple treatments, but we focus on scalar k_{it} within a discrete set $\mathcal{K} = \{0, 1, \ldots, K\}$.

We write the potential outcomes for each individual i in time period t under treatment path $\vec{k} = (k_1, \ldots, k_T)$ as

$$Y_{it}(\vec{k}) = Y_{it}(k_1, \ldots, k_T)$$

Implicit in this notation is that there are no treatment "spillover" effects, i.e., the potential outcome for unit i at time t depends only on i's treatment path. Finally, we define the potential outcome matrix $\boldsymbol{Y_i}$ for individual i. For a fixed ordering of the potential treatment paths $\mathcal{K} = \{\vec{k}_1, \vec{k}_2, \ldots\}$, the (i, j)-th entry of $\boldsymbol{Y_i}$ is given by $Y_{i,t=j}(\vec{k}_i)$. For example, consider a two-time period model with a binary treatment possible in each period and an individual i with observed treatment path $\vec{K}_i = (k_{i1} = 0, k_{i2} = 1)$. Then we write i's potential outcome matrix $\boldsymbol{Y_i}$ as follows:

$$\boldsymbol{Y_i} = \begin{bmatrix} Y_{i,t=1}(0,0) & Y_{i,t=2}(0,0) \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ Y_{i,t=1}(1,0) & Y_{i,t=2}(1,0) \\ Y_{i,t=1}(1,1) & Y_{i,t=2}(1,1) \end{bmatrix} = \begin{bmatrix} ? & ? \\ Y_{i,t=1}(0,1) & Y_{i,t=2}(0,1) \\ ? & ? \\ ? & ? \end{bmatrix}$$

The unobservable potential outcomes for i are denoted by the question marks. They define the **identification problem**: because unit i has treatment history $\vec{K}_i = (0, 1)$, the outcomes for i in period t at different treatment histories are unobservable.

Bounded Deviations: Identification

To begin, we specify the treatment effect estimand of interest. We study the *average treatment effect on the treated* ("ATT") for units on two treatment paths, which we define as follows: consider for treatment paths $\vec{k}$ and $\vec{k}'$ the average treatment effect in time t on units treated with history $\vec{k}$ as:

$$ATT(t, \vec{k}, \vec{k}') = \mathbb{E}[Y_{it}(\vec{k}) - Y_{it}(\vec{k}') | \vec{K}_i = \vec{k}]$$

This is the average treatment effect from switching to path $\vec{k}'$ for units on path $\vec{k}$ at time t. As discussed above, the identification problem is that $Y_{it}(\vec{k}')$ is unobservable for units with observed treatment history $\vec{K}_i = \vec{k}$.

Consider two distinct treatment paths $\vec{k}$ and $\vec{k}'$, and let the first time period the treatment histories do not agree be denoted $T_{div} = \text{argmin}_t(k_t \neq k'_t)$. We use information from the pre-divergence periods (i.e., time periods before T_{div}) and the following assumption to impute counterfactual outcomes.

Assumption 10. *[Average Bounded Deviation] Let $T_{div} =_t (k_t \neq k'_t)$ for distinct treatment paths $\vec{k}$ and $\vec{k}'$. Then the following holds for each i, $t \geq T_{div}$, any $(\vec{k}, \vec{k}')$ pair, and some $C \in \mathbb{R}_+$:*

$$\left|\mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}(\vec{k}')|\vec{K}_i = \vec{k}']\right| \leq C \cdot \max_{t<T_{div}} \left|\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}']\right|$$

where $C = C(t, \vec{k}, \vec{k}')$ can depend on the time period and both treatment paths.

This assumption says the counterfactual average outcome from treating with $\vec{k}'$ instead of $\vec{k}$ for units with observed treatment path $\vec{K}_i = \vec{k}$ is "near" the observed outcome for units on treatment path $\vec{k}'$. How near? Specifically, the counterfactual average for the $\vec{k}$-treated group is within C times the the maximum absolute difference in means for units on the two paths before the paths diverge.[4]

