

Essays on Government Policy in Real Estate Markets

David Joseph Munroe

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2014

©2014

David Joseph Munroe

All Rights Reserved

ABSTRACT

Essays on Government Policy in Real Estate Markets

David Joseph Munroe

This dissertation uses administrative data to study regulatory issues in the American real estate market. The first chapter studies spillovers from home foreclosures in Cook County, Illinois. Random assignment of foreclosure cases to judges allows for estimation of the causal effect of foreclosure (relative to a foreclosure case being dismissed) on neighboring foreclosure filings and housing transactions. When a property forecloses, the local housing market is disrupted—prices fall and more lower quality homes sell—and neighbors are more likely to end up in default and going through the foreclosure process.

The second chapter examines how discontinuously applied transfer taxes distort the market for real estate sales in New York and New Jersey. These transfer taxes distort not only the price of real estate transactions that occur near the discontinuity, corresponding to sellers bearing the entire incidence of the tax, but also the volume of sales that occur—productive transactions that would occur if the tax were not discontinuous disappear from the market.

The third and final chapter estimates the market-level response of home equity loans to two discontinuous mortgage policies—the home mortgage interest deduction, and real estate appraisal regulations in the Financial Institutions Reform Recovery and Enforcement Act. The estimates therein imply that home equity debt is very responsive to both the after-tax interest rate as well as lender underwriting requirements.

Contents

1 Foreclosure Contagion: Measurement and Mechanisms

(with Laurence Wilse-Samson)	1
1.1 Introduction	3
1.2 Judicial Foreclosure in Illinois	10
1.3 Data	13
1.4 Empirical Strategy	17
1.4.1 Measuring Local Contagion and Prices	20
1.4.2 Instrumental Variables Approach and First Stage Regression	22
1.4.3 Interpretation of the Two-Stage Least Squares Estimate	27
1.5 Neighborhood-Level Effects of Completed Foreclosure	30
1.5.1 Contagion in Foreclosure Filings	31
1.5.2 Contagion in Completed Foreclosures	35
1.5.3 Housing Markets	37
1.6 Evidence of Contagion Mechanisms	40
1.6.1 Distinguishing Borrower and Lender Response	40
1.6.2 Foreclosure Contagion and Negative Equity	44
1.6.3 Foreclosure Contagion and Information	48
1.7 Conclusion	51
1.8 Appendix	64
1.8.1 Data Appendix	64
1.8.2 Adjusting Price Data for Property Quality Using Repeat Sales	69
1.8.3 Monotonicity of Instrument	70
1.8.4 Nonlinearities in Foreclosure Contagion	71

1.8.5	Lender Response to a Completed Foreclosure	73
1.8.6	Estimates by Proxies for Social Connectedness	77
2	Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate	
	Market (with Wojciech Kopczuk)	103
2.1	Introduction	105
2.2	Policy	111
2.3	Data	113
2.4	Theoretical Framework	115
2.4.1	Bargaining	116
2.4.2	Equilibrium	119
2.4.3	Measuring the Impact on the Price Distribution	121
2.4.4	Implications for Efficiency of the Equilibrium Allocation	129
2.4.5	Econometric Implementation	130
2.5	Distortion to the Price Distribution	132
2.6	Unraveling: Market Distortions Local to the Threshold	139
2.7	Global Market Distortions	141
2.8	Conclusions	144
2.9	Appendix	163
2.9.1	Bargaining model	163
2.9.2	Proportional tax	166
2.9.3	Data Appendix	169
2.9.4	Robustness of incidence estimates	173
3	Response of Home Equity Debt to Mortgage Policy: Evidence from a Kink	
	and a Notch	197
3.1	Introduction	199
3.2	Policies	204

3.2.1	The Home Mortgage Interest Deduction	204
3.2.2	FIRREA Appraisal Requirements	206
3.3	Data	207
3.4	Methodology	209
3.4.1	Estimating Response to Policy from a Kink	209
3.4.2	Estimating Response to Policy from a Notch	212
3.4.3	Implementation	214
3.5	Results	218
3.5.1	Graphical Evidence of Response	218
3.5.2	Baseline Estimates and Robustness	220
3.5.3	Response of Debt to the Mortgage Interest Deduction	223
3.5.4	Response of debt to FIRREA Appraisal Requirements	228
3.6	Conclusion	229
3.7	Data Appendix	243
	References	247

List of Tables

1.1	Descriptive Statistics: Pre-Treatment Characteristics	54
1.2	Descriptive Statistics: Outcomes	54
1.3	Balance of Covariates	55
1.4	Complier Characteristics: Ratio of Subgroup First Stage Estimate to Overall First Stage	55
1.5	Baseline Contagion Estimates: 2SLS Coefficient of Effect of Completed Fore- closure on Given Outcome in Given Year	56
1.6	Constant Sample Contagion Estimates	57
1.7	Estimates in Years Before Filing	58

1.8	Contagion in Completed Foreclosures	59
1.9	Baseline Housing Market Estimates	60
1.10	Contagion Among Loans with Lenders Implicated in Independent Foreclosure Review Settlement	61
1.11	Contagion Estimates By Proxy for Borrower Equity	62
1.12	Contagion Estimates By Lender Identity	63
1.13	First Stage Regression of Foreclosure on Propensity to Foreclose	83
1.14	Contagion Estimates Measured Since Case Filing	83
1.15	Controlling for Length of Case	84
1.16	Robustness of Contagion Estimates By Specification	85
1.17	Robustness of Estimates by Sample	86
1.18	Contagion in Any New Filing, Omitting Each Filing Year	87
1.19	Contagion in Total New Filings, Omitting Each Filing Year	88
1.20	Baseline Estimates for 0.25 Mile Radius	89
1.21	Estimates for Loans with Multiple Claimants	90
1.22	Contagion in Completed Foreclosures/Dismissals Filed After Decision	91
1.23	Robustness of Price Estimates By Specification	92
1.24	Price Effects, Omitting Each Filing Year	93
1.25	Sub-Sample Price Effects	93
1.26	Contagion Among Loans with Lenders Implicated in Independent Foreclosure Review Settlement	94
1.27	Contagion among Borrowers in Positive vs. Negative Equity	95
1.28	Interaction with Positive Price Growth	96
1.29	Response Among Borrowers with Conventional Mortgages	97
1.30	Contagion among Filings with -10% – 10% Equity versus < -10% or > 10%	98
1.31	Contagion by Proxy for Neighborhood Homogeneity	99
1.32	Contagion by Condo Status	100

1.33	Contagion Among Borrowers with the Same Lender	101
1.34	Neighborhood-Level Aggregate Contagion Regressions	102
2.1	Real Estate Transfer Tax Schedules	148
2.2	Sample Statistics for Taxable Sales	148
2.3	Median Price of Taxable Sales Over Time	149
2.4	REBNY Listings Sample Statistics	150
2.5	Response to Mansion Tax, by Region and Years Since Construction	151
2.6	Heterogeneity in Response by Notch and Sub-Sample	152
2.7	Mansion Tax: Listings	153
2.8	Mansion Tax: NYC	192
2.9	NYC Mansion Tax: Placebos	193
2.10	Local Incidence Over Time	194
2.11	Predicted Price Discounts	195
2.12	Predicted Price Dispersion	196
3.1	Home Loan Statistics (1995–2008)	238
3.2	Baseline Estimates and Sensitivity	239
3.3	Round Number Placebos	240
3.4	Estimates by Quartile of AGI	241
3.5	State-Year Regressions of $1 - \phi_k$ on Percent-Change in After-Tax Interest Rate	242
3.6	Implied Elasticity of Debt to Interest Rates at \$100,000	242

List of Figures

1.1	Foreclosure Cases Over Time in Cook County, IL	53
1.2	Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates	78
1.3	Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates	79
1.4	Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates	80

1.5	Contagion Estimates by Income Quartile	81
1.6	Contagion Estimates by Social Connection Proxy	82
2.1	Distribution of Taxable Sales in New York State	154
2.2	Distribution of Taxable Sales in New York City	155
2.3	Distribution of NJ Sales Pre- and Post-Mansion Tax	156
2.4	New Jersey Monthly Sales Above \$990,000	157
2.5	Incidence and Gap Concepts	158
2.6	Bunching at the notch and efficient allocation	158
2.7	Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only) . .	159
2.8	Median & 75th Percentile Price Discounts by Initial Asking Price	160
2.9	Predicted Price Discounts by Initial Asking Price (Relative to \$1,000,000) . .	161
2.10	Predicted Dispersion of Log of Sale Price by Initial Asking Price (Relative to \$1,000,000)	162
2.11	Distribution of Sales in New York City around the \$500,000 RPTT tax notch	176
2.12	Distribution of Taxable Sales in New York State	177
2.13	New Jersey Monthly Sales Above \$990,000	178
2.14	Distribution of Monthly Sales in New Jersey (\$900k – \$1M)	179
2.15	NJ Local Incidence Over Time	180
2.16	Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only) . .	181
2.17	Distribution of Real-Estate Listing Prices in NYC (All REBNY-Listed Prop- erties)	182
2.18	Distribution of Sale Price by Initial Asking Price (with Quantile Regression)	183
2.19	Probability that Listed Property Sells by Initial Asking Price	184
2.20	Median Days to Sale by Initial Asking Price	185
2.21	Probability of Selling Without REBNY by Initial Asking Price	186
2.22	Median Price Discount by Final Asking Price	187
2.23	Decomposition of Mean Price Discounts by Initial Asking Price	188

2.24	Dispersion of Sale Price, Conditional on First Asking Price	189
2.25	Dispersion of Sale Price, Conditional on Last Asking Price	190
2.26	Predicted Dispersion of Log of Sale Price (Median Regression in First Stage)	191
3.1	Data Over time	231
3.2	Conceptual Kink Figures	232
3.3	Conceptual Notch Figures	233
3.4	Histogram of Home Loans	234
3.5	Log of Histogram of Home Loans	235
3.6	Kink Estimates Over Time	236
3.7	\$250k Notch Estimates Over Time	237
3.8	County Start Dates	245
3.9	Loans over Time by Sample	246

Acknowledgements

Firstly, I would like to acknowledge the support and advice of my dissertation committee. I am thankful to Wojciech Kopczuk for exceptional guidance. Wojciech introduced me to empirical public economics and provided invaluable feedback on my research. His critical eye and enthusiasm always pushed me to do my best work. I am also indebted to Ethan Kaplan for holding me to, and helping me achieve, high empirical standards. Like Wojciech, Ethan gave me the confidence to trust my instincts and my training. Douglas Almond, Miguel Urquiola, and Tomasz Piskorski all provided exceptional feedback and encouragement.

I am fortunate to have been a part of a rich academic environment at Columbia. In addition to the guidance of my dissertation committee, I received valuable comments from current and former Columbia faculty including Pierre-Andre Chiappori, Janet Currie, Ilyana Kuziemko, Chris Mayer, Christian Pop-Eleches, Bernard Salanie, Eric Verhoogen, and Till von Wachter. I am grateful for the feedback of my fellow students Corinne Low, Katherine Meckel, Mike Mueller-Smith, Giovanni Paci, Petra Persson, Maya Rossin-Slater, Jessica Van Parys, and Reed Walker, and to Thuy Lan Nguyen, Wataru Miyamoto, Jeong Hwan Lee, and Tao Li for helping me get through the first year of the program. In addition to sharing a passion for empirical public economics with me, Ferran Elias and Ben Marx provided hours of valuable discussion. The origin of my interest in real estate topics can be traced directly to my colleague, friend, and co-author Laurence Wilse-Samson; this dissertation would look very different were it not for his support. The social support and friendship of Jonathan Dingel, Donald Ngwe, Hyun Oh, Sebastien Turban, and Harold Stolper ensured that I did not work too hard and occasionally enjoyed myself.

I am grateful for the resources that have been made available to me as a doctoral student. I would like to acknowledge the financial support of the Social Sciences and Humanities

Research Council of Canada, the Lincoln Institute of Land Policy, and the Program for Economic Research at Columbia University. Access to incredible data was made possible by Chris Mayer and the Paul Milstein Center for Real Estate at Columbia Business School, Chicago-based Record Information Services Inc., and the Real Estate Board of New York and RealPlus, LLC. Special thanks to Alicia Horwath of Record Information Services and Thomas Croke of RealPlus for assistance with the data.

Finally, I would like to thank my family. My parents, Robert and Sheila, have been incredibly supportive throughout a decade of post-secondary schooling. They have always encouraged me to follow my interests and do my best. I am forever indebted to my incredible wife, Chrissy. Chrissy has made me feel loved through many work-related difficulties and countless research-filled evenings and weekends. She is a constant source of fun and joy in my life. This thesis would not have been possible without her inspiration and support.

Dedication

To Chrissy, for believing in me.

1 Foreclosure Contagion: Measurement and Mechanisms

(with Laurence Wilse-Samson¹)

¹Department of Economics, Columbia University. lhw2110@columbia.edu.

Abstract

This paper shows that completed foreclosures cause neighboring foreclosure filings. We estimate this relationship using administrative data on home foreclosures and sales in Cook County, IL, instrumenting completed foreclosures with randomly assigned chancery-court judges. A completed foreclosure causes 0.5 to 0.7 additional foreclosure filings within 0.1 miles, an effect that persists for several years. Contagion is driven by borrowers on the margins of default, not those severely at risk. We find evidence that borrowers learn about lender behavior from neighboring foreclosures. Finally, a foreclosure causes an increase in housing sales among relatively low-quality properties.

1.1 Introduction

The housing bubble and crisis of the last decade has resulted in an unusually large number of foreclosures in the United States. Completed foreclosures—when a mortgage borrower does not make payments on their loan and the lending institution claims the mortgaged property—increased dramatically starting in 2007 from 404,849 properties per year, peaking at 1.05 million completed foreclosures in 2010.² The length and severity of this crisis have increased academic interest in the consequences of home foreclosures and have raised questions about how and why foreclosures spread (e.g., Guren and McQuade (2013)).

In this paper we ask whether home foreclosures are contagious: does one completed foreclosure increase the probability that geographically neighboring borrowers end up in the foreclosure process? The answer to this question informs our understanding of home foreclosures, borrower and lender behavior, and appropriate policy toward mortgages and foreclosure procedures. Foreclosure contagion is suspected of exacerbating the housing crises during the Great Depression and the recent financial crisis (Campbell (2013)). Identifying and understanding contagion in foreclosures will provide a better understanding of how and why such crises spread. Furthermore, the presence of contagion is relevant to policy makers concerned with mitigating the spread of home foreclosures.

Our chief contribution is to develop a randomly assigned instrument for foreclosures, which we apply to administrative data to achieve credible, policy-relevant estimates of foreclosure contagion. To our knowledge, ours is the first study to use a randomly assigned instrument to study the local effects of foreclosures. In Chicago (Cook County), IL, where foreclosure cases are decided in court, we use the randomization of new cases to fixed groups of judges as an instrument for a completed foreclosure. Intuitively, our estimates compare the neighborhoods around two types of properties going through the foreclosure process (i.e.,

²We use “foreclosure filing” to refer to the initiation of the foreclosure process by the lender, and “completed foreclosure” to refer to a foreclosure proceeding ending with the mortgaged property being sold at auction. However, lenders are not always successful in foreclosing on a home, and so not all filings end in completed foreclosure—we refer to such unsuccessful foreclosure attempts as “dismissals.”

situations in which a borrower is in default and the lender wants to claim the home as collateral): properties randomly assigned to “difficult” judges that, as a result, are foreclosed upon and sold at auction versus properties randomly assigned to “lenient” judges that dismiss the foreclosure case. Since our empirical strategy necessarily relies upon the comparison of neighborhoods around homes in default that do and do not end in foreclosure, our estimates speak directly to the policy question of how strongly lenders should be incentivized to renegotiate delinquent loans.³ This instrument allows us to surmount the endogeneity of home foreclosures—a key empirical challenge—present in existing studies of foreclosure externalities that primarily rely on fixed effects analyses of very local neighborhoods (e.g., Campbell et al. (2011), who study how real estate sale prices are influenced by foreclosures within 0.1 miles).

We develop a novel data set that matches administrative records of foreclosure *court cases* to records on foreclosure *filings and auctions* for Cook County. This county, which contains most of the city of Chicago, is the second-most populous in the U.S. and was relatively hard hit by the housing crisis: the surrounding MSA experienced the 12th largest decline in city-wide housing prices between 2007 and 2011, while 5.2% of the 1.9 million households in Cook County experienced a completed foreclosure. Our data covers the universe of foreclosure filings and completed foreclosures in Cook County between 2004 and 2011, allowing us to leverage the random assignment of foreclosure judges while observing the precise location of the associated property. We also use administrative data on residential housing sales to assess whether a completed home foreclosure lowers neighboring housing values.