This assumption has "identifying power" because it relates the unobservable counterfactual outcomes—i.e., the conditional expectation $\mathbb{E}[Y_{it}(\vec{k}')|K_i = \vec{k}]$—to observable averages. Under this assumption, the average treatment effect on the treated (ATT) is partially identified. That is, we may obtain upper and lower bounds on the ATT. The bounds are given by the following observable quantities for fixed C:

$$\overline{ATT(t, \vec{k}, \vec{k}')} = \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \left(\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}'] - C \cdot \text{Pre-Div}(\vec{k}, \vec{k}')\right) \quad \text{(D.1)}$$

$$\underline{ATT(t, \vec{k}, \vec{k}')} = \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \left(\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}'] + C \cdot \text{Pre-Div}(\vec{k}, \vec{k}')\right) \quad \text{(D.2)}$$

where the maximum pre-divergence deviation is given by:

$$\text{Pre-Div}(\vec{k}, \vec{k}') = \max_{t<T_{div}} \left|\mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}] - \mathbb{E}[Y_{it}|\vec{K}_i = \vec{k}']\right|$$

[4]Note that $C = C(t, \vec{k}, \vec{k}')$ can depend on the time period t and treatment paths $(\vec{k}, \vec{k}')$. In what follows, we consider fixed C-values that are constant across all post-divergence time periods. But, in principle, further assumptions on C as a function of the treatment paths $\vec{k}, \vec{k}'$ or the time period can be incorporated. For example, in some cases there may be good reasons (e.g., theory, past studies, etc.) to believe the bound should grow in time. That is, C should be a weakly increasing function of t.

In this notation, the upper bound is given by $\overline{ATT(t, \vec{k}, \vec{k}')}$ and the lower bound is given by $\underline{ATT(t, \vec{k}, \vec{k}')}$. The method is illustrated below, with descriptive caption.

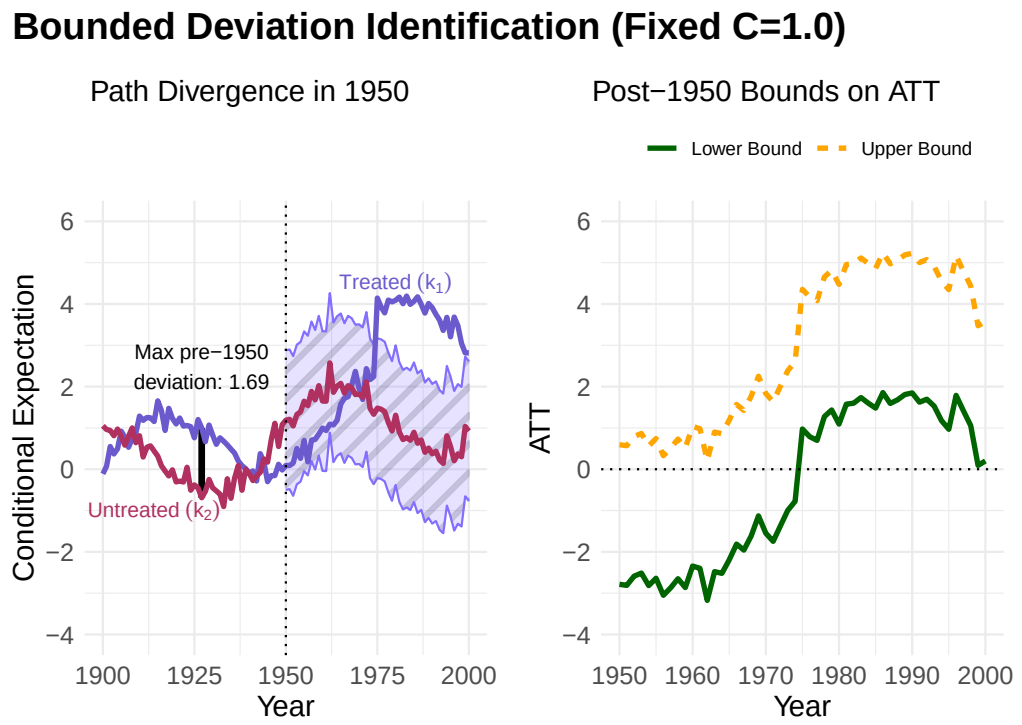

Figure D.12: Bounded Deviations (Fixed $C = 1$)

Note: In the left panel, the two lines are $\mathbb{E}[Y(\vec{k})|\vec{K}_i = \vec{k}]$ for two treatment paths: $\vec{k}_1$ (blue) and $\vec{k}_2$ (red). The black dotted vertical line is the path divergence time period (T_{div}). The blue-striped region is the counterfactual region for blue units (i.e., on treatment path $\vec{k}_1$) if they had instead experienced the red unit treatment path ($\vec{k}_2$). The maximum pre-divergence deviation is 1.69 and is shown by the black vertical line around 1925. In the right panel, y-axis is $ATT(t, \vec{k}_1, \vec{k}_2) = \mathbb{E}[Y_{it}(\vec{k}_1) - Y_{it}(\vec{k}_2)|\vec{K}_i = \vec{k}_1]$. The black dotted horizontal line is plotted for $ATT = 0$. The upper and lower bounds on the ATT for each post-1950 time are plotted as dashed-yellow and solid-green lines, resp.

Comparison with DID

For those familiar with the vast economics literature on difference-in-differences (DID), we also compare our Bounded Deviations method to the simpler DID standard. We consider here the simplest possible setting where there are only two groups (the treated and control) that may receive a single dose of a binary treatment in the year 1950.[5] In the left panel of Appendix Figure D.13, we show the standard DID

[5]For example, this could be: "Did Congress ban the sale of good i in 1950?" If the answer is yes, then good i is part of the treated group. If the answer is no, then good i is part of the control group.

strategy to identify causal effects in this simple setting. The solid-blue line is the observed average outcome in each year for treated units. The solid-red line is the observed average outcome for the control group for each year.

How to estimate the causal effect of treatment on the treated group? To estimate the treatment effect, the DID-strategy assumes that the counterfactual potential outcome would've trended in parallel to the control group. This is called the *parallel trends* assumption and it is the basis for a vast empirical economics literature.[6] In Appendix Figure D.13, the dashed-blue line shows this counterfactual, unobserved potential outcome for units that actually received treatment in the left panel. This is the outcome that treated units would have experienced (on average) if they had not been treated. Notice that the dashed-blue and solid-red lines are parallel.

The treatment effect on the treated is the difference between the solid-blue and dashed-blue lines. The difference is the difference between the average treated outcome and the average control outcome for units that actually received treatment. In the left panel of Appendix Figure D.13, the treatment effect—called the *average treatment effect on the treated* ("ATT")—is shown by the black arrow labeled "Treatment Effect." Notice that in this example, the ATT is constant across the post-1950 years, but this is not a requirement of the DID strategy.

[6] See Goldsmith-Pinkham (2026).

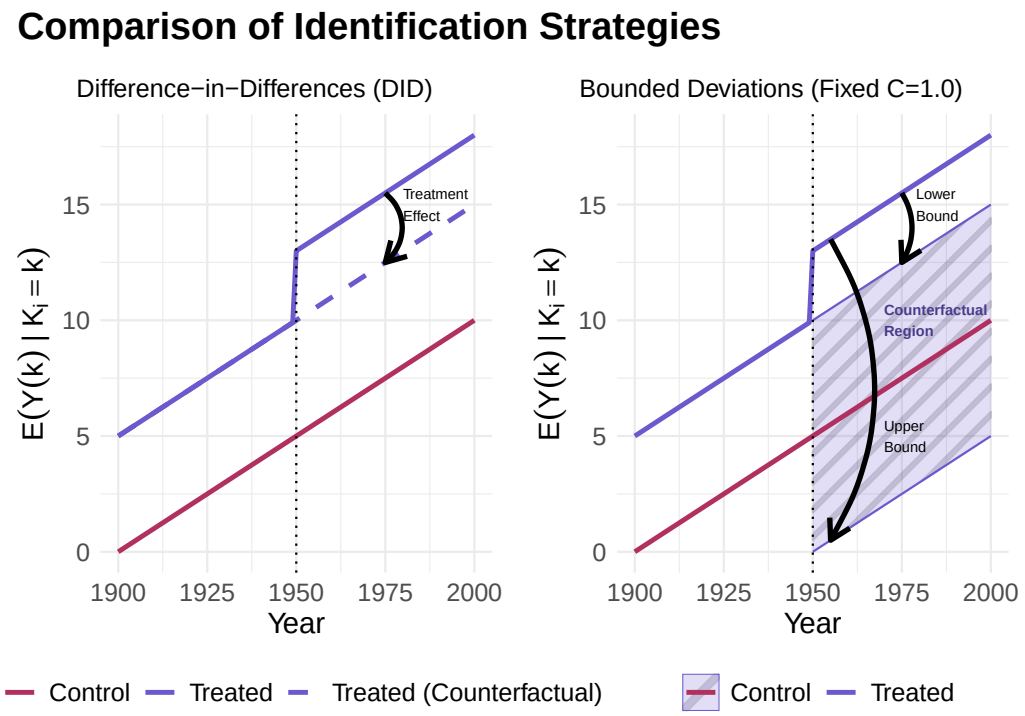

Figure D.13: Bounded Deviations & DID Comparison

Note: The black dotted vertical line is the path divergence time period (T_{div} = 1950). The blue-striped region is the counterfactual region.