Concrete evidence on foreclosure-related externalities and contagion has been elusive, owing to empirical challenges (Frame (2010)). Home foreclosures are known to be corre-

³There is a developed literature that uses judicial bias as an instrument, as we do herein, including: Kling (2006) (sentencing propensities of judges to instrument for incarceration length); Autor and Houseman (2010) (job placement rates of non-profit contractors to instrument for receiving temporary help jobs); Chang and Schoar (2006) (judicial fixed effects to measure judge-debtor-friendliness); Dobbie and Song (2013) (judge discharge rates to instrument for bankruptcy protection); Doyle (2007) (placement frequency of child protection investigators to instrument for foster care); and Maestas et al. (2013) (allowance rates of disability examiners to instrument disability insurance receipt).

lated with neighborhood characteristics and changes in housing prices and macroeconomic circumstances (Mian et al. (2011); Mian and Sufi (2009)). Existing studies that find negative housing price effects of foreclosure have relied primarily on local analyses that explicitly control for property and neighborhood characteristics (Campbell et al. (2011); Immergluck and Smith (2006); Schuetz et al. (2008); Pennington-Cross (2006); Leonard and Murdoch (2009); and Lin et al. (2007)) or repeat-sales analyses (Harding et al. (2009) and Gerardi et al. (2012)). Similarly, existing studies finding evidence of foreclosure contagion rely either on local analyses (Towe and Lawley (2013)) or aggregate analyses controlling for neighborhood and zip code characteristics (Goodstein and Lee (2010)). Few studies have taken a quasi-experimental approach to identifying the externalities associated with foreclosure.⁴ One notable exception is Anenberg and Kung (2014), who find a drop in real estate listing prices immediately after foreclosed properties are listed on the market. To our knowledge, ours is the first study to use a randomly assigned instrument to estimate contagion and local price effects of home foreclosures.

We find evidence of foreclosure contagion using our instrumental variables strategy that compares neighborhoods with completed foreclosures to neighborhoods where foreclosure cases are dismissed. Relative to dismissal, a completed foreclosure raises the probability of any new foreclosure filing within 0.1 miles by 10% per year and leads to about 0.5 new filings per year. This foreclosure contagion effect is robust and persistent, lasting for three to four years after the case is decided. Additionally, our estimates show that substantial contagion in foreclosure filings occurs even in neighborhoods with no recent foreclosures—the “first” completed foreclosure in a neighborhood substantially increases foreclosure filings in the following years. Interestingly, there is no evidence of contagion at the height of the crisis (2009–2011); contagion is present primarily during the peak and initial decline of the Chicago housing market (2004–2008). We interpret this temporal pattern as evidence that contagion

⁴Mian et al. (2012) exploit changes at state borders in policy toward foreclosure (in particular, the distinction between judicial vs. non-judicial states) to instrument foreclosures, although this instrument is not randomly assigned (for example, Pence (2006) shows substantial changes in housing market conditions at the boundaries between judicial and non-judicial foreclosure states).

operates through borrowers who are on the margin of the default decision (2004–2008), rather than those in dire straits (2009–2011). A neighboring completed foreclosure may not be very meaningful to a borrower who is already in negative equity (thus, relatively insensitive to a foreclosure-induced loss of property value) or when foreclosures are common-place (and so the marginal foreclosure conveys little information).

Contagion is not limited to new foreclosure filings—a completed foreclosure increases the number of neighboring completed foreclosures as well. This result suggests that contagion is costly and plays a role in the geographic spread of foreclosures. In particular, we find that, on average, each completed foreclosures induces an additional 1.5 completed foreclosures within four years. Taken literally, our estimates suggest that in the absence of the contagion externality, Chicago would have experienced more than 50% (roughly 43,000) fewer completed foreclosures between 2004 and 2010.⁵ We also find that a completed foreclosure increases the number of completed foreclosures even among neighboring cases that had already begun.

While our estimates suggest that a completed foreclosure lowers neighboring residential sale prices, our estimates are largely driven by selection into sale. Within the first year of a case ending in a foreclosure, *relative to a case that ends in dismissal*, the average price of neighboring housing sales drops by up to 40%. However, using a repeat-sales methodology to adjust for property quality, our estimates of this effect fall substantially to zero. We interpret this as evidence that a completed foreclosure disrupts the housing market in terms of the types of homes that sell, causing a larger share of lower-quality homes to transact at correspondingly lower prices than the average home in the neighborhood. At the same time, due to the small size of our housing sales sample and the resulting imprecision of our housing price estimates, we cannot rule out a small negative effect of completed foreclosure on neighboring home values (holding quality constant).

We show evidence consistent with the commonly held belief that foreclosure contagion

⁵These figures are based on a back of the envelope calculation. Our estimates suggest that each completed foreclosure causes more than 1 additional completed foreclosure within three years, and so an absence of contagion suggests 50% fewer completed foreclosures.

is driven by an increase in borrowers defaulting on their loans in response to a neighboring foreclosure, rather than lenders filing for foreclosure against already delinquent borrowers. There is substantial evidence that lenders and mortgage servicers—third parties employed by creditors to manage loans—indiscriminately favor pursuing foreclosure on delinquent mortgages, rather than modification (see discussions in Adelino et al. (2009); Foote et al. (2008) and Levitin and Twomey (2011)). We argue that in the absence of borrower-driven contagion, lenders would not exhibit positive foreclosure contagion. Given that we do observe positive contagion provides evidence that borrowers respond. Moreover, there is substantial contagion even among mortgages serviced by lenders known for automating foreclosure procedures and who are, thus, unlikely to respond to very local market conditions.

There are two prominent explanations of why a completed foreclosure will increase the probability that neighbors default on their own mortgages. The first hinges on a completed foreclosure lowering neighboring home values, thus increasing the likelihood that borrowers are in negative equity or “underwater” on their loans—i.e., owing more than the mortgaged property is worth (Campbell and Cocco (2011); Campbell (2013); and Goodstein et al. (2011)). As one becomes further underwater on one’s loan, the incentive to default on the mortgage increases: the loss on the asset grows relative to the costs associated with foreclosure (primarily moving costs and a drop in credit score). The second explanation is that a completed foreclosure transmits information to neighbors (Guiso et al. (2013); Towe and Lawley (2013)). Specifically, a completed foreclosure may send a signal to neighbors about the future of the neighborhood (influencing the expected value of the property to the borrower), or about the foreclosure process itself (e.g., neighbors may learn about the likelihood of a mortgage modification if they default on their loans).

Contagion is driven not by borrowers in severe negative equity, but by borrowers who are on the margin of being underwater. We use information about loan principal and outstanding balance to construct a proxy for borrowers being underwater and find that a completed foreclosure induces additional new filings among non-underwater loans only. Moreover, this

effect is concentrated among loans who are on the margin of being underwater with outstanding debt at the time of filing within 10% of the initial principal. For example, a neighboring foreclosure may act as a “wakeup call” for borrowers in positive equity, sending a strong signal about the current value of their property, the future of the neighborhood, and/or information about the foreclosure process itself. On the other hand, those who are very underwater on their loans may already be well informed about the foreclosure process and the consequences thereof and/or have a sufficiently large negative equity position that an additional loss in value is negligible with respect to the default decision.

Contagion varies substantially depending on whether neighbors have mortgages serviced by the same lender, which we interpret as evidence that information—in this case, information about lenders—plays an important role in reducing contagion. Specifically, we find that when a completed foreclosure occurs there are significantly fewer new foreclosure filings among loans serviced by the same lender than loans serviced by different lenders. This difference may be driven by borrowers learning more and different information from the experience of neighbors whose loans are serviced by the same institutions: the neighboring foreclosure may send a signal about their lender’s behavior, lowering the perceived probability of a successful renegotiation of the loan and, thus, reducing strategic incentives to default.

Our results suggest that policies that keep delinquent borrowers in their homes, for example by encouraging lenders to modify delinquent loans, may reduce the spillovers associated with home foreclosures. We are able to speak to this question since our empirical strategy identifies foreclosure spillovers for the set of marginally delinquent loans for whom the idiosyncrasies of the overseeing judge matter (i.e., the cases that are likely to comply with the instrument)—these are the cases most likely to be influenced by policy interventions. In principle, our results provide support for government policies that encourage modification, such as the Treasury’s Home Affordable Modification Program (HAMP) (even if there have been considerable (and well documented) problems in implementation). For example, as of December 2013 HAMP has achieved 27,525 successful modifications of delinquent mort-

gages in the Chicago MSA. However, this understates the true benefit of the program—our contagion estimates imply that HAMP prevented an additional 44,040 completed foreclosures (and the spillovers associated with these).⁶ Extrapolating our contagion estimates to the national level, by successfully modifying 1.1 million delinquent mortgages, HAMP has prevented an additional 1.76 million completed foreclosures. Of course any program must weigh general equilibrium considerations—these include the effects on ex-ante incentives for loan origination, and incentives for default by other borrowers. Our finding that contagion is minimal or even negative among borrowers with the same creditors, may provide evidence that these incentives for strategic default matter—we interpret this finding as demonstrating that borrowers update the probability of a modification downward and are discouraged from defaulting on their loans. As such, a policy that raises the cost of default (e.g., achieving a reduction in loan principal only by going through bankruptcy, as suggested by Levitin (2009)), or a more direct policy that targets the vacancy and neglect associated with REO properties, may be preferred.

The rest of the paper proceeds as follows: in Section 1 we outline the judicial foreclosure process in the state of Illinois and randomization of judges to cases in Cook County. Section 2 sketches our data sources—administrative records on court cases, geocoded administrative records on foreclosure filings, and deed transfer records. In Section 3 we outline the empirical strategy, which exploits the random assignment of judges for quasi-experimental identification, the results of which we present in Section 4. Section 5 explores possible mechanisms, and Section 6 concludes.

⁶HAMP numbers are from the December 2013 Making Home Affordable Program Performance Report from the U.S. Department of the Treasury (available at www.treasury.gov/initiatives/financial-stability/reports/Pages/Making-Home-Affordable-Program-Performance-Report.aspx). The progress report identifies 46,183 permanent loan modifications for the Chicago-Joliet-Naperville MSA. The report indicates a 40.4% redefault rate (within four years) among HAMP modified loans, leaving 27,525 successful modifications. Summing our estimates of contagion in completed foreclosures from Table 1.8 suggests that each completed foreclosure induces, on average, an additional 1.6 completed foreclosures (and $1.6 * 27,525 = 44,040$).

1.2 Judicial Foreclosure in Illinois

Cook County, Illinois, provides a good context in which to study foreclosure contagion. Firstly, it was badly affected by the foreclosure crisis. Between 2002 and 2011, the county saw 302,166 foreclosure proceedings initiated by lenders (“foreclosure filings”), and 134,924 completed foreclosures. These trends are illustrated in Figure 1.1. There is a sharp increase in the number of foreclosure cases filed in Cook County (left axis) from about 1,000 per month in 2004 to more than 3,000 filings per month in 2008. At the same time, foreclosure proceedings became more likely to end in a completed foreclosure: the completed foreclosure rate (right axis) jumps from 45% for cases filed in 2004 to 65% in 2008. Secondly, the foreclosure process in Cook County, IL, goes through the court system, allowing us to instrument a foreclosure outcome using random assignment of judges to cases.

In Illinois, as in many so-called “judicial foreclosure” states, lenders must take delinquent borrowers to court in order to claim a mortgaged property. When a borrower has missed three mortgage payments (i.e., is in default), a lender or the third party servicing the mortgage may initiate the foreclosure process by filing for foreclosure on the associated property with the chancery court (we refer to this event as a **foreclosure filing**). If after ninety days the borrower has not made up all missed payments, the trial begins and the lender’s attorney must establish that the borrower: has borrowed money from the lender; has signed a mortgage note promising the property as collateral; and is behind on payments. At the same time, the borrower may mount a defense, for example by disputing any of these facts or claiming that the lender has violated lending laws (e.g., the Truth-in-Lending Act). After hearing the arguments, the presiding judge decides the case, either dismissing the foreclosure action or filing a judgment of foreclosure. If the case is dismissed, the borrower typically continues to reside in the home. If a judgment of foreclosure is filed, then the case proceeds to a foreclosure auction, which we refer to as a **completed foreclosure**.⁷ If the sale price

⁷Following a foreclosure judgement, a redemption period begins during which the borrower may pay off the entire outstanding mortgage plus late fees, attorney fees, court costs, and taxes. The redemption period ends either three months after the judgment or seven months after the initial foreclosure complaint is served,

does not cover the outstanding balance of the mortgage then the borrower is still considered in debt to the lender, although it is common for lenders to forgive this remaining debt. In the vast majority of cases (around 95% for Cook County), the lending institution purchases the property at auction for the amount of the outstanding loan—in doing so the lender need not record a loss on their balance sheet.⁸

A dismissal may refer to several possible outcomes, most of which result in the property remaining occupied by the borrower. First, if a borrower makes all missed payments within 90 days of the filing, then the case is dismissed and the mortgage is reinstated. Second, rather than continuing to pursue an ongoing foreclosure case, the lender may modify the terms of the mortgage to make payments more affordable to the borrower. Third, the lender may “lose” the case by failing to adequately establish non-payment of the mortgage or that they are owed the debt, or the borrower’s defense may be successful. Fourth, a case may be dismissed because the lender does not take action in pursuing the foreclosure. Fifth, a lender may accept a deed-in-lieu of foreclosure, in which the borrower forfeits the home to the lender without going through the courts. Finally, the borrower may negotiate a short sale of their home: the lender accepts the proceeds from the sale of the home as payment for the mortgage. Deed-in-lieu of foreclosure and short sales are generally not an option when there are multiple liens on the property, a fact we exploit to confirm that our results are not driven by these outcomes in which delinquent borrowers lose their property (Agarwal et al. (2011)).⁹ In our data we cannot distinguish which of these outcomes occurs; we only know whether the case ends in dismissal or completed foreclosure. However, with the exception of a deed-in-lieu of foreclosure or a short sale, in all of these dismissal outcomes the house

whichever is later.

⁸See statistics for Cook County compiled by the Woodstock Institute: blog.cookcountyil.gov/economicdevelopment/wp-content/uploads/2012/11/Wodstock-Institute-Foreclosure-Filings-2007-2012.pdf

⁹Anecdotally, deed-in-lieu of foreclosure and short sales are uncommon in Cook County for the reason that, in both cases, creditors are typically taking a loss, while mortgage servicers will accrue lower fees (relevant in cases where the property is being managed by a mortgage servicer): in Illinois, by accepting a deed-in-lieu of foreclosure the lender must forgive all debt, while short sales typically transact at a price below the outstanding debt (Ghent and Kudlyak (2011)).

remains occupied by the borrower.

Foreclosure cases are randomly assigned to a case calendar, which restricts the set of judges that will ever hear an action on the case. A case calendar is a weekly schedule of court-room/judge pairings, usually made up of two or three judges. Judges typically only hear cases associated with their case calendar. Similarly, since chancery court cases are only assigned to one calendar, only the associated judges will oversee an action on that case. When a case is filed, it is assigned a unique case number, sorted by property type (single-family home, condominium, commercial property, etc.), and randomly assigned to a case calendar.¹⁰ As of 2010, there were 12 chancery court case calendars hearing foreclosure cases (there are additional calendars that hear only other chancery court cases). Judges are assigned to case calendars each year by the Chief Judge of the Circuit Court of Cook County.

There are several ways that a Cook County judge might influence the outcome of a foreclosure case, which is necessary for the validity of our instrument.¹¹ Firstly, the judge has discretion to determine how long a defendant has to find a lawyer and mount a defense. Secondly, even if a defendant does not mount a defense, the judge determines whether or not the lender successfully establishes that the borrower is behind on payments and that the debt is owed to the lender. Establishing these points is not trivial. Throughout the foreclosure crisis, there have been accounts of mistakes and wrongdoing in the prosecution of foreclosures (Kiel (2012)), including failures of banks to produce proper documentation or lenders initiating foreclosure proceedings without reviewing the history of the loan (“robosigning”). Similarly, it is up to the judge to evaluate a borrower’s defense, for example by determining

¹⁰Random assignment of cases to case calendars is performed by the Chancery Court computer system. As described on the Chancery Court’s FAQ page, “When a case is filed in the Law Division it is randomly assigned via a computer program to a calendar letter. You may contact the Information Desk in Law Division to obtain the Judge’s information associated with the calendar letter.” See www.cookcountyclerkofcourt.org/?section=FAQSPage

¹¹There is substantial evidence of judicial bias in many settings: Anderson et al. (1999) illustrate important *differences across judges in decision-making*—sometimes suggestive of bias (e.g., Abrams et al. (2008) or Yang (2012)), and sometimes more generally based on “personal assessments” of case-specific information (Iaryczower (2009)). Berdejo and Chen (2010) present evidence suggestive of unconscious judicial bias—illustrating priming effects on judges of wars (which suppress dissents)—as well as more partisan behavior before Presidential elections.

whether a mortgage is legal in the first place (e.g., is not in violation of (predatory) lending laws). Anecdotal evidence suggests that judges vary substantially in their leniency on these issues.¹²

In what follows, we use the random assignment of foreclosure cases to case calendars to instrument the outcome that the case ends in foreclosure. As discussed above, judges may influence the outcome of a case. At the same time, the case calendar to which a foreclosure case is randomly assigned determines the possible judges who will ever hear the case. If judges vary sufficiently in their biases toward foreclosure, then the case calendar to which a case is assigned may influence whether a case ends in foreclosure or dismissal.¹³ Thus, our identification relies on the comparison of two delinquent borrowers going through the foreclosure process, one of whom is randomly assigned to a “lenient” case calendar and ends in dismissal, while the other is randomly assigned to a “strict” calendar and ends in foreclosure. To implement this study we require data on Cook County foreclosure cases, including case calendar assignment and the case outcome (foreclosure or dismissal).