Our method in this simple setting is depicted in the right panel of Appendix Figure D.13. The observed average outcomes are, as before, plotted in the solid-blue and solid-red lines. The difference here is that we allow for the counterfactual (i.e., untreated) average outcome for the treated group—which before was the dashed-blue line—to not evolve in parallel with the control group average outcome. Indeed, we allow for the counterfactual average outcome to differ from the parallel counterfactual in the DID provided it remains "close" to the control group outcome.[7] In the right panel of Appendix Figure D.13, this is shown by the blue-striped "Counterfactual Region." Our method allows the counterfactual (i.e., untreated) average outcome for the treated group to lie above or below the control group average outcome (solid-red line) so long as it remains within the blue-striped region.

The treatment effect on the treated (ATT) under our method is bounded above and below. It is bounded below by taking the difference between the observed average treated outcome and the largest possible counterfactual average outcome. In the

[7]For much more discussion of how to relax the parallel trends assumption, see, e.g., Rambachan and Roth (2023).

right panel of Appendix Figure D.13, this is shown by the arrow labeled "Lower Bound." The upper bound is found similarly: take the difference between the observed average treated outcome and the smallest possible counterfactual average outcome (labeled "Upper Bound" in Appendix Figure D.13). Finally, notice that the bounds in our example nest or contain the DID estimated treatment effect. That is, if parallel trends does hold, then the DID-based estimate is equal to the Bounded Deviations lower bound (for $C = 1$).

Estimation & Inference

We also compute all bounds as a function of C, which facilitates checking how conclusions change as a function of C. The estimation strategy proceeds by first partitioning the units by observed treatment paths. These paths will make up the counterfactual paths of interest. Next, for any two path pair of interest $(\vec{k}, \vec{k}')$ we compute the divergence period T_{div} and the maximum deviation in time periods before T_{div}. We then estimate the counterfactual path using the identification assumption for different C values. Although the primary focus of the methodology here is on identification, we note that, where possible, one can estimate the conditional expectations non-parametrically and use bootstrap techniques to compute standard errors. Inference is a difficult problem, however, in cases where there is only one treated unit (as in synthetic control), so we leave further estimation and inference issues for future work.

D.5 Bounded Deviations Tuning Parameter

Appendix D.5 show how the Bounded Deviations identification strategy depends on the tuning parameter C. As shown in C. Manski and Pepper (2018), researchers can study how sensitive their causal estimates are to tuning parameters. We show how this identification tuning parameter affects the upper and lower bounds in Appendix Figure D.14. The upper and lower bounds are plotted in the right and left panels, resp., for different values of C. In the plot, we consider varying C from being very small (yellow colored lines) to very large (dark blue lines). Smaller C values result in the upper and lower bounds being very close. For example, the yellow lower bound and yellow upper bound lines do not differ by much. By contrast, large C values result in larger counterfactual regions. This is shown by how far apart the lower and upper bounds are when $C = 3$ (i.e., dark blue lower bound and dark blue upper bound are far apart).