1.3 Data

We use geocoded administrative data for Cook County from three sources: Cook County chancery court records, foreclosure filings, and deed transfer records. Publicly available chancery court records for 2004–2010 provide us with details of each foreclosure case, including the information necessary to construct our instrument: the case calendar to which the case is randomly assigned and the outcome of the case (dismissal or foreclosure). To study neighborhood outcomes, however, we need to know the location of borrowers’ homes.

¹²The *Washington Post* observes, for example, that “[in] Suffolk and Nassau counties on Long Island and Kings County... which have among the highest rates of foreclosure in the state and where the 81 judges handling foreclosures have become infamous over the past few years for scrutinizing paperwork ... *the level of tolerance for document mistakes varies from judge to judge ...*” (emphasis added). “Some judges chastise banks over foreclosure paperwork”, *Washington Post*, 9 November 2010.

¹³Of course, once assigned to a case calendar, the judge within that calendar who hears a case may is not necessarily random. For this reason, we use the case calendar (group of potential judges) as the unit of randomization to ensure orthogonality of the instrument to unobservable characteristics. As long as there is sufficient bias across judges to ensure that the case calendar (group) to which a case is assigned influences the outcome, then the instrument remains valid.

To this end, we match each chancery court foreclosure case record to the associated foreclosure filing record (2002–2011), which has been provided to us by Chicago-based Record Information Services, Inc. (RIS). These records allow us to observe new foreclosure filings that occur around any given delinquent homeowner’s property, which we use to study foreclosure contagion. To observe how completed foreclosures affect housing markets—prices and sales volumes—we rely on deed transfer records (1995–2008) provided to us by the Paul Milstein Center for Real Estate at the Columbia Business School. These records allow us to observe the state of the housing market around each property going through the foreclosure courts. Finally, we bolster the information about each neighborhood using data from the 2000 Decennial Census and Zillow housing price indices.

The Cook County chancery court makes public all court records, which include details on each foreclosure case. We manually collected data for each of the 217,230 chancery court cases filed between January 1 2004 and June 30 2010 from the court’s public electronic docket.¹⁴ Each record identifies the case number (a unique identifier assigned by the court), the type of case (e.g., foreclosure vs. other chancery case), the plaintiff (lending institution or mortgage servicer), the defendant, and the case calendar. The records also include every action on the case (and corresponding date), although the action descriptions are terse.

We rely on foreclosure filings from RIS to identify the location of properties going through the foreclosure process. The RIS data span all 307,209 foreclosure cases filed between 2002 and 2011 in Cook County. These records contain the same variables as the online chancery court records, except RIS does not collect the case calendar. However, RIS also collects information not included in the court records, such as the type of property (single-family, condo, etc.), details about the mortgage (type of mortgage, original loan principal, outstanding balance at time of foreclosure filing), any additional lien holders identified on the filing, and the address and latitude/longitude of the home under foreclosure. Finally, RIS also collects a record of each foreclosure auction between 2002 and 2011 (168,577 in total). This

¹⁴www.cookcountyclerkofcourt.org

allows us to conclusively observe a foreclosure outcome and associated date.

We match the chancery court records to the RIS foreclosure filings by case id.¹⁵ The resulting data set covers 174,187 foreclosure filings in Cook County filed between January 2004 and June 2010. For each record, we observe the date the case is filed, whether and when the case is dismissed or foreclosed, the location of the home under foreclosure, and the above-mentioned details of the property and mortgage. We consider a case as ending in completed foreclosure if the RIS records indicate that a foreclosure auction occurs for that property and mortgage, and we consider the auction date to be the end of the case. We consider a foreclosure case as being dismissed if it does not have an associated foreclosure auction and if the chancery court data records a dismissal action, where we take the date of the dismissal action as the relevant “dismissal date” (see the Data Appendix for details of these variable definitions). For our analysis, we drop 847 filings associated with Veterans Affairs mortgages (VA), 12,755 filings made during the Cook County foreclosure moratorium of 2009,¹⁶ and 12,365 filings made during the first or last year in which a case calendar hears foreclosure cases, as cases may be non-randomly assigned as the calendar makes the transition.¹⁷ Our results are robust to these sample restrictions.

The majority of cases end in a completed foreclosure, while a small fraction of cases are unresolved due to right-censoring. As can be seen in Table 1.1, which provides descriptive statistics (imposing the above-mentioned sample restrictions), 90,653 (61.2%) cases have an

¹⁵ See the Data Appendix for more details on the cleaning process.

¹⁶Cook County enacted a moratorium on new foreclosure filings on April 16, 2009 to last through September 1, 2009. This moratorium applied to all new filings except those in which the borrower agreed not to mount a defense prior to filing. The effect of the foreclosure moratorium can be seen in Figure 1.1: we see that the number of cases filed dips sharply during the moratorium (left axis), while the foreclosure rate jumps up (right axis). Interestingly, it seems the moratorium finished early—new filings spike before September 1, 2009.

¹⁷Case calendars have been added over time to ease the burden on existing calendars. In our data, we observe the addition of six new calendars to the foreclosure roster and the phase-out of 16 calendars (that move from hearing foreclosure cases to hearing exclusively other chancery cases). Unfortunately, the details of these phase-in and phase-out processes are not well publicized and we observe unusually low case assignment to these calendars during the phase-in periods. Our concern is that as new calendars are introduced to the foreclosure process they are restricted in the type of foreclosure cases that they hear. Indeed, including cases assigned to new case calendars in our sample brings our IV estimates closer in line with the OLS estimates, suggesting some non-random assignment to newly introduced calendars.

associated foreclosure auction, 50,140 end in dismissal (33.8%), and the remaining 7,427 foreclosure cases remain undecided due to right-censoring. The average length of a case is about 373.6 days, although this is significantly longer for cases that end in foreclosure (428.7 days vs. 274.9 for dismissals). Since the Cook County chancery court records are up to date as of the date of collection (early 2012), and the RIS foreclosure auctions are up to date through 2011, we do not observe the end of particularly long cases. This is especially true for cases filed in 2009 and 2010, from which 79.08% of the undecided cases originate. We omit these undecided cases from our analyses (as well as cases for which we observe the decision, but do not have data on our outcomes for that year).

Among dismissals, we see that only 12.0% of the borrowers “redefault”, suggesting that the dismissal outcome does not merely delay a completed foreclosure. We define redefault as a new foreclosure filing occurring against the same loan after the first case has been decided; this definition excludes future defaults to the same borrower on different loans and future defaults from different borrowers at the same property. Since dismissed cases make up our counterfactual in studying the neighborhood-level effects of completed foreclosures, this low rate of redefault is reassuring—in most instances dismissing a case does not merely delay the foreclosure (for example, while the lender finds a missing mortgage note), but provides a concrete resolution of the mortgage default within the time frame that we observe.

To study the neighborhood-level effect of completed foreclosure on housing sales and prices we rely on deed transfer records for Cook County from 1995–2008. These records cover the universe of real estate transactions and indicate the date of sale, the sale price, and the property type (residential, commercial, etc.). We restrict this data to residential real estate transactions between 2000 and 2008,¹⁸ which leaves us with 862,215 residential real estate sales. The mean sale price for these transactions is \$276,401, while the median is \$215,000. We geocode the transactions using the reported property address (using Yahoo! Placefinder), allowing us to observe transactions near properties associated with foreclosure

¹⁸See the Data Appendix for more details.

cases.

Finally, we add data from the 2000 Decennial Census matched by tract and the IRS Statistics of Income (SOI) and Zillow (matched by zip code). The Census provides us with details (as of 2000) on the population density, race, and median of the census tract in which each property is located. The IRS Statistics of Income provide a measure of zip-code-level income (mean adjusted gross income) derived from aggregated tax returns. These data are available for the 1998, 2001, 2002, and 2004–2008 tax years (for 2003, we use the mean of 2002 and 2004, while for 2009+ we use the observed adjusted gross income in 2008). Zillow provides zip-code-level housing price indices for 2000–2011.

1.4 Empirical Strategy

Our primary objective is to estimate whether and to what extent a completed home foreclosure is contagious, which we define in terms of the question: does one *completed* foreclosure cause new foreclosure *filings*? To this end, we compare the number of new foreclosure filings in neighborhoods around properties going through the foreclosure process that end in a completed foreclosure to properties that end in dismissal. An obvious concern is that there is non-random selection into completed foreclosure (versus dismissal); we deal with this endogeneity and omitted variable bias by instrumenting a completed foreclosure using the random assignment of foreclosure cases to chancery court case calendars.

For each property that goes through the foreclosure courts, we measure all outcomes annually within an x -mile radius of the property. We measure outcomes relative to the date that the case is decided (either the date of the foreclosure auction or the date of the court action in which the case is dismissed). For case i , let $d(i)$ be the time period in which the case is decided and $Y_{i,d(i)+t}$ be the outcome for property i measured within an x -mile radius of the property, t periods from the decision date. In practice, we measure time in terms of years: $d(i)$ is the year in which case i is decided, $d(i) + 1$ is the year after the case is decided, and so

on.¹⁹ In our baseline specification we use a 0.1-mile radius around each property, although our results are not sensitive to taking smaller or larger radii (of the same order of magnitude). We choose this radius both for comparability with existing literature (e.g., Campbell et al. (2011)) and to reduce the extent of observations with overlapping neighborhoods (see the discussion in the following paragraph). As an example of how we construct our outcomes, one measure of contagion we consider is the number of new foreclosure filings within a 0.1-mile radius of each property every year since the case is decided.

One consequence of using foreclosure cases (rather than mortgages) as the unit of observation is that the neighborhoods around cases may overlap. Ideally, we would use mortgages as the unit of observation and relate the probability of each mortgage defaulting to the number of neighboring foreclosures (instrumented by the expected number of neighboring foreclosures). However, in our data we only observe mortgages when they have an associated foreclosure filing; we do not observe mortgages that never enter into the foreclosure process. With our specification, cases within $2x$ miles of one another will have overlapping neighborhoods of observation. The completed foreclosure (versus dismissal) treatment will be imperfectly assigned—e.g., the neighborhood around a case that is dismissed may overlap with the neighborhood around a case that ends in completed foreclosure.²⁰ However, our instrument is still randomly assigned and will not be correlated with neighborhood overlap. Additionally, the geographic clustering of our standard errors (discussed below) will account for correlated shocks between observations with overlapping neighborhoods. Finally, we test the sensitivity of our contagion estimates to this potential overlap by restricting our sample to cases with no neighboring foreclosures or foreclosure filings within recent years of the case’s decision (we discuss this exercise in Section 1.5.1).

To achieve our goal of comparing cases filed at the same time that have different outcomes

¹⁹We have also tried months and quarters. However, since home sales and foreclosure filings in small geographic areas are low-frequency events, estimates using these finer units of time end up being low-powered and imprecise.

²⁰The extent of overlap is not trivial given the high volume of foreclosure filings in Chicago during the crisis and the possibility of foreclosure contagion. The median observation in our sample has four neighboring filings within 0.2 miles in the two years prior to the end of the case and one neighbor within 0.1 miles.

(owing to the random assignment of case calendars) we include several sets of fixed effects in our baseline specification. Filing-month fixed effects, $M_{m(i)}$, where $m(i)$ is the filing month associated with case i , allow us to compare foreclosure and dismissal among cases filed at roughly the same time (and, as explained below, we construct our instrument at the filing-month level). However, cases filed in the same month may be decided in different years. Since we do not want our estimates to be based on the comparison of cases decided in drastically different times (e.g., the onset of the financial crisis in 2008 versus the peak of the boom in 2006) we include year-of-observation fixed effects, $\psi_{d(i)+t}$.²¹ In our baseline specification, we also include property-type fixed effects, Φ_i (single-family home, condo, etc.), as cases are sorted by property type prior to randomization to case calendar. Finally, we include a vector of covariates, X_i (loan principal at origination, a dummy variable for the lender/plaintiff being a “large” plaintiff (six largest plaintiffs each representing ≥ 7000 filings), a dummy variable for the plaintiff having a “large” attorney (three largest attorneys each representing $\geq 10,000$ cases), whether the census tract has an above-median share of white residents, a set of dummy variables for the quartile of median census-tract income, and census tract population density). While these controls improve precision, our estimates are robust to excluding both the property-type fixed effects and the covariates. The resulting relationship we estimate is:

$$Y_{i,d(i)+t} = \beta_0 + \beta_1 F_i + \beta X_i + M_{m(i)} + \Phi_i + \psi_{d(i)+t} + u_{i,d(i)+t} \quad (1)$$

where F_i is an indicator for case i ending in foreclosure. Our goal is to estimate β_1 from Specification 1 separately for each value of $t \in \{0, 1, 2, 3, 4, 5\}$ for contagion and $t \in \{0, 1, 2\}$ for price and sales effects (due to data limitations).

²¹One concern is that the length of the case is itself endogenous. We have explored this in several ways: in Table 1.14 we estimate the baseline effects measuring the outcome as of the date that the foreclosure case is filed (rather than decided). While this leads to somewhat noisier estimates (the treatment is diluted by cases that have not yet been decided) the results are generally consistent with our baseline estimates. In Table 1.15 of the appendix we add controls for the length of the case and find that our contagion results hold.

We cluster our standard errors along two dimensions: filing month and census tract (Cameron et al. (2011)). Clustering on filing month captures correlation due to macroeconomic trends—cases filed in the same month may experience similar shocks. Since we also expect correlation between properties that exist in the same geographic area, we cluster at the census tract level. One issue with multi-way clustering that we occasionally encounter is invalid negative variance terms (and a non-positive-definite variance matrix). As suggested in Cameron et al. (2011), we conservatively address this by taking the maximum of the standard errors clustered only on filing month, clustered only on census tract, and clustered on filing month and census tract (and the minimum of the corresponding first-stage F-statistics).

1.4.1 Measuring Local Contagion and Prices

We define two outcomes to test for contagion. Firstly, we consider an indicator for whether any new foreclosure filing occurs within x miles of property i in year $d(i) + t$ —how does a completed foreclosure affect the probability of observing any new foreclosure filing? This outcome is of interest when there are few filings in the neighborhood and speaks to the question of whether a completed foreclosure influences the general state of nearby mortgages. However, this measure is of limited interest in neighborhoods or time periods with high filing rates—when all neighborhoods have at least one filing, a completed foreclosure will have no effect on the probability of any new filing. Secondly, we consider the count of new foreclosure filings within x miles of property i in year $d(i) + t$ —how does a completed foreclosure affect the total number of new filings? In both cases, we omit new foreclosure filings at the same address or associated with the same foreclosure case, but at a different address (e.g., a loan taken with multiple properties as collateral). We also consider the effect of a completed foreclosure on the probability of any and total number of neighboring completed foreclosures.

We also examine the effect of a completed foreclosure on local housing prices, although our estimates are hampered by sample size. While important in its own right, understanding

the pecuniary externality associated with completed foreclosure helps assess whether loss of home equity as a channel of foreclosure contagion. We take the log of the average sale price of all properties that sell within the x -mile radius of property i in the year of observation $d(i) + t$. Importantly, we omit the delinquent property itself to ensure that our price estimates are not influenced by an own-price discount of foreclosure (as found by Campbell et al. (2011)). This measure does not account for selection into sale—the types of homes that sell after a completed foreclosure may be different from the types of homes that sell after a dismissal. This selection may drive any observed price effects. Unfortunately, a hedonic approach is not possible since we do not observe property characteristics in our data.

While we cannot estimate selection into sale in terms of the types of homes that sell, we can observe whether the volume of sales itself changes. For each property going through the foreclosure process, i , we take as an outcome the count of sales, as of $d(i) + t$, that have occurred within x miles of property i since the year of decision, $d(i)$. We omit sales at the delinquent property i itself to avoid the mechanical effect of foreclosure on sales. A change in the quantity of sales in response to a completed foreclosure suggests that some sellers (or buyers) are selecting into or out of the market. At the same time, observing no response of sales volume does not prove that there is no selection.

To further explore selection into sale, we study the subset of repeat-sales (about 44% of the sample) in our data to adjust for fixed property characteristics. We estimate the quality-adjusted home value by netting out property-specific fixed effects—details are in the appendix. Our estimates of the effect of a completed foreclosure on the log of the mean quality-adjusted price will not be biased by selection into sales under two assumptions. Firstly, we assume that there is not differential occurrence of repeat sales around properties that end in completed foreclosure vs. dismissal (i.e., no selection into sample). Secondly, we assume that the property characteristics that determine sale price are not changing differentially for properties near a completed foreclosure vs. a dismissal (i.e., the error in the repeat-sales adjustment is invariant to the case outcome). Our repeat-sales estimates may

suffer from imprecision due to measurement error in the repeat-sales adjusted measure of home value.

The means of the various outcomes, as displayed in Table 1.2, suggest foreclosure contagion and a foreclosure price effect. These averages are constructed using all concluded cases (with the above-mentioned sample restrictions) observed annually for up to five years after the case decision for contagion outcomes, and up to two years after for price and sales; each observation represents a case-year. The means in the upper panel of Table 1.2 suggest foreclosure contagion: on average, there are 0.435 fewer new foreclosure filings per year around properties whose cases are dismissed than properties associated with foreclosures. There is also evidence that completed foreclosures disrupt the housing market (see the lower panel): properties that end in foreclosure see a higher volume of neighboring sales (3.099 per year relative to 2.962 near dismissed homes). At the same time, these sales occur at a lower average price—\$157,181.90 vs. \$184,212.50—although this difference is not apparent in the repeat-sales-adjusted price. While these descriptive statistics suggest negative externalities of home foreclosure, these comparisons of means suffer from omitted variable bias and endogeneity of home foreclosure.

1.4.2 Instrumental Variables Approach and First Stage Regression

There are several reasons home foreclosures may be endogenous to neighborhood-level characteristics. A completed foreclosure is not a random event—it is the product of the choice of a borrower to default on a loan, the choice of a lender to pursue a foreclosure, and the actions of the associated attorneys and judges. The borrower default decision may be influenced by local housing prices, the type of mortgage a borrower has, and the borrower’s financial position (both in terms of balance sheet, and cash flow). For example, foreclosures may be more likely to occur in neighborhoods with lower housing price levels and negative price growth (Campbell and Cocco (2011)). Similarly, the lender’s decision to pursue a foreclosure versus a loan modification depends on the home value, the probability that the borrower re-defaults

on a modified loan, the probability that the borrower brings him/herself out of delinquency without a modification, and, if the loan is serviced by a company that is not the creditor, the potential fees associated with foreclosure (Foote et al. (2008), Levitin and Twomey (2011)).

Descriptive empirical evidence suggests that observable borrower and neighborhood characteristics are correlated with home foreclosures. Table 1.1 shows means for various covariates broken down by case outcome, where the fourth column contains the p-value on the test of equality between the foreclosure and dismissal. Cases that end in foreclosure are significantly less likely to be single-family homes (59.0% vs. 68.4%), more likely to have a plaintiff that is a “large institution” (47.3% vs. 47.2%) or have a plaintiff represented by a “large attorney” (68.7% vs. 68.2%), are less likely to have a conventional fixed-rate mortgage (65.3% vs. 66.5%), and tend to be in neighborhoods with lower median income (43,748.26 vs. 46,409.49), a lower share of white residents (43.3% vs. 46.8%), and a lower population density. While studies have attempted to control for omitted variable bias using very local fixed effects analyses (see summaries in Foote et al. (2009), Towe and Lawley (2013)), ours is the first study to directly address the endogeneity of home foreclosure with a randomly assigned instrument.

We use a measure of the propensity to foreclose for each chancery court case calendar as an instrument for completed foreclosure. We construct our instrumental variable to capture the notion of judicial bias—the judges on some case calendars are more likely to foreclose than others, all else equal—by taking the “jackknife” or “leave-one-out” foreclosure rate for each case calendar, as is common in studies that use judicial random assignment as an instrumental variable (e.g., Kling (2006), Doyle (2007), Dobbie and Song (2013)). Specifically, for each case i , filed in month $m(i)$ and randomly assigned to calendar k , we take the foreclosure rate among all other cases j filed in that month and assigned to that calendar:

$$Z_i = \frac{\sum_{j \in K_{m(i)}, j \neq i} F_j}{n(K_{m(i)}) - 1} \quad (2)$$

where $K_{m(i)}$ is the set of all cases filed in month $m(i)$ and assigned to calendar k , $n(K_{m(i)})$ is the cardinality of set $K_{m(i)}$, and $F_j = 1$ if case j ends in a completed foreclosure. A case calendar with “strict” judges whose cases end often in foreclosure will have a high value of the instrument, Z_i , while a calendar with “lenient” judges will have a low value. By omitting case i when constructing the instrument, we ensure that we are not regressing the outcome of the case on itself (resulting in a mechanical correlation in the first stage). Calculating this instrument at the filing-month level accommodates changing case-calendar rosters and attitudes of judges over time.²² Failing to account for these changes may violate monotonicity of the instrument.

Our first-stage regression relates an indicator for a case ending in foreclosure to our measure of case-calendar strictness. For each case, we regress an indicator for the case ending in foreclosure (F_i) on the instrument. As with the second stage described in Specification 1, we include filing month fixed effects, $M_{m(i)}$, property-type fixed effects, Φ_i and year of observation fixed effects ($\Psi_{d(i)+t}$). The resulting first-stage regression is:

$$F_i = \alpha_0 + \alpha_1 Z_i + \alpha X_i + M_{m(i)} + \Phi_i + \Psi_{d(i)+t} + v_i \quad (3)$$

We rely on the usual instrumental variables assumptions: the instrument influences the outcome of the foreclosure case (instrument relevance), the instrument is randomly assigned (instrument exogeneity), the instrument does not itself influence neighborhood outcomes except through foreclosure of the house in question (exclusion restriction), and an increase in the instrument is associated with an increase in the probability of the case ending in foreclosure (monotonicity). We check the instrument relevance by examining the first-stage F-statistic for all of our regressions (presented in the Tables discussed in subsequent sections) and find a strong relationship between the instrument and completed foreclosure (Table 1.13 in the Appendix presents the coefficients from the baseline first-stage relationship: a one percentage-point increase in the case-calendar foreclosure rate increases the probability of a

²²For example, we see in Figure 1.1 that the foreclosure rate changes over time.

completed foreclosure by 0.556 percentage points).

If the rules of the Chancery Court are followed, then the instrument should be randomly assigned and appear independent of case characteristics. We run two sets of regressions to check the assumption that the instrument, Z_i , is exogenous. First, we regress Z_i on a set of pre-treatment covariates (controlling for property type and filing month):

$$Z_i = \gamma_0 + \gamma X_i + M_{m(i)} + \Phi_i + e_i \quad (4)$$

where Z_i is the instrument, X_i is a vector of fixed or pre-treatment property and case characteristics, and $M_{m(i)}$ and Φ_i are filing month and property type fixed effects. Random assignment (conditional on filing month and property type) implies that none of the covariates predict the value of the instrument ($H_0 : \gamma_i = 0$) and nor do they jointly determine the value of the instrument ($H_0 : \gamma_1 = \gamma_2 = \dots = \gamma_k = 0$). Second, we regress each of these covariates on a full vector of case calendar dummies:

$$X_{ji} = \rho_0 + \sum_k \rho_k \kappa_{ki} + M_{m(i)} + \Phi_i + u_i \quad (5)$$

where X_{ji} is a given pre-treatment characteristic j observed for case i , and κ_{ki} is a vector of calendar-specific dummy variables such that $\kappa_{ki} = 1$ if case i is assigned to calendar k . We then test the joint significance of these dummy variables: $H_0 : \rho_1 = \rho_2 = \dots = \rho_k = 0$.

The first column of Table 1.3 presents the coefficient estimates from Specification 4 and the p-value for the joint significance test of the covariates. We see no evidence of systematic correlation between pre-treatment covariates and the instrument, and cannot reject the hypothesis that the covariates are jointly insignificant. The second column displays the p-value for the joint significance test of case calendar dummies for Specification 5, where the outcome variable is given by the row. Again, there is no systematic relationship between case calendar assignment and pre-treatment covariates, with the exception of loan principal.²³

²³Given that we are conducting 19 significance tests in this table, we would expect to observe significance at the 5% level about once (or 0.95 times). Moreover, the relationship between loan principal and the

Importantly, we see no relationship between the instrument and total filings and total completed foreclosures (our two main contagion outcomes) in the year prior to the case being filed. We conclude that, conditional on filing month and property type, case calendars are randomly assigned.

The assumption that the instrument does not itself influence neighborhood-level outcomes is reasonable. The outcomes we are studying are the result of the decisions of those not involved in the court case (e.g., neighboring home owners). Moreover, while foreclosure cases span many months, defendants will have minimal direct contact with the presiding judges.

Finally, we find no evidence of a failure of monotonicity. The assumption maintains that a higher value of the instrument—i.e., being assigned to a stricter case calendar—weakly increases the probability of foreclosure for all cases. One can imagine a prejudiced judge who is lenient toward delinquent wealthy borrowers, for example, but push for foreclosure against defendants of lower social class. Then if there are disproportionately more of one type of borrower, a higher value of the instrument will not mean a higher probability of foreclosure for all cases. We explore this possibility by relating group-level foreclosure rates (e.g., foreclosure rate among cases in predominantly white vs. non-white census tracts) within each case calendar to the overall foreclosure rate for each calendar and find that foreclosure rates for sub-groups are all increasing with the overall case calendar foreclosure rate. A discussion of these results can be found in the appendix.²⁴

instrument is economically small: a one percentage-point increase in the case-calendar-specific foreclosure rate is associated with a drop of 0.0045 standard deviations in the instrument.

²⁴As suggested by Mueller-Smith (2013), we have also estimated our baseline specification by constructing the instrument separately for various sub-groups. If monotonicity is violated, then these results may differ substantially. While we do not see a substantial difference in our baseline results (see Table 1.16 in the Appendix), these “monotonicity-robust” estimates are imprecise; as splitting the data into filing-month/characteristics cells often yields few observations per cell.

1.4.3 Interpretation of the Two-Stage Least Squares Estimate

Our estimate captures the local average treatment effect (LATE) for foreclosure cases in which judges are influential, compounds the effect of all subsequent completed foreclosures caused by the initial foreclosure, and is representative of neighborhoods with many foreclosure filings. The estimate does not represent the effect of a completed foreclosure relative to a mortgage that is in good standing; rather, the estimate represents the effect of a completed foreclosure relative to the effect of a foreclosure case being dismissed. We argue that this parameter is of interest to policy makers.

Firstly, as discussed in Doyle (2007), if there are heterogeneous treatment effects the parameter identified by a judicial random assignment instrumental variable (or in Doyle’s case, rotationally assigned case workers) is the LATE for “marginal” cases—those where the judge is likely to have an influence. Intuitively, there are cases that will always end in foreclosure and cases that will always end in dismissal; the set of “compliers” with our instrument are the marginal cases where the judges on the case calendar have influence on the outcome.²⁵ We find that the characteristics of the sub-population of loans that comply with our instrument are consistent with cases on the margin of foreclosure or dismissal, representing individuals who have a higher ability to pay than the typical delinquent borrower, but are facing difficult circumstances that could be mitigated through loan modification. We stratify the sample along several margins: tract-level quartile of income (from the 2000 Decennial Census), whether the loan is from a “large” lender, whether the mortgage is conventional, whether the zip code experiences positive price growth in the year that the case is filed, and a proxy for whether the property is worth less than the loan (“underwater”).²⁶ Our goal is to proxy characteristics of borrowers who are likely to benefit from loan modification; creditors

²⁵We can conceive of situations where judges will not matter. Sophisticated borrowers may always be able to renegotiate the terms of their mortgages (and a dismissal of the case), regardless of who the judge is. Other borrowers may resign themselves to walking away from their home and mortgage and choose not to appear in court at all.

²⁶We define a proxy for a borrower being underwater as whether or not the outstanding debt at filing is larger than the initial loan principal.

may be more willing to modify in such situations (Adelino et al. (2009)), making them more responsive to judicial input. For each sub-sample, G , we estimate the first-stage:

$$F_i = \alpha_0 + \alpha_{1G}Z_i + \alpha X_i + M_{m(i)} + \Phi_i + \Psi_{d(i)+t} + v_i \quad \forall i \in G \quad (6)$$

We then take the ratio of the estimate of the first-stage relationship for group G to the estimate for the full sample from Specification 3: $\frac{\alpha_{1G}}{\alpha_1}$. As described by Angrist and Pischke (2008), the ratio of the sub-group-specific first stage to the full-sample first stage represents the relative likelihood that a complier belongs to the given subgroup.

We interpret our estimates of these ratios, presented in Table 1.4, as demonstrating that compliant cases are likely to be on the margin of completed foreclosure. The upper panel shows that compliers are more likely to be in the upper two quartiles of income than the general population of foreclosure cases. Taking income as a measure of a borrower’s ability to repay their loan, these estimates suggest that compliers are more likely than the typical borrower to be able to resume payments if the case is dismissed.

At the same time, the compliant sub-population may benefit from a mortgage modification. Compliant borrowers are less likely to be in a zip code with positive price growth. Falling house prices may be largely responsible for the default crisis (Mayer et al. (2009)): borrowers may be in default because they expect to lose money on their mortgages as housing prices fall and the value of the asset drops below the cost of the debt. A modification reducing the loan principal or the interest rate may reduce the anticipated loss, making default and foreclosure less appealing. At the same time, compliant borrowers are not in dire straits—they are less likely to be underwater on their loans, so a modification may be more effective (home value is not so low that the mortgage is a lost cause) and may result in smaller losses to lenders than modification of more severely underwater loans. Additionally, compliers are less likely to have conventional loans. There is some suspicion that unconventional mortgages are responsible for many defaults during the crisis. For example, borrowers with low “teaser” interest rates or balloon payments may have been expecting to refinance

their loans to avoid higher monthly payments, but found themselves without this option during the financial crisis. In such cases, a modification may be particularly effective (by mimicking the effect of a refinance).²⁷ Finally, it is interesting to note that the differential characteristics of the compliant population appear borrower specific—compliers are no more or less likely to have a loan from a “large” lender.

Secondly, our LATE estimate does not simply identify the effect of a single completed foreclosure, but compounds the effects of all subsequent induced foreclosures. If foreclosures are contagious, then a completed foreclosure will lead to subsequent foreclosure filings. In turn, some of these filings will become completed foreclosures and themselves cause new filings. Since our empirical strategy compares the neighborhoods around cases in the foreclosure courts each year after the case is decided and does not control for the effects of subsequent foreclosures, our estimates will compound the effects of these subsequent foreclosures. We believe that this parameter is relevant to policy makers since it represents broad consequences of the marginal foreclosure, although is imperfect, as the 0.1 mile radius will only partially capture the effect of successive foreclosures.

Similarly, if there is foreclosure contagion, our LATE represents neighborhoods with several completed foreclosures. There is selection into our sample: we only observe neighborhoods around properties going through the foreclosure courts. If completed foreclosures induce subsequent filings resulting in additional observations in our data, neighborhoods with previous completed foreclosures will be over-represented in our sample. This does not affect the validity of the instrument—case calendars are still randomly assigned—but influences the interpretation of the LATE. Nonetheless, we find that our contagion estimates persist when we restrict the sample to properties with no foreclosures in the past two years within 0.1 miles.

Finally, our estimates are conditional upon a foreclosure filing having occurred in the neighborhood. Our empirical strategy and data set necessarily rely on comparing neigh-

²⁷ At the same time, there is debate about the importance of unconventional loans in the default decision (c.f., Mayer et al. (2009)), so this channel may be less relevant.

borhoods around properties that are already going through the foreclosure process. Our estimates will not account for any externalities associated with a borrower default or a foreclosure filing. Many have argued that it is a completed foreclosure and subsequent real-estate ownership of the associated property that drives foreclosure-related externalities. While we cannot speak to any spillovers from borrower default, our estimates provide a well-identified answer to whether there are negative spillovers associated with the completed foreclosure itself.

The LATE represented by our estimates is a relevant parameter for the policy question of how best to address the problems of delinquent borrowers. Policymakers concerned with foreclosures can focus on several stages of the lending process: how easy it is to originate/obtain mortgages, how to prevent borrowers from defaulting, and what to do once a borrower has defaulted. Our parameter, which is estimated conditional on foreclosure filing, focuses directly on the latter question.²⁸ Moreover, the LATE is relevant for cases on the margin of foreclosure and dismissal, and who are influenced by foreclosure court judges. These cases are also likely to be influenced by policies discouraging foreclosure on delinquent loans.

1.5 Neighborhood-Level Effects of Completed Foreclosure

We find robust evidence of foreclosure contagion that persists over several years. Neighborhoods around a completed foreclosure are 10% more likely to have at least one foreclosure filing in a given year relative to neighborhoods around a dismissed property and experience around 0.5 to 0.7 more total filings per year. We also find that residential properties that transact around completed foreclosures do so at a price discount (on the order of 30–40%), although this effect may be largely explained by negative selection into sale.

²⁸Of course, the usual partial-equilibrium caveat applies: any change to foreclosure policy may affect ex-ante incentives (e.g., Mayer et al. (2011)) and housing market outcomes (e.g., Pence 2006), which are not captured in our reduced-form estimates.