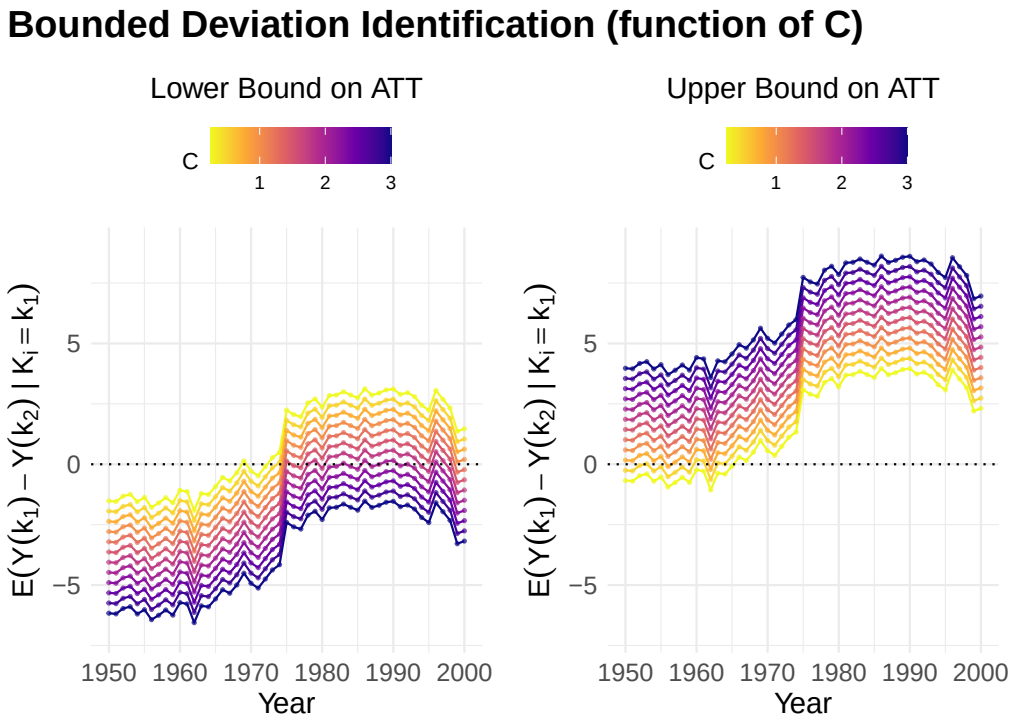

Figure D.14: ATT Bounds as a Function of C

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$.

D.6 Additional Results Plots

Appendix D.6 contains additional results plots. The figures are ATT bounds plots as above for different values of the tuning parameter C. Please note in each plot the cohorts that are being compared. Unlike the ATT bounds discussed in-text, these are comparisons between treated cohorts.

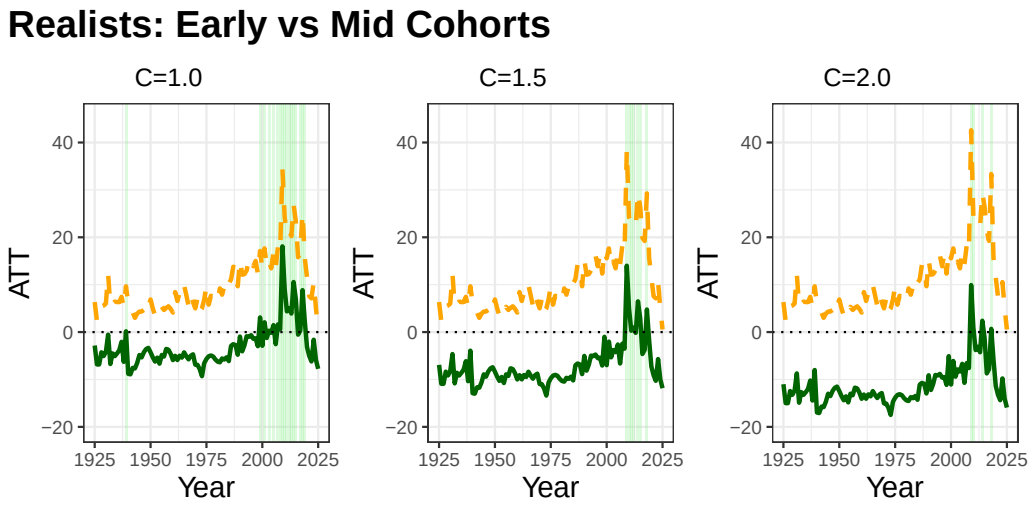

Figure D.15: Realists: Early vs Mid Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