1.5.1 Contagion in Foreclosure Filings

Our estimates demonstrate that completed foreclosures are contagious. Table 1.5 presents our baseline 2SLS estimates of the effect of a completed foreclosure on the probability of observing any neighboring foreclosure filing in a year and on the annual count of neighboring foreclosure filings within 0.1 miles of the at-risk property. The 2SLS estimates show that a completed foreclosure increases the probability of observing any new filing within 0.1 miles by 0.052 percentage points in the year of the decision (a 7.4% increase in the mean for all dismissed cases). This effect increases over time to 8.2 percentage points (11.7%) in the second year after the decision, 9.0 percentage points (12.8%) in the third, and 24.7 percentage points (35%) in the fourth year out. Similarly, the 2SLS estimates show that a completed foreclosure causes 0.54 to 0.70 new foreclosure filings per year in the year the case is decided and the following three years. This contagion represents a 25–32% increase in total annual filings relative to an average of 2.161 filings per year around dismissed properties. Note that the instrument is strong in the year of the decision through the second year after the decision (F-stats around 200), although is relatively weak three, four and five years out owing to the smaller sample for these periods.

Our contagion estimates are generally not sensitive to the specification, sample, or geographic measurement of the outcome. The results are robust to excluding the covariates, omitting the property fixed effects, dropping cases decided in the summer months (the court automatically dismisses inactive cases during this time), using a monotonicity-robust instrumental variable, using the full sample and including the foreclosure moratorium, omitting each filing year one by one, and dropping neighborhood-years with foreclosure filings above the 99th percentile (see Tables 1.16, 1.17, 1.18, and 1.19 in the Appendix). We also estimate our baseline results measuring outcomes within 0.25 miles of the delinquent property and find that contagion (and price effects, discussed below) persist—See Table 1.20.²⁹ Finally, we

²⁹The estimates for price generally decline as we increase the radius. We find minimal effect for any new foreclosure filing—this is due to the fact most cases have at least one new filing per year within this expanded radius (and so there is little variation in the outcome). On the other hand, the effects for total filings tend

confirm that our estimates are driven by dismissals where the defendant retains possession of the property (rather than deeds-in-lieu of foreclosure or short sales). We estimate our baseline results on the sample of cases in which the plaintiff identifies that there are additional liens against the property. Although less precise, the point-estimates for this sample are comparable to the full-sample estimates and are not significantly different (see Table 1.21 in the Appendix).

We further explore the validity of our estimates by applying the same 2SLS procedure to our contagion outcomes measured in the three years prior to the case being filed. If our instrumental variable is truly randomly assigned, we should not expect to see any effect of a case ending in foreclosure before the case has even started. We present these “pre-treatment” estimates in Table 1.7. Reassuringly, when instrumented by case calendar leniency, a case ending in foreclosure appears to have no relationship to local housing prices prior to the start of the case—the point estimates are close to zero and generally insignificant.

We examine the cumulative effect of a completed foreclosure in order to appreciate the full extent of contagion. Rather than using as an outcome the number of new foreclosure filings per year for each year since the decision, we instead consider the total number of new filings since the decision. These estimates are presented in the second panel of Table 1.5 and show that a completed foreclosure leads to a significant divergence in foreclosure filings relative to a dismissal. As noted above, in the year of the decision a completed foreclosure causes 0.691 new filings. However, neighborhoods around completed foreclosures have experienced 2.09 more foreclosure filings by the second year after the decision, and 6.45 more filings by the fourth year after the case ends. One completed foreclosure may have a substantial impact on the composition of a neighborhood, at least in the short and medium term.

to be larger. There are two explanations for this. First, as the radius expands, the compounding effect of neighboring completed foreclosures (induced by the initial observation) grows—the larger radius allows for capturing more of the neighborhood around neighboring foreclosures. Second, as the radius grows the total base number of properties that might file for foreclosure grows. Increasing the radius 2.5 times increases the area of the neighborhood 6.25 times, thus dramatically increasing the potential number of properties that may be affected by the completed foreclosure. However, the estimates themselves only grow by a factor between about 2 and 4, suggesting that the effect declines with distance.

One concern with our findings is that they are specific to neighborhoods that are experiencing a wave of foreclosures. Firstly, our period of study (2004–2011) is largely made up of the housing crisis. Secondly, since we only observe neighborhoods where a foreclosure filing has occurred, and since we do find that foreclosures are contagious, there is likely selection into our sample—foreclosure filings (and, thus, observations in our data) are likely to be in neighborhoods with recent completed foreclosures. From a policy perspective, it is especially important to understand the cumulative impact of the first foreclosure in a neighborhood.

A completed foreclosure is contagious even in a neighborhood that has not experienced a foreclosure in recent years. We restrict our sample to cases where there have been no completed foreclosures within 0.1 miles in the two years prior to the decision (the results are similar if we restrict to cases with no filings within two years) and estimate the cumulative contagion effect of a completed foreclosure, presented at the bottom of the second panel of Table 1.5. We find clear evidence that completed foreclosures are contagious even in neighborhoods with no other recent completed foreclosures, although these results are less precise than when we use the full sample owing to a smaller sample size: a completed foreclosure leads to 1.3 more filings by the end of the first year after the decision, and almost four more filings by the third year out. Even the first completed foreclosure in a neighborhood has externalities. Moreover, these results suggest that the contagion we observe is not an artifact of selection into the sample (cases induced into the sample because of neighboring prior completed foreclosures).

Similarly, these results suggest that overlapping treatment and control neighborhoods (i.e., properties around a case that ends in dismissal may also be within the neighborhood of a case that ends in foreclosure) is not a serious problem with our baseline estimates. No new filings within 0.1 miles in the past two years reduces the potential for such overlapping neighborhoods. That contagion persists in these neighborhoods is reassuring. Of course, this is an imperfect test: cases that have no new filings within 0.1 miles in the past two years have a median of 1 new filing per year within 0.2 miles (whose 0.1 mile radius neighborhood

will overlap). However, the extent of overlap for such cases is not too large. For example, the overlapping region of a case that is 0.15 miles away represents 14.4% of the total neighborhood area. When we restrict the sample to cases with no new filings within 0.2 miles in the past two years, thus eliminating all potential for overlap, we suffer from small sample size. The (unreported) point estimates still suggest foreclosure contagion (around 0.6 new filings per year in the first three years), but the estimates are very imprecise.³⁰

Finally, we find that contagion is strongest during the peak of the housing bubble and beginning of the crash, and disappears at the height of the foreclosure crisis. We estimate the baseline specification restricting the data to a constant sample of cases for which we observe a full three years after the decision.³¹ We present estimates for the constant sample (cases decided in 2004 through 2008 and observed for three years after the decision) in Table 1.6. These estimates show strong evidence of contagion—a completed foreclosure increases the probability of any neighboring filing by between 6 and 17 percentage points. Similarly, a completed foreclosure during this period increases the total number of new filings by 1.2 in the year of the decision, increasing to 1.6 the year after, and dropping down to 1.1 and 0.7 new filings in the subsequent two years. Interestingly, contagion is stronger for this sample than for the baseline sample. In the bottom panel of Table 1.6 we present estimates of contagion for the complementary sample—cases decided in 2009 through 2011, which we observe for two or fewer years after the decision. There is little evidence of contagion during this time period; the point estimates are very small (-0.018 to 0.027 for any filing, 0.099 to 0.239 for new filings). It is not surprising that during this period of heightened foreclosure

³⁰An additional concern relates to the length of cases—as seen in Table 1.1, cases ending in dismissal are significantly shorter than cases that end in completed foreclosure. A possible explanation is that foreclosure externalities are driven by borrower behavior while in default, and the effect is larger for cases ending in completed foreclosure since these cases are longer. To rule out this explanation, we estimate our baseline 2SLS estimates, adding flexible controls for the length of the case. We try three different sets of controls—log of the number of months, a quadratic in number of months, and dummy variables for the number of quarters of length—and present these results in Table 1.15 in the Appendix. These estimates show contagion effects that are comparable to our baseline estimates, although the addition of these length-of-case controls reduces the precision of the estimates.

³¹This also addresses the issue of interpretation with the baseline estimates in the top panel of Table 1.5 that the sample changes each year (e.g., for cases decided in 2010, we only observe the year of decision and one year out).

activity we find no effect on any new filing per year—if most neighborhoods are already experiencing a foreclosure filing (regardless of neighboring foreclosure cases), then there will be little movement in this outcome. However, the fact that we find no effect of a completed foreclosure on the total number of new filings suggests that foreclosure contagion is not a strong force during this time period.

We interpret this finding of differential contagion by time period as evidence that foreclosure contagion acts on marginal borrowers—those on the threshold of being able to stay in their homes. In particular, at the peak of the housing bubble and beginning of the bust, it may be that a completed foreclosure sends a stronger signal to mortgage holders at risk of default—for example, a signal about the future of the neighborhood and local property values or conveys information about the foreclosure process. Conversely, during the height of the crisis (2009–2011) it may be that mortgage holders are already well informed about the state of their mortgages and the foreclosure process itself. Indeed, neighborhoods in the latter period experienced 16% more filings in the year the case is decided than in the earlier period. Similarly, in Section 1.8.4 of the Appendix, we extend our empirical method to relate counts of new filings to lagged counts of completed foreclosures for small neighborhoods. We find a smaller, although non-zero, contagion effect in the latter period, and also find evidence that the marginal contagion effect of a completed foreclosure diminishes as a neighborhood experiences more foreclosures.

1.5.2 Contagion in Completed Foreclosures

To better understand the costs of foreclosure contagion, we look for contagion in completed foreclosures. Above, we established contagion in foreclosure filings—the result of new borrower defaults and lenders pursuing foreclosure action (we take up the discussion of these two actions in more depth in Section 1.6). However, if we do not see contagion in completed foreclosures, then contagion in filings is unlikely to be a large contributor to the spread of a foreclosure crisis. Moreover, the costs of new filings that end in dismissal are,

perhaps, smaller than the costs of new completed foreclosures (for example, owing to pecuniary externalities of completed foreclosure, moving costs associated with the displacement of homeowners, etc.).³²

We find contagion in completed foreclosures. We estimate the baseline contagion IV regressions replacing the outcomes with an indicator for any neighboring completed foreclosure (within 0.1 miles of the property in the given year since the case is decided) and the count of completed foreclosures. We present these estimates in Table 1.8 and find that a completed foreclosure moderately increases the probability of observing any neighboring completed foreclosure (by 13.8 percentage points three-years out). Moreover, there is a notable increase in the number of neighboring completed foreclosures: one completed foreclosure causes between 0.28 and 0.56 additional completed foreclosures annually (or between 40 and 93 percent off of the mean). Thus, contagion appears to play an important role in the spread of foreclosures; mitigating completed foreclosures may reduce the depth and costs of a housing crisis.

The timing of contagion in completed foreclosures suggests that borrowers and/or lenders who are already involved in the foreclosure process respond to nearby events. Given how long the foreclosure process takes—from default to filing to completed court case—it may seem strange that we find contagion in completed foreclosures in the year of the decision. However, borrowers and lenders may respond at any stage of the foreclosure process. For example, a neighboring completed foreclosure may influence the effort a borrower puts into fighting an ongoing foreclosure case by conveying information about the costs of fighting foreclosure or the probability of a successful loan modification.³³

³²One difficulty in studying contagion in completed foreclosures is that the response may be driven by judges. While this is not an issue when studying contagion in foreclosure filings—an event that depends only on the actions of the borrower and lender—judges have an influence over the outcome of a foreclosure case. We cannot explicitly rule out judge behavior as driving contagion in completed foreclosures. However, we do not expect judge contagion to be a dominant force—this would require judges to be well informed about recent events in the neighborhoods around the delinquent properties associated with their cases, which we find unlikely given the volume of cases and judicial random assignment (judges do not specialize in neighborhoods).

³³Of course, this need not be a question of effort—given a binary choice (with no effort necessary) between keeping their home and walking away, delinquent borrowers may be swayed toward the latter by neighboring completed foreclosures (e.g., the neighboring event reduces the borrower’s perception of the value of the neighborhood).

Indeed, we find that completed foreclosures influence cases that are ongoing. For each case, i , we split neighboring completed foreclosures into two groups: completed foreclosures among cases filed before case i is decided and cases filed after case i is decided. By focusing on the former group, we can observe how the outcome of case i influences ongoing foreclosure cases. We define as an outcome the count of completed foreclosures among cases filed before the decision (and an indicator for any completed foreclosure in this group as another outcome) and estimate our contagion model—the results are presented in the lower panel of Table 1.8.³⁴ The estimates show that a completed foreclosure causes 0.11 to 0.60 new completed foreclosures per year among cases that were filed before the decision in the first two years after the decision (these cases drive the contagion in the first two years, during which there is minimal contagion among newly filed cases). These results provide evidence that contagion acts not only through influencing the behavior of borrowers and lenders before filing (e.g., encouraging borrower default), but also by changing how borrowers/lenders approach an ongoing foreclosure case.

1.5.3 Housing Markets

Our baseline 2SLS estimates, presented in the first panel of Table 1.9, suggest that a completed foreclosure lowers the average neighboring sale price over several years. The columns of Table 1.9 present the baseline price and sales effects for the year in which the case is decided, and one and two years after. The estimates suggest that a completed foreclosure depresses neighboring residential sale prices by 12.7% in the year of the foreclosure, 41.1% in

³⁴We present similar estimates for the complementary sample—completed foreclosures among cases filed after case i is decided—in Table 1.22 of the Appendix. These estimates (in the upper panel) show a *negative* effect of a completed foreclosure on neighboring completed foreclosures for the year of decision and the year after, with a positive effect by the 2nd year out. However, the negative effect is not a result of more cases being dismissed—estimates in the lower panel show no (or sometimes negative) effect of a completed foreclosure on neighboring dismissals for this sample. Thus, it appears that the negative estimate for completed foreclosures is a result of cases taking longer to be decided after a neighboring completed foreclosure, perhaps because those cases that are induced into filing by the neighboring foreclosure are different in some sense (e.g., more complicated, borrower puts up a stronger defense, etc.).

the year after and 35.8% two years out.³⁵ However, the precision of these estimates suffers from a smaller sample size than our contagion results (we only observe housing sales through June 2008 and only observe prices when a home sells): they are only significant in the year after the decision and the first-stage F-statistics are around 20. Nonetheless, these estimates are not sensitive to the same robustness checks described in Section 1.5.1 (see Tables 1.17, 1.23, and 1.24 in the appendix). Similarly, we see no pre-filing relationship between local housing prices and the eventual outcomes of the cases (see Table 1.7).

Our 2SLS estimates show suggestive evidence that a completed foreclosure influences the volume of residential housing transactions in a neighborhood. The point estimates of the effect of a completed foreclosure on the cumulative number of neighboring residential sales since the case decision shows a large increasing trend—while this effect is not statistically significant, it appears as though a completed foreclosure may induce additional home sales. Keep in mind that we omit sales of the delinquent home itself from the count of sales—this is not a mechanical increase in sales due to foreclosure auctions and REO sales. This increase in number of sales raises the question of whether the drop in sale price after a completed foreclosure is caused by selection into sale of lower quality (and thus, lower price) homes or a drop in the value of neighboring properties (conditional on quality). In particular, if this drive to sell in response to a completed foreclosure is stronger among those with lower-quality houses (for example, because they are in a more precarious financial situation), then we would expect a decline in average neighboring sale price.

Our repeat-sales-adjusted price results suggest that there is negative selection into residential sales after a completed foreclosure, which may explain much of the negative sale-price effect associated with completed foreclosure. As explained in Section 3.3, we use a repeat-sales methodology to adjust reported sale prices for property quality. We then estimate the effect of a completed foreclosure on the log of the average neighboring quality-adjusted sale

³⁵That the effect becomes more pronounced one year out is consistent with the theory that foreclosures lower neighboring prices through a disamenity effect, and with existing studies that find that the housing-price effects of foreclosure are driven by the supply-effect of the property being listed, keeping in mind that banks generally do not list foreclosed properties on the market immediately.

price. The point estimates of the effect of home foreclosure on the log of the mean repeat-sales adjusted price tend to be small: 0.059 in the year of the decision, 0.003 the following year, and are not significantly different from zero; controlling for property quality yields smaller price effects of completed foreclosure.³⁶ By pooling the unadjusted (i.e., baseline) and repeat-sales-adjusted price regressions (allowing all fixed effects to vary by group) and testing the cross-equation restriction that the effects of completed foreclosure are different (i.e., $H_0 : \beta_{\text{raw prices}} = \beta_{\text{repeat sales}}$), we find that the treatment effects for adjusted prices are significantly smaller (in absolute value) than for the unadjusted prices in the year after the decision (p-values for this test are presented in the adjusted-price panel of Table 1.9). We examine the relationship between completed foreclosure and quality more explicitly by using as an outcome the log of the mean price at previous sale (adjusted for year-of-sale effects) for all neighboring repeat sales observed in the given year. If completed foreclosures induce more low-quality properties to sell, we would expect these properties to have sold at a lower price in the past. We find a large negative, although insignificant, effect of completed foreclosure on the price at previous sale (see the bottom of Table 1.9).

We conclude that there is negative selection into sale—when a neighboring foreclosure occurs, the properties that do sell tend to be of a lower quality. The difference between our repeat-sale adjusted and unadjusted baseline sale price estimates, along with the earlier estimates that suggest an increase in sales volume after a completed foreclosure, suggest that lower-quality homes are more likely to transact after a completed foreclosure. This may be the case if owners of lower quality homes in a neighborhood are those who have lower income or wealth. Given a signal that the neighborhood is declining (i.e., a neighboring completed foreclosure), these owners may be eager to sell before the neighborhood “falls apart” in order to avoid the liquidity shock, both from the difficulty of selling an underwater property and

³⁶In Table 1.25 in the appendix, we confirm that the difference between the baseline and repeat-sales adjusted estimates is not driven by a selected sample of repeat sales. We estimate the baseline price effects for two measures of neighboring unadjusted price: mean log sale price for all non-repeat sales, and mean log (unadjusted) sale price for all repeat sales. In both cases, the point estimates are comparable to the baseline price effects—large and negative—although the smaller sample size reduces precision.

the inability to borrow against an underwater home. Of course, we cannot conclude that all foreclosure-related pecuniary externalities are driven by this selection into sale. Existing studies of foreclosure-price externalities (e.g., Campbell et al. (2011)) find estimates one order of magnitude smaller than our own, and, although our quality-adjusted point estimates are generally zero or positive, given the imprecision of our estimates, we cannot rule out that there are negative effects of completed foreclosure on the quality-adjusted value of neighboring properties.

1.6 Evidence of Contagion Mechanisms

While our reduced-form estimates provide clear evidence of contagion, interpreting the causes of contagion is difficult since a foreclosure filing is the result of a joint decision of the borrower, who defaults on his/her loan, and a lender or servicer that chooses to pursue foreclosure on the borrower's home. Both parties may be influenced by a neighboring completed foreclosure, which may lower the value of the asset and alter the incentives for the borrower to repay and the lender to file for foreclosure. We study heterogeneity in the treatment effect of foreclosures—across loan types, lender identity, and market conditions—to gain insight into the importance of borrower and lender behavior.

1.6.1 Distinguishing Borrower and Lender Response

Our goal is to distinguish whether contagion is driven by borrowers or by lenders. We argue that lenders would exhibit no response to a local completed foreclosure or even anti-contagion—fewer foreclosure filings in response to a completed foreclosure. This provides a weak test for the presence of borrower-driven contagion. We then explicitly test for borrower-driven contagion by examining contagion among loans from lenders known for automation of foreclosure proceedings (who are, thus, unlikely to respond to very local market conditions).

Lenders and mortgage servicers—third parties who are paid by creditors to collect mortgage payments and manage defaults—have three options when dealing with a delinquent

borrower: do nothing, pursue foreclosure, or renegotiate the loan. Through foreclosure, lenders acquire the mortgaged property and gain the benefit of the value of the property (servicers collect fees for managing the foreclosure). Lenders or servicers may instead modify the terms of a mortgage with the goal of making the payments more affordable—e.g., by reducing the principal or interest rate on a loan. The benefit of modification is ensuring that loan payments continue, at the cost of a lower lifetime value of repayment. Finally, when faced with a borrower who is not making payments, a lender or servicer can always do nothing, although this is generally not optimal.³⁷ Existing empirical studies have found that mortgage modification is uncommon (Adelino et al. (2009), Ding et al. (2009)). There are several (non-exclusive) explanations for why lenders and servicers prefer foreclosure to modification: borrowers may redefault on a modified loan (Adelino et al. (2009)); frequent modification encourages strategic default among borrowers who can pay, but would benefit from modification (Foote et al. (2008)); foreclosing and purchasing the property at auction allows creditors to delay recognizing a loss on their balance sheets; modification requires coordination and agreement among all creditors (e.g., multiple holders of a mortgage-backed security or a second lien on the property), some of whom may need to accept a loss (Gelpert and Levitin (2009)).

These incentives to file for foreclosure against delinquent borrowers are unlikely to respond to neighboring completed foreclosures (unless borrowers themselves change their behavior), making lender-driven contagion unlikely. In the Appendix, we adapt the simple framework of Foote et al. (2008) and Adelino et al. (2009) to explore how the lender’s incentive to modify or foreclose might change with a neighboring completed foreclosure. Intuitively, suppose that a completed foreclosure does not influence the behavior of neighboring borrowers (i.e., no borrower-driven contagion), but lowers neighboring home values.³⁸ A lender would be

³⁷Inaction may be a good strategy if renegotiation and foreclosure are costly to the lender and if there is a high probability that the default will “self cure”—i.e., the borrower will resume making payments. Additionally, accounting rules adopted in April 2009 allow creditors to keep a delinquent loan at face value on balance sheets if there is a reasonable chance that the loan may be repaid.

³⁸If a completed foreclosure has no influence on neighboring home prices and if we maintain the assumption that borrowers are unresponsive to a completed foreclosure, then lenders and servicers should not respond—

weakly less inclined to foreclose since the collateral is worth less relative to the value of the modified loan, although the lender may still foreclose for the reasons discussed in the preceding paragraph. Moreover, mortgage servicers are not likely to respond to a neighboring foreclosure at all. Servicer compensation in foreclosure depends on fees incurred during the foreclosure process and not on the value of the home itself—servicers are indifferent to the value of the collateral. On the other hand, since servicer compensation on a loan in good standing is typically a function of the monthly payments collected, modifying a mortgage entails lower future payments to the servicer and servicers do not typically recover fees for modifying a loan. Thus, when faced with a delinquent loan, pursuing foreclosure is generally more valuable to servicers, and the decision to take this action is orthogonal to the value of the property. This is consistent with the discussion of Levitin and Twomey (2011) who note that servicers have no stake in the value of the property under consideration, and will take the action that maximizes the fees they collect, which is typically foreclosure.³⁹ In summary, under the assumption that borrowers do not respond to completed foreclosures, a neighboring completed home foreclosure should (weakly) discourage foreclosure filings—the collateral is less valuable to the lender and servicers are indifferent to home value.

The above-outlined servicer and lender incentives provide a simple test for the hypothesis of no **borrower-driven** contagion, which we reject. Specifically, this null hypothesis can be rejected if we observe positive contagion (since lenders would display weak anti-contagion). That we observe contagion in Table 1.5 suggests that borrowers respond. Note, that while this confirms the presence of borrower contagion, this does not rule out lender response—when borrower default probabilities are changing, lender response is ambiguous.

To further establish that contagion is driven primarily by borrowers, we study contagion

the conditions of the neighboring loans are unchanged.

³⁹Levitin and Twomey (2011) also point out that, although foreclosure is preferable to modification from the servicer’s perspective, there is incentive to delay the foreclosure process. Servicers must forward missed mortgage payments to the creditor and are repaid when the borrower resumes payment or the property is sold at foreclosure auction. However, servicers charge late fees that are paid when payments resume or the foreclosure occurs. Levitin and Twomey (2011) show that servicers can benefit by waiting several months to accrue late fees before beginning the foreclosure process. Since this wait period is unlikely to be influenced by a neighboring completed foreclosure, it does not relate to contagion.

among loans held by lenders and servicers who are known to have automated foreclosure filing processes. We identify foreclosure filings by all lenders/servicers investigated in the Independent Foreclosure Review Settlement conducted by the Office of the Comptroller of the Currency and the Board of Governors of the Federal Reserve System.⁴⁰ We then redefine our contagion outcomes in terms of new foreclosure filings among automating lenders only and re-estimate our baseline 2SLS contagion models. The new outcomes are: the count of new foreclosure filings among automating lenders and an indicator for any new foreclosure filing from these lenders. Because this restriction reduces our power—we are studying filings among a subset of lenders, which are lower frequency events—we pool all post-decision years ($d(i) + t$ where $t \in \{0, 1, 2, 3, 4\}$) and restrict the contagion effect to be constant across all years (see Table 1.26 in the Appendix for annual estimates). Clustering our standard errors on the census-tract dimension accommodates correlation between the same property across different years of observation. If contagion is driven largely through lender response, then we should find no effect of a completed home foreclosure on the number of foreclosure filings among lenders likely to be automating foreclosures. At the same time, if contagion is driven primarily through borrower response, then our estimates should be comparable to our baseline full-sample contagion estimates.

We find that contagion is primarily borrower-driven—contagion in foreclosure filings persists when studying loans among automating lenders. Table 1.10 presents the pooled estimates for all filings (upper panel) and for filings among automating lenders. One concern is that the total number of loans in a neighborhood from automating lenders is necessarily smaller than the total among all lenders (as in the baseline estimates). If new filings are proportional to the existing number of loans, then we would expect a larger response in the

⁴⁰Throughout the mid- and late-2000s, many lenders and servicers adopted automated foreclosure filing procedures (Levitin and Twomey (2011)). In many cases the delinquent borrower’s situation and background were not given close consideration. As investigations have revealed, employees of several large mortgage servicers and financial institutions falsely testified that they had personally inspected delinquent borrowers’ information, even though the processing speeds made it impossible for this to be true (this was the so-called “robosigning” controversy, settled for \$25bn in April 2012 between the federal government, 49 state attorneys general and the five largest servicers). See Kiel, P. “The Great American Foreclosure Story: The Struggle for Justice and a Place to Call Home”, ProPublica, 10 April 2012.

total count of new filings for the baseline than for automating lenders. To address this, we focus on the estimates of the log of total filings in the third column. We observe similar contagion effects for the two outcomes: an increase in foreclosure filings of 9.4% per year in the baseline sample versus an increase of 8.4% among filings from automating lenders.⁴¹

We also observe similar contagion in the probability of observing any new filing (0.062 versus 0.039). Thus, given the incentives of lenders and servicers and the finding that contagion persists among automating lenders, we conclude that contagion is mainly borrower driven.

1.6.2 Foreclosure Contagion and Negative Equity

We investigate whether foreclosure contagion is driven by mortgages held by individuals who face high debt relative to the value of their property. Mortgage default theory (e.g., Campbell and Cocco (2011), Deng et al. (2000)) suggests that borrowers will only default when the value of their home falls below the balance of their mortgage, putting the borrower “underwater.” Intuitively, if the market value of a home is greater than the outstanding debt, a homeowner who is having difficulty making mortgage payments may sell the property and use the proceeds to pay off the debt. Conversely, underwater borrowers who are having difficulty making payments do not have the option to sell. To the extent that foreclosures lower or send a signal about neighboring home values, foreclosure contagion may operate through pushing borrowers (further) underwater (Campbell (2013)).⁴² We use a proxy for equity—relating the size of the initial loan to the outstanding debt at foreclosure filing—to

⁴¹Similarly, while the point estimates for total filings are different, 0.74 for the baseline and 0.115 for the automating lenders, relative to the mean number of filings for each group (2.44 filings per year within a 0.1 mile radius for all filings, 0.37 for automating lenders), these estimates are remarkably similar (30.3% versus a 29.7%).

⁴²The benefit of foreclosure is muted when borrowers have more than one mortgage taken against the property (e.g., a mortgage and a home equity loan). In Illinois, mortgages are “recourse” debt—if the foreclosure auction generates less revenue than the value of the outstanding debt, the borrower still owes the creditor the balance. In the majority of cases the lender purchases the property in the amount of the outstanding debt to avoid writing down a loss (i.e., the property becomes REO), in which case borrowers no longer owe. However, a borrower is still on the hook for any additional liens on the property. Thus, a foreclosure will not always render borrowers debt free.

determine whether contagion is driven by individuals who are likely to be underwater on their loans. Understanding which borrowers are most influenced by a neighboring completed foreclosure sheds light on the mechanism of contagion and the borrower default decision more generally.

We examine whether foreclosure contagion is more prevalent among loans with large outstanding debt relative to the initial balance. For each foreclosure filing, we proxy the borrower being underwater at the time of filing by whether or not the lender’s claim against the borrower (i.e., the outstanding debt at filing) is larger than the initial loan principal. Although we observe the lender’s claim at the time of filing, we do not observe a direct measure of the property value and so we proxy the value using initial loan principal.⁴³ For each foreclosure case we split the count of neighboring foreclosures in two—filings among borrowers that are underwater according to our proxy ($N_{i,u,d(i)+t}$) and filings in positive equity ($N_{i,p,d(i)+t}$). We define contagion outcomes in terms of these two counts: $Y_{i,j,d(i)+t} = 1 [N_{i,j,d(i)+t} > 0]$, $Y_{i,j,d(i)+t} = N_{i,j,d(i)+t}$, and $Y_{i,j,d(i)+t} = \log(N_{i,j,d(i)+t})$ for $j \in \{u, p\}$, and estimate our baseline 2SLS specification jointly allowing the time-specific fixed effects to vary with the type of filings under consideration:

$$Y_{i,j,d(i)+t} = \beta_0 + \beta_u \cdot F_i \cdot 1[j = u] + \beta_p \cdot F_i \cdot 1[j = p] + \beta X_i + \Phi_i + M_{j,m(i)} + \psi_{j,d(i)+t} + u_{i,d(i)+t} \quad (7)$$

Pooling the estimates for both outcomes allows us to test the null hypothesis that a completed foreclosure has the same effect on underwater borrowers as borrowers in positive equity ($H_0 : \beta_u = \beta_p$).⁴⁴ We assume that, were it not for their differing loan to value ratios, borrowers in positive and negative equity respond similarly to a neighboring completed

⁴³We have experimented with alternative definitions of underwater, for example by making assumptions about loan-to-value ratios at origination to determine home value and adjusting this using local price indices or comparing outstanding balance at filing to local home values. Our results are not very sensitive to these changes.

⁴⁴Using log count of filings and the probability of any filing for each group avoids the problem that we do not observe the base number of loans in each category. Additionally, having a randomly assigned instrument ensures that, in expectation, neighborhoods with completed foreclosures have similar numbers of underwater borrowers as neighborhoods with dismissals.

foreclosure, and that lender behavior does not vary along these margins (i.e., any observed heterogeneity in estimates is driven by whether or not the neighbors are underwater on their loans).

We find that contagion is more prevalent among borrowers who are *not* in negative equity. We present the 2SLS estimates of Equation 7 in the upper panel of Table 1.11. We pool five years of post-decision observations ($t \in \{0, 1, 2, 3, 4\}$), restricting the treatment to be constant across all years (we present the yearly estimates in Table 1.27 of the Appendix). These estimates show that a completed foreclosure increases the probability of observing any non-underwater filing by 6.6 percentage points and total non-underwater filings by 9.6 percent, while having little effect on filings from underwater borrowers (0.023 and -0.3%). The difference between the two groups is significant.⁴⁵

We suspect that this finding that contagion is more prominent among borrowers who are not underwater is driven by borrowers who are on the margin of negative equity. Since being in negative equity is generally considered a necessary condition for default, it is unlikely that contagion is driven by borrowers who have lots of equity in their homes. At the same time, using loan principal at origination to proxy housing value in defining underwater and non-underwater loans may have some error on the margins—we may be classifying borrowers who are “just” underwater as being in positive equity. In this case of misclassification, our estimates suggest that the response may be coming from those who are in slight negative equity, rather than those who are very underwater. Either way, our estimates do not imply that those who are severely underwater will not default on their loans, but that they are not responsive to neighboring foreclosures, whereas those who are on the margin of negative

⁴⁵We also examine heterogeneity by local price growth, loan type and by zip-code income level (a proxy for ability to pay). Deng et al. (2000) argue that the decision to default depends on expectations of home value. We adjust our baseline contagion specification by interacting the foreclosure effect with an indicator for positive zip-code-year price growth, but find no evidence that price growth matters for contagion—see Table 1.28 in the Appendix. Campbell and Cocco (2011) argue that the probability of default is decreasing in ability to pay (income relative to loan payments) and that incentives for default vary depending on the structure of loan payments. We proxy ability to pay with zip-code-level median income (from the IRS SOI), but find no discernible difference in contagion by income quartile. These results are plotted in Figure 1.5 of the appendix. We also find no significant difference between the response among neighboring borrowers with conventional fixed-rate mortgages and those with alternative mortgage products—see Table 1.29.

equity are responsive. For example, a borrower who is severely underwater may default regardless of what happens to their neighbors. However, for a borrower with debt close to 100% of their home value, a nearby foreclosure may be quite important: not only might it push them into negative equity, but it might also convey information about the foreclosure process, lender behavior, or the future of the neighborhood (signals of less import to a severely underwater borrower who has thought through the default decision).

We further explore if contagion is driven by individuals on the margin of negative equity by comparing our 2SLS contagion estimates for filings among borrowers who have close to zero equity in their home (as measured by our proxy) to contagion for filings among borrowers with lots of equity or who are severely underwater. We define filings on the margin of underwater as those where the absolute difference between loan principal at origination (ρ) and the outstanding debt at filing (d) is less than 10% of principal: $\frac{\rho-d}{\rho} \in [-0.1, 0.1]$. As above with underwater and non-underwater filings, we split the count of neighboring filings in two—marginal and non-marginal filings—and estimate the analogue to Specification 7 for these two groups. The results, presented in the lower panel of Table 1.11, demonstrate that contagion is driven by borrowers on the margin of negative equity: a completed foreclosure induces a 6.3 percentage point increase in the probability of observing any filing for those on the margin versus an insignificant 0.8 for other filers, and a 10.7% increase in the number of new filings among those on the margin, versus an insignificant drop of 3.5% (in all cases the difference between the two groups is significant). Thus, we conclude that contagion is operating through borrowers who are on the margin of negative equity, rather than those who are reasonably well off or in dire straits. This interpretation is consistent with our finding from Section 1.5.1 that contagion is pronounced for cases decided in 2004–2008 when borrowers are less likely to be severely underwater, but not for cases decided at the depth of the foreclosure crisis (2009–2011) when foreclosures are common and many mortgage holders experience financial difficulty.