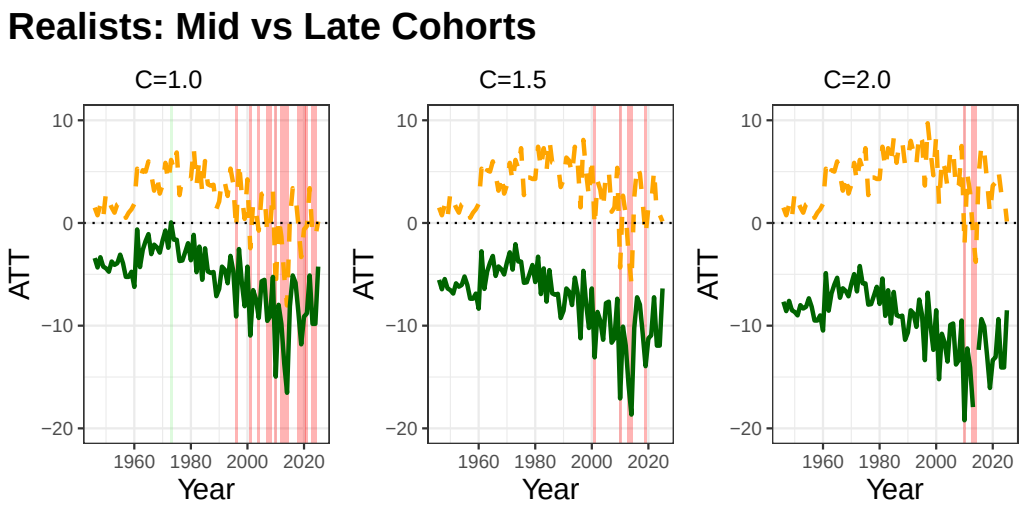

Figure D.16: Realists: Mid vs Late Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

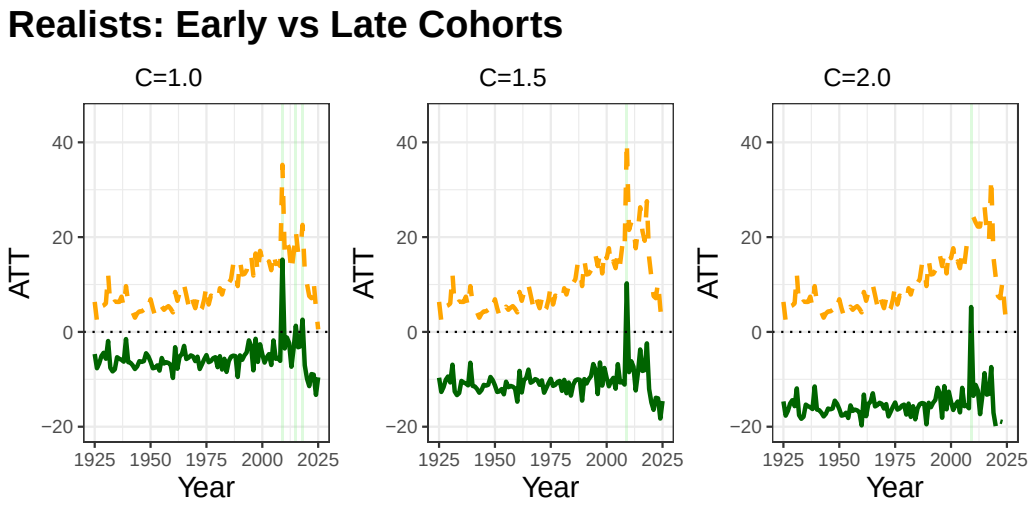

Figure D.17: Realists: Early vs Late Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

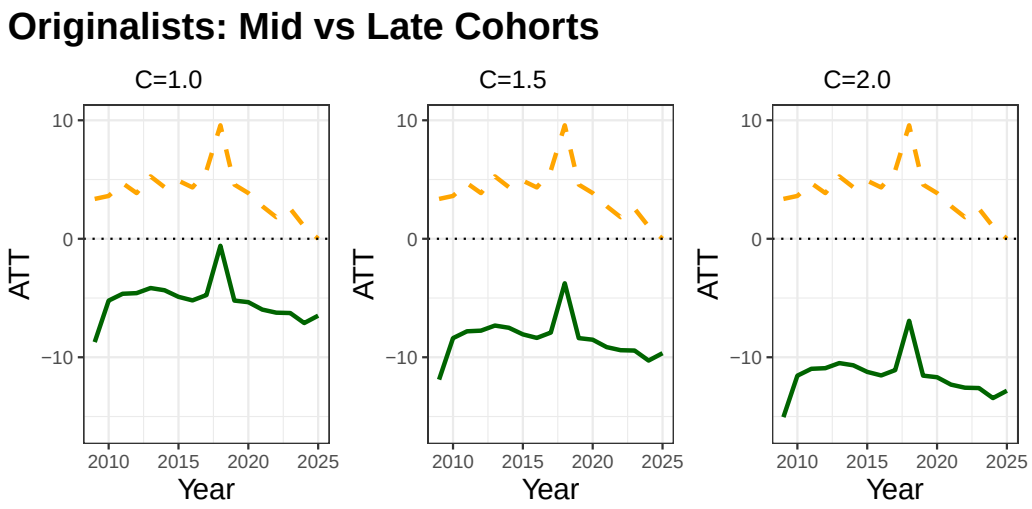

Figure D.18: Originalists: Mid vs Late Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