1.6.3 Foreclosure Contagion and Information

Contagion may also occur because borrowers learn about the foreclosure process, including the behavior of lenders and servicers, by observing neighboring foreclosure cases. Recent survey evidence suggests that borrowers learn from defaults and foreclosures within their social networks (Guiso et al. (2013)). For example, a neighbor may learn about the foreclosure process, including the costs of default (and completed foreclosure), how long the process takes, and the probability of a positive resolution (e.g., mortgage modification) by observing his/her neighbor’s experience.⁴⁶ A priori, it is not clear whether the foreclosure event would increase or decrease the probability of neighboring filings (contagion vs. anti-contagion). For instance, a foreclosure may lower neighbors’ perception of the probability of renegotiating one’s loan by defaulting, thus lowering the expected value of default and discouraging this behavior (Mayer et al. (2011) find evidence of borrowers defaulting in response to an increase in the probability of modification).

We investigate a specific social network—neighbors with loans from the same lender.⁴⁷ Individuals with the same lender may be more likely to discuss the foreclosure (and/or mortgage renegotiation) process with one another. At the same time, a successful (or failed) mortgage renegotiation provides a stronger signal to individuals with loans from the same institution; it may be that a neighbor’s foreclosure discourages default among those with loans from the same servicer/lender by lowering the perceived probability of a modification. To test for the presence of learning-based contagion, we test whether contagion is stronger/weaker among neighboring loans from the same lender. Under the assumptions

⁴⁶Guiso et al. (2013) also find that social norms may matter for contagion. Many (82.7%) respondents to their survey feel they have a moral obligation to repay their debts. But a neighboring default may weaken respondents’ sense of moral responsibility and increase their probability of default. We do not expect this social morality channel to be particularly strong in our context, as our estimates compare neighborhoods around properties where the borrower has already defaulted—the event that sends a signal about the moral obligation to repay.

⁴⁷We have also examined heterogeneity in the treatment effect across neighborhoods where we expect weaker and stronger social networks. We study how contagion varies with geography-based proxies for neighbor connectivity, including racial homogeneity, population density, and housing type (drawing on the notion of social capital outlined by Glaeser and Sacerdote (2000)), although we find no systematic relationship between contagion and these proxies. We include a discussion of these estimates in the Appendix.

that i) individuals are aware of (at least some of) their neighbors' lenders and ii) lenders are not reacting to local conditions, then the difference between contagion among same-lender and different-lender loans provides evidence of the importance of learning from neighboring foreclosures.⁴⁸ Note that this signal may be present as long as a borrower knows who his/her neighbor's lender is and observes the outcome of the case. Since a borrower may learn about lender behavior by simply observing the outcome of his/her neighbor's case, heterogeneity in contagion along this margin provides evidence that borrowers learn from their neighbors' experiences, but not necessarily that communication through social networks matters.

To test for a difference in contagion between same-lender and different-lender borrowers, we estimate our 2SLS contagion effects for each subset of neighboring filings. We split the count of neighboring foreclosures in two—filings with the same lender listed as the plaintiff ($N_{i,s,d(i)+t}$) and filings with any other lender listed as the plaintiff ($N_{i,o,d(i)+t}$), and redefine our outcomes in terms of these two counts: $Y_{i,j,d(i)+t} = 1 [N_{i,j,d(i)+t} > 0]$, $Y_{i,j,d(i)+t} = N_{i,j,d(i)+t}$, or $Y_{i,j,d(i)+t} = \log(N_{i,j,d(i)+t})$ for $j \in \{s, o\}$. We then estimate our baseline 2SLS specification jointly allowing the time-specific fixed effects to vary with lender type:

$$Y_{i,j,d(i)+t} = \beta_0 + \beta_s \cdot F_i \cdot 1[j = s] + \beta_o \cdot F_i \cdot 1[j = o] + \beta X_i + \Phi_i + M_{j,m(i)} + \psi_{j,d(i)+t} + u_{i,d(i)+t} \quad (8)$$

Pooling the estimates for both outcomes (same-lender filings and other-lender filings) allows us to test the null hypothesis that a completed foreclosure has the same effect on mortgages held by the same lender as mortgages held by other lenders ($H_0 : \beta_s = \beta_o$).

The 2SLS estimates of β_s and β_o from Specification 8, presented in Table 1.12 (pooling all years; see Table 1.33 for yearly estimates), suggest that lender-specific local networks matter. The estimates show that contagion is primarily driven by foreclosure filings among loans held by different lenders—0.061 percentage point increase in the probability of any

⁴⁸Our same-lender and different-lender estimates are robust to restricting the sample to lenders known for automating the foreclosure process, suggesting that lenders are not driving the patterns we observe in these estimates.

new filing on a loan from a different lender versus an (insignificant) increase of 0.006 for loans from the same lender, 0.770 additional foreclosure filings from different lenders versus a drop in filings of 0.184 among loans held by the same lender, and an increase of 0.085 in the log of total filings among different lenders versus 0.008 for filings on loans from the same lender. In the first two cases, the difference between the estimates are statistically significantly different from zero.

We interpret these same-plaintiff results as evidence that borrowers learn about lenders from the experience of their neighbors.⁴⁹ Given that there is general contagion among loans from other lenders, it appears that borrowers experience anti-contagion when a neighbor from the same lender ends up in a completed foreclosure. A possible explanation for this difference is that the neighboring foreclosure sends different information to different individuals: contagion among those with different lenders is consistent with the foreclosure lowering neighboring home values or sending a broad signal about the general direction of the neighborhood (i.e., the value of the neighborhood is deteriorating). Anti-contagion among borrowers with the same plaintiff is consistent with borrowers revising downward the probability of a positive outcome (i.e., not losing their home) when a neighboring borrower with the same lender is unsuccessful. This lowered expectation of a positive outcome decreases the value of (strategic) defaults.⁵⁰

⁴⁹Another possibility is that lenders are aware of the externalities of completed foreclosures and do not pursue new filings in neighborhoods where they have had a recent successful foreclosure (while other lenders do not react). However, we find this to be a less likely explanation for two reasons. Firstly, we still observe, on average, a positive and significant number of foreclosure filings around foreclosed properties in the year of and years after the foreclosure decision. Thus, it does not appear that lenders are avoiding foreclosure externalities altogether. Secondly, the anti-contagion effect among filings from the same lender persist for several years after the decision, which represents a long time for lenders to be waiting to avoid foreclosure externalities (while other lenders continue to file in spite of these externalities).

⁵⁰Given our instrument, it is also conceivable that borrowers learn about the impact of a given judge on foreclosure outcomes. However, we do not expect that foreclosure contagion is driven by a completed foreclosure revealing information about a specific judge. Firstly, given the random assignment of judges to foreclosure cases, the probability that a given borrower ends up with a given judge is low—thus, the neighboring borrower’s expectations about the outcome of default should not change substantially. Secondly, if learning about judges is the primary driver of foreclosure contagion, then we would not expect to observe differential contagion among individuals with the same lenders.

1.7 Conclusion

We provide clean estimates of the effect of a completed foreclosure on neighboring residential sale prices and on neighboring foreclosure filings in Cook County, IL. We exploit a randomly assigned instrument—the set of judges who hear a foreclosure case—to compare neighborhood-level outcomes around a delinquent property that ends in completed foreclosure to a delinquent property whose foreclosure claim is dismissed.

We find robust evidence of foreclosure contagion. A completed foreclosure leads to about 0.5 to 0.70 more foreclosure filings per year within 0.1 miles and increases the probability of observing any neighboring foreclosure filing by about 10%. Moreover, a completed foreclosure causes between 0.25 and 0.5 new completed foreclosures per year. These contagion effects persist for at least four years after the case is decided.

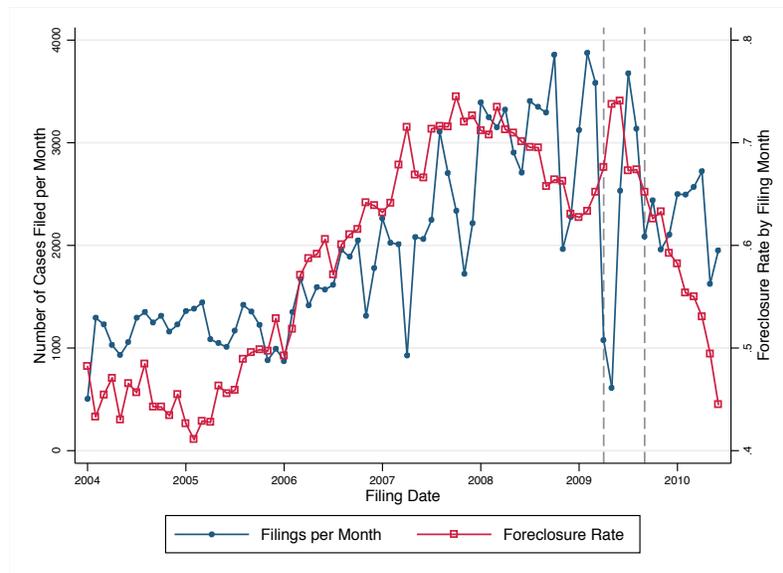
Contagion is primarily driven by borrowers who are on the threshold of default. We find contagion among loans held by lenders who are known to automate foreclosure filings and are likely unresponsive to very local conditions, which we interpret as evidence that foreclosure contagion is driven by borrowers. Contagion is strongest not during the depths of the crisis, but at the end of the housing boom and beginning of the crash. Moreover, contagion is most prevalent among borrowers who are on the cusp of being underwater and not those who are severely underwater on their loans. We interpret this as evidence that a neighboring foreclosure has the greatest impact on borrowers who are on the margin of defaulting, and not those who are severely at risk. Finally, we find that contagion is minimal when borrowers have the same lender, perhaps because the neighboring foreclosure sends a signal about their lender's behavior, lowering the perceived probability of a successful renegotiation of the loan, thus reducing strategic incentives to default.

We find evidence that completed foreclosures disrupt local housing markets. After a completed foreclosure, the mean residential sale price dips by as much as 40%. However, this drop is largely explained by selection into sale—in the wake of a completed foreclosure, the composition of residential sales is skewed toward lower quality (and thus, lower price)

homes. Nonetheless, we cannot rule out that foreclosures do influence the value of homes, conditional on quality.

While our instrumental-variables method provides clean identification of the effect of foreclosure, the resulting estimates are of a particular parameter that helps to inform foreclosure policy. Our estimates represent the effect of a completed foreclosure on the neighborhoods around properties that are most likely to be influenced by foreclosure judges. These are the cases that are most likely to be influenced by policy. At the same time, our estimates compare the neighborhood-level effect of a completed foreclosure relative to a delinquent mortgage that does not end in foreclosure. This is the relevant parameter for assessing policy that addresses how easily and how often lenders should be able to foreclose on delinquent borrowers. Finally, while there may be concerns about the external validity of our estimates, which are derived from a housing crisis in one of the worse-hit cities, it is in exactly such circumstances that policymakers and economists must worry most about foreclosure contagion. In sum, our estimates of foreclosure contagion suggest room for policy that seeks alternative solutions for delinquent borrowers.

Figure 1.1: Foreclosure Cases Over Time in Cook County, IL



Notes: Monthly count of new foreclosure filings in Cook County over time (left axis) and share of cases filed in a given month that end in a completed foreclosure (right axis). Dashed vertical lines indicate Cook County suspension of all new foreclosure filings starting April 16, 2009, except for so-called “consent foreclosures:” foreclosure filings in which lender and borrower had already agreed to foreclosure prior to filing. This “moratorium” was scheduled to end on September 1, 2009, although appears to have ended earlier, given the spike in filings prior to Sept. 1.

Table 1.1: Descriptive Statistics: Pre-Treatment Characteristics

	Mean			P-Value
	All	D^\dagger	F^\dagger	$(H_0 : D = F)$
Case Resolved	0.950	1.000	1.000	.
Ends in Foreclosure	0.612	0.000	1.000	.
Days to Decision	373.554	274.931	428.665	0.000
Prob Redefault	0.049	0.120	0.013	0.000
Single-Family Property	0.623	0.684	0.590	0.000
Conventional Mortgage	0.647	0.665	0.653	0.000
Loan Principal	237328.100	219520.200	237733.000	0.892
Complaint Amount	229621.000	211420.400	231905.600	0.449
Large Plaintiff	0.471	0.472	0.473	0.065
Large Attorney	0.685	0.682	0.687	0.032
Median Income (tract) ^{††}	44859.720	46409.490	43748.260	0.000
Share White (tract) ^{††}	0.449	0.468	0.433	0.000
Population (tract) ^{††}	5434.558	5472.077	5411.712	0.000
N	148220	50140	90653	

Notes: Data from matched court records and foreclosure filings (one obs per case) with baseline sample restrictions as described in text.

[†] D = dismissed cases, F = completed foreclosures.

^{††}Data from 2000 Census (tract-level).

Table 1.2: Descriptive Statistics: Outcomes

	Means			P-Value	N
	All	D^\dagger	F^\dagger	$(H_0 : D = F)$	
Neighboring Filings	2.423	2.161	2.596	0.000	475127
Any Neighboring Filing	0.734	0.702	0.755	0.000	475127
Neighboring Foreclosures	0.748	0.603	0.844	0.000	475127
Any Neighboring Foreclosure	0.365	0.326	0.390	0.000	475127
Neighboring Sales	3.038	2.962	3.099	0.006	133176
Mean Neighboring Sale Price	169326	184213	157182	0.000	81371
Mean Repeat-sale Adjusted Price	122896	123240	122621	0.205	56306

Notes: Outcome variables (0.1 mile radius) measured annually for five years (two years for sales outcomes) after case is decided (observation = one case-year).

[†] D = dismissed cases, F = completed foreclosures.

Table 1.3: Balance of Covariates

	Coefficient [†]	<i>p</i> Value ^{††}
Adj. Rate Mortgage	0.0005100 (0.0006240)	0.3837
Loan Principal ^{†††}	-0.0000033 (0.0000025)	0.0343**
Large Plaintiff	-0.0005790 (0.0003840)	0.9438
Large Attorney	-0.0007020 (0.0007410)	0.1370
Median Income (tract) ^{†††}	-0.0002780 (0.0001730)	0.6566
Share White (tract)	-0.0000439 (0.0005620)	0.4025
Population Density (tract)	0.0454000 (0.0355000)	0.1364
Price (Zip code) ^{†††}	0.0000541 (0.0000410)	0.8564
Total Foreclosure Filings ^{††††}	0.0000266 (0.0000570)	0.3578
Total Completed Foreclosures ^{††††}	0.0001580 (0.0001230)	0.5866
<i>p</i> Value	0.4280	
<i>N</i>	143276	

Notes: **Indicates significance at the 5% level. Standard errors are the max of the SE clustered on census tract, clustered on filing month, and multi-way clustered on tract and filing month. [†]Coefficient from regression of instrument (case-calendar-filing-month foreclosure rate) on given pre-treatment covariates, controlling for filing month and property type fixed effects. *P* value in first column is from a joint significance test for the given covariates. ^{††}Given covariate (for that row) is regressed on full set of case calendar dummies (plus filing month and property type fixed effects); *p* value for a joint significance test of the case calendar dummies. ^{†††}In \$10,000 of dollars. ^{††††}Outcomes as defined in Section 1.4 measured the year before the case is filed.

Table 1.4: Complier Characteristics: Ratio of Subgroup First Stage Estimate to Overall First Stage

Quartile of Income [†]			
1	2	3	4
0.784	1.011	1.060	1.066
Loan Characteristics ^{††}			
Large Lender	Conventional Mortgage	Positive Zip Code Price Growth	Underwater
1.002	0.896	0.795	0.907

Notes: [†]Income quartile is given by the tract-level quartile of median tract income from the 2000 Decennial Census. ^{††}Large lender is an indicator for the plaintiff being one of the six most prominent banks in the sample, each representing $\geq 10\%$ of filings. Zip-code-level annual price growth is taken from Zillow housing price indices for Cook County. Underwater is a proxy for the outstanding balance of the mortgage being greater than the estimated value of the home, as described in the text.