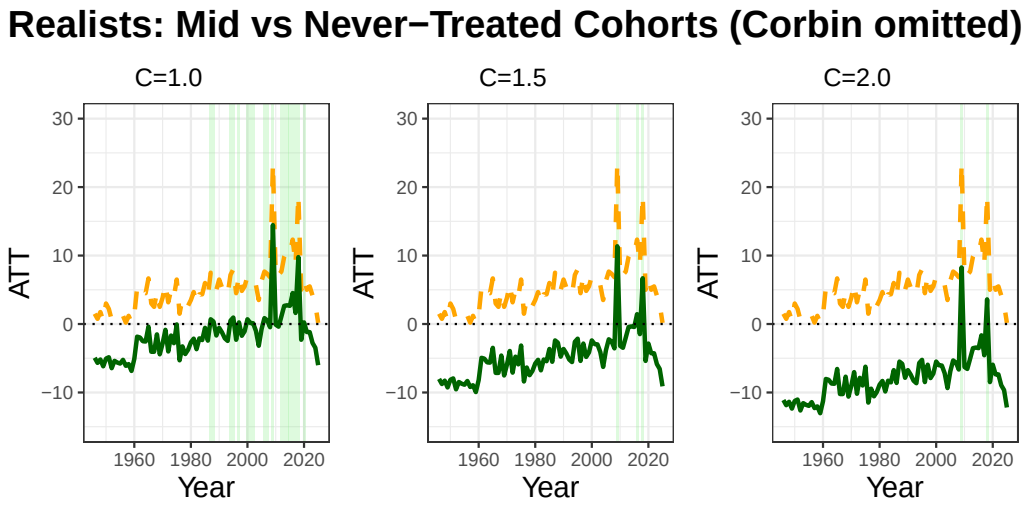

Figure D.19: Realists: Mid vs Never-Treated Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.

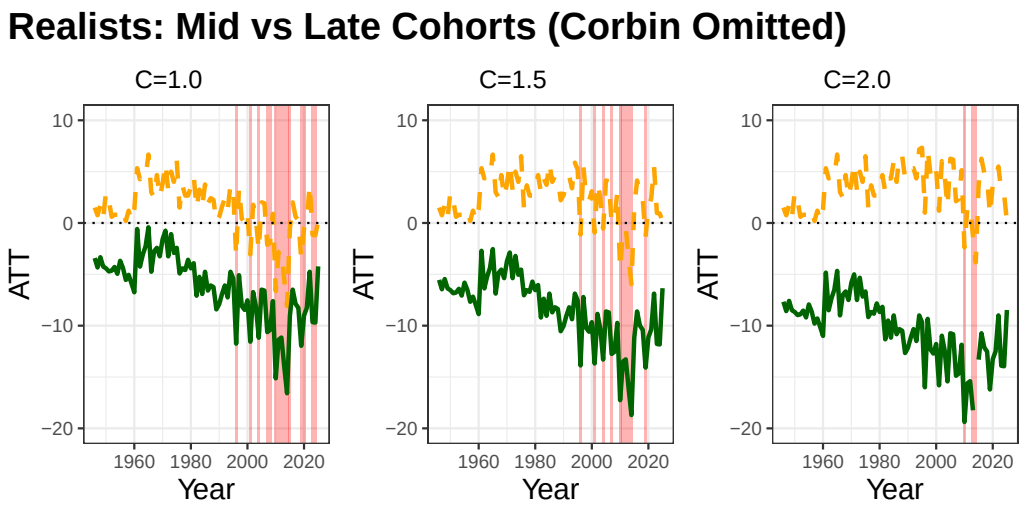

Figure D.20: Realists: Mid vs Late Cohorts

Note: In each panel, the black dotted horizontal line is plotted for $ATT = 0$, the ATT upper bound is the dashed-yellow line, and the ATT lower bound is the solid-green line. The years for which the upper and lower bounds are both positive are shaded green. The years for which the upper and lower bounds are both negative are shaded red.