Table 1.5: Baseline Contagion Estimates: 2SLS Coefficient of Effect of Completed Foreclosure on Given Outcome in Given Year

Years After End of Case (t):	0	1	2	3	4	5
Baseline Estimates						
$Y_{id(i)+t}$ = Any Filing per Year	0.052*	0.012	0.082***	0.090*	0.247**	0.140
	(0.028)	(0.027)	(0.027)	(0.053)	(0.112)	(0.133)
1st-stage F	238.300	224.200	205	24.620	19.060	12.970
N	130199	118566	93143	67379	41958	23831
$Y_{id(i)+t}$ = Total Filings per Year	0.691*	0.670*	0.536***	0.657**	1.551	0.538
	(0.393)	(0.368)	(0.183)	(0.319)	(0.983)	(0.987)
1st-stage F	238.300	224.200	205	24.620	19.060	12.970
N	130199	118566	93143	67379	41958	23831
Cumulative Count of Filings						
$Y_{id(i)+t}$ = Cumulative Filings	0.691*	1.395*	2.090***	4.522***	6.446**	5.610
	(0.393)	(0.732)	(0.794)	(1.169)	(3.105)	(4.610)
1st-stage F	238.300	224.200	205	24.620	19.060	12.970
N	130199	118566	93143	67379	41958	23831
$Y_{id(i)+t}$ = Cumulative Filings	0.622	1.266*	1.594**	3.820***	4.076	2.693
No Recent Foreclosures	(0.380)	(0.667)	(0.703)	(1.429)	(2.918)	(4.611)
1st-stage F	172.500	159.700	138	20.910	12.790	7.515
N	71925	66995	55105	43186	27395	14948

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates of the effect of completed foreclosure on given outcome (measured within 0.1 miles of the property in the given year since the case is decided), on an indicator for the case ending in foreclosure (instrumented by the leave-one-out case-calendar-filing-month-specific foreclosure rate), filing month, property type, and year of observation fixed effects, and case-level controls:share of tract that report race as white in 2000 decennial census, income quartile from decennial census, whether plaintiff is a “large plaintiff” (six largest plaintiffs each representing ≥ 7000 filings) or attorney is a “large attorney” (three largest attorneys each representing $\geq 10,000$ cases), whether mortgage is adjustable rate, size of initial loan, and census tract population. Cumulative filings represent the total number of new filings since decision through the given year.

Table 1.6: Constant Sample Contagion Estimates

Years After End of Case (t):	0	1	2	3	45
Constant Sample Observed for 3 Years after Case Decision (Decided in 2004–2008)					
$Y_{id(i)+t}$ = Any Filing per Year	0.094** (0.048)	0.059 (0.046)	0.168*** (0.061)	0.090* (0.053)	
1st-stage F	24.620	24.620	24.620	24.620	
N	67379	67379	67379	67379	
$Y_{id(i)+t}$ = Total Filings per Year	1.243** (0.552)	1.567*** (0.395)	1.056*** (0.391)	0.657** (0.319)	
1st-stage F	24.620	24.620	24.620	24.620	
N	67379	67379	67379	67379	
Cases Observed for 2 Years or Fewer after Decision (Decided in 2009–2011)					
$Y_{id(i)+t}$ = Any Filing per Year	0.027 (0.037)	-0.018 (0.038)	0.026 (0.038)		
1st-stage F	72.890	69.410	66.700		
N	62820	51187	25764		
$Y_{id(i)+t}$ = Total Filings per Year	0.239 (0.447)	0.099 (0.515)	0.202 (0.263)		
1st-stage F	72.890	69.410	66.700		
N	62820	51187	25764		

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates of the effect of completed foreclosure on given outcome (measured within 0.1 miles of the property in the given year since the case is decided), on an indicator for the case ending in foreclosure (instrumented by the leave-one-out case-calendar-filing-month-specific foreclosure rate), filing month, property type, and year of observation fixed effects, and case-level controls:share of tract that report race as white in 2000 decennial census, income quartile from decennial census, whether plaintiff is a “large plaintiff” (six largest plaintiffs each representing ≥ 7000 filings) or attorney is a “large attorney” (three largest attorneys each representing $\geq 10,000$ cases), whether mortgage is adjustable rate, size of initial loan, and census tract population. Constant sample includes only cases for which we observe three years post decision.

Table 1.7: Estimates in Years Before Filing

Years Before Case is Filed	3	2	1
Any Filing	0.003 (0.029)	0.009 (0.029)	0.018 (0.022)
1st-stage F	400.000	400.000	400.100
N	140672	140672	140672
Total Filings per Year	0.020 (0.087)	0.084 (0.122)	0.214 (0.139)
1st-stage F	400.000	400.000	400.100
N	140672	140672	140672
Any Completed Foreclosure	-0.027 (0.023)	0.045* (0.025)	-0.004 (0.030)
1st-stage F	159.400	182.100	164.300
N	140672	140672	140672
Total Completed Foreclosures per Year	-0.057* (0.033)	0.075 (0.065)	0.110 (0.071)
1st-stage F	159.400	182.100	164.300
N	140672	140672	140672
log(price)	0.007 (0.038)	0.037 (0.042)	0.007 (0.049)
1st-stage F	164.900	166.500	136.400
N	109123	102220	84126

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, with outcomes measured in the given year prior to the case being filed.

Table 1.8: Contagion in Completed Foreclosures

Years After End of Case	0	1	2	3	4	5
	Baseline Estimates					
Any Completed Foreclosure	-0.048 (0.031)	0.002 (0.030)	0.045 (0.038)	0.138*** (0.051)	0.033 (0.106)	0.174 (0.144)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
Total Completed Foreclosures	0.280* (0.164)	0.435*** (0.160)	0.312*** (0.117)	0.558*** (0.185)	0.243 (0.316)	0.063 (0.334)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
	Contagion in Completed Foreclosures Filed Prior To Decision					
Any Completed Foreclosure	-0.043 (0.031)	0.135*** (0.028)	0.073*** (0.019)	0.025 (0.016)	-0.010 (0.017)	0.003 (0.015)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
Total Completed Foreclosures	0.303* (0.164)	0.609*** (0.118)	0.111*** (0.028)	0.031 (0.019)	-0.010 (0.017)	0.003 (0.015)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure outcome are as in Table 1.5.

Table 1.9: Baseline Housing Market Estimates

Years Since End of Case		0	1	2
log(price)	2SLS	-0.127 (0.105)	-0.411** (0.203)	-0.358 (0.395)
	1st-stage F	22.500	18.840	14.390
	N	43079	26047	12241
log(price) (repeat-sales)	2SLS	0.059 (0.066)	0.003 (0.146)	-0.251 (0.201)
	1st-stage F	11.130	6.314	7.355
	p-value	0.107	0.030	0.763
	N	30482	17916	7904
Total Sales (Cumulative Over Years)	2SLS	0.220 (0.914)	6.562 (4.574)	9.964 (9.255)
	1st-stage F	24.620	19.060	12.970
	N	67379	41958	23831
log(price at last sale) [†] (repeat sample)	2SLS	-0.031 (0.116)	-0.373 (0.230)	-0.186 (0.323)
	1st-stage F	21.000	12.400	13.910
	N	29620	17597	7852

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure outcome are as in Table 1.5. Sales outcomes represent residential transactions within the given radius and time period only (total sales represents the total count of sales since the decision year), while price represents the mean sale price of these transactions. Repeat-sales adjusted prices are estimated as outlined in the text; associated p-value is for the cross-equation test of equality between the repeat-sales adjusted and unadjusted (first panel) price effects. [†]Price at previous sale (for repeat sales) adjusted for annual price growth. Fewer observations here since sample is restricted to the latter of all repeat sales (need to observe a previous sale).

Table 1.10: Contagion Among Loans with Lenders Implicated in Independent Foreclosure Review Settlement

Outcome Variable		Any Filing	Total Filings	Log Filings
Baseline [†]	Effect of Completed Foreclosure	0.062*** (0.022)	0.735** (0.320)	0.094* (0.050)
	First-stage F	131.900	131.900	147.100
	<i>N</i>	311116	311116	235325
Automating Lenders [†]	Effect of Completed Foreclosure	0.039* (0.022)	0.115 (0.073)	0.084* (0.049)
	First-stage F	131.900	131.900	138.000
	<i>N</i>	311116	311116	86054

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS specification as in Table 1.5, although all five years post-decision are pooled and the effect of completed foreclosure is fixed.

[†]Outcomes are measured based on new foreclosure filings by lenders implicated in the Office of the Comptroller of the Currency and the Board of Governors of the Federal Reserve System Independent Foreclosure Review Settlement, while “Baseline” includes new foreclosure filings among all lenders. Both samples are restricted to cases decided between 2007 and 2010 (during which automation of foreclosure filings is thought to be most common).

Table 1.11: Contagion Estimates By Proxy for Borrower Equity

Outcome Variable	Any Filing	Total Filings	Log Filings
Neighboring Filings from Underwater versus Non-Underwater Borrowers			
Effect of Foreclosure on Filings of Non-Underwater Borrowers [†]	0.066*** (0.020)	0.696*** (0.230)	0.096** (0.043)
Effect of Foreclosure on Filings of Underwater Borrowers [†]	0.023 (0.019)	-0.066 (0.062)	-0.003 (0.035)
First-stage F	84.220	84.220	86.940
<i>p</i> val for difference between groups	0.053	0.000	0.006
<i>N</i>	902490	902490	441734
Neighboring Filings from -10% – 10% Equity versus < -10% or > 10%			
Effect of Foreclosure on Filings with Equity < -10% or > 10% ^{††}	0.008 (0.017)	-0.159*** (0.054)	-0.035 (0.035)
Effect of Foreclosure on Filings with Equity between -10% and 10% ^{††}	0.063*** (0.020)	0.789*** (0.245)	0.107** (0.045)
First-stage F	84.220	84.220	89.350
<i>p</i> val for difference between groups	0.017	0.000	0.001
<i>N</i>	902490	902490	416141

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, pooling observations for five years after the case is decided (fixing the effect of completed foreclosure to be constant).

[†]Outcomes are for two separate counts of new foreclosure filings by either borrowers in negative equity or borrowers in positive equity (as defined in the text). P-value tests the significance between the two responses (positive vs. negative equity)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

^{††}Outcomes are for two separate counts of new foreclosure filings by either borrowers where the difference between outstanding debt and loan principal is within 10% of initial loan principal and all other borrowers (see definition in the text). P-value tests the significance between the two responses—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

Table 1.12: Contagion Estimates By Lender Identity

Years Since Decision	Any Filing	Total Filings	Log Filings
Filings from Different Lender [†]	0.061*** (0.022)	0.770*** (0.245)	0.085** (0.041)
Filings from Same Lender [†]	0.006 (0.015)	-0.184*** (0.056)	0.008 (0.045)
First-stage F	87.490	87.490	93.560
<i>p</i> val for difference between groups	0.052*	0.000***	0.109
<i>N</i>	796112	796112	330880

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, pooling observations for five years after the case is decided (fixing the effect of completed foreclosure to be constant).

[†]Outcomes are for two separate counts of new foreclosure filings by either different lenders than that in the observed case or the same lender. P-value tests the significance between the two responses (same lender vs. other lender)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

1.8 Appendix

1.8.1 Data Appendix

Cleaning Court Records We collected all chancery court case records filed between January 2004 and June 2010 (inclusive). We extract from each record the associated case number and the case calendar to which the case is assigned. The records also contain a list of case actions, the lawyer who initiated this action, the associated judge, and the date. We extract this list of actions (simple text descriptions, e.g., “Amend complaint or petition - allowed” or “Dismiss by stipulation or agreement”) and the corresponding dates.

We identify a case as ending in a dismissal if an action occurs containing one of the following descriptions: "mortgage foreclosure motion plaintiff dismissed", "mortgage foreclosure voluntary dismissal, non-suit or dismiss by agreement", "mortgage foreclosure motion defendant dismissed", "mortgage foreclosure dismissed for want of prosecution", "dismissed for want of prosecution", "general chancery - dismissed for want of prosecution", "general chancery - voluntary dismissal, non suit, dismiss by agreement", "mortgage foreclosure voluntary dismissal, non-suit or dismiss by agreement", or "mortgage foreclosure judgment for defendant"; or an action containing any of the following: "case dismissed", "voluntary dismissal", "declaratory judgment voluntary dismissal", "dismiss entire cause" and not "denied", or "dismiss by stipulation or agreement" and not "denied". For dismissed cases, we consider the end of the case to be the date of this “dismissal” action (in the case of multiple such actions, we take the final).

Cleaning RIS Data Record Information Services, Inc. provided us with details of foreclosure filings and foreclosure auctions for Cook County from 2002 through 2011. RIS is a private data provision company that collects publicly available records on all foreclosure filings in the five counties of Chicago. RIS employees manually input data on each foreclosure filing. From the foreclosure filings, we extract the associated chancery court case number, unique loan ID, the filing date, details of the associated loan (origination date, principal at

origination, outstanding claim at time of foreclosure filing, a general indication of mortgage type—conventional, adjustable rate, etc.), details of the associated property (latitude and longitude, census tract, zip code, property type—condo, single family, etc.), and the parties involved (defendant name, plaintiff—the lender or servicer—identity, plaintiff law firm).

We identify a case as ending in a completed foreclosure if there is an associated foreclosure auction record in the RIS data. For completed foreclosures, we use the date of the foreclosure auction as the end-date of the case. If there is both a dismissal action in the court records and an associated foreclosure auction, we consider the case to have ended in a completed foreclosure, although, our results are not sensitive to this decision. Relatedly, there is a field in the RIS data that indicates the outcome of the auction, including if the auction is canceled. Since this information is missing for half of the years and since it is not indicated why a cancellation occurs, we do not code canceled auctions as dismissals in our analysis sample. Again, however, our baseline results are not sensitive to coding canceled auctions as dismissals. We consider a borrower to have redefaulted if a foreclosure filing is brought against the same loan ID. Note that this will not count new filings at the same property for different loans (e.g., if the home owner has filings against mortgage and a home equity loan, we count these as distinct filings). However, if a second filing against the same loan ID occurs within 180 days, we consider this to be the same case, taking the first date as the true filing date (and merging info from the two filings). The descriptive statistics in Table 1.1 show a non-zero number of redefaults (1.3%) among loans that end in foreclosure. This is likely due to miscoding in the RIS data—for example, a foreclosure auction is scheduled and recorded by RIS, the case is dismissed before the auction takes place (and RIS misses this) and the borrower subsequently redefaults on the loan. Our results are not sensitive to discarding these observations.

We construct a unique ID for each plaintiff and attorney as follows. For plaintiffs and attorneys who are on more than ~ 100 cases, we manually checked the names for consistency and constructed a unique ID number. We then identified “large plaintiffs” as those plaintiff-

IDs associated with greater than 7000 cases, and “large attorneys” as those attorney-IDs associated with greater than 10000 cases. We identify by name all plaintiffs implicated in the Independent Foreclosure Review Settlement conducted by the Office of the Comptroller of the Currency and the Board of Governors of the Federal Reserve System.

We match the RIS and chancery court records by case number. We discard non-matches, which arise due to several factors: non-foreclosure chancery court cases (e.g., name changes, mechanic’s liens) will not appear in the RIS foreclosure filings; differing date ranges between the two data sources (2004–2010 for court records, 2002–2011 for RIS); and differing geographies (RIS data includes some cases in neighboring counties).

Cleaning Census Data We merge in the following census-tract-level data from the 2000 Decennial Census: median tract income, population, land area, and share of population that identifies as each census-designated race. We construct tract-level income quartiles (i.e., what quartile of median income does a given tract fall into), an indicator for being a predominantly white tract (share white is greater than the median share), population density (and associated quartiles), and a Herfindahl-Hirschman index for each race (i.e., the sum of the squared share of each race in the tract). We merge the census tract data to the RIS data using census tract FIPS codes.

Cleaning Deeds Records These data are collected from the county recorder for Cook County, IL. The records were collected by an anonymous private firm and made available to us by the Paul Milstein Center for Real Estate at the Columbia Graduate School of Business. These data include the date of each sale property transaction, the type of property, the address, the price of the sale, and an indicator for the property being residential. We drop all transactions with sale price or address missing. We drop duplicate records—multiple sales with identical sale prices that occur at the same property within 30 days of one another. We keep only residential sales. We geocode these deeds records based on the property address using Yahoo! Placefinder.

Defining Outcomes Using the cleaned and matched RIS and court records, for each foreclosure case we calculate the distance between the associated property and the properties associated with all other foreclosure filings. We then count the number of new filings around each property (i.e., within the given radius; 0.1 miles in the baseline) in each calendar year, omitting from the count new filings at the same property or filings associated with the same case—e.g., a given loan may be tied to multiple properties, which we do not want to include in the count. If there are multiple foreclosure filings at a neighboring property, we include each of these in the count (although we have found that our results are not sensitive to treating these as a single new filing). For each calendar year, we then construct an indicator for there being any new foreclosure filing. We then identify the year that the case is decided and define our contagion outcomes relative to that year: number of filings and indicator for any filing in the year the case ends, number and indicator one year after the case ends, two years after, and so on. We follow the same procedure to get a count of new completed foreclosures in each calendar year: calculate the distance between the property associated with each case and each completed foreclosure (from the RIS auction data), and count the auctions that occur within the given radius in the given calendar year (where this date is based on the auction date).

We also construct several sub-counts of new foreclosure filings. Foreclosure filing records in the RIS data have a field reporting additional lien holders that are listed on the foreclosure claim (reporting additional lien holders is optional for plaintiffs). For each filing we create an indicator for additional lien holders. We then construct the same contagion outcomes, but only counting filings with multiple lien holders. Similarly, we construct our contagion outcomes using the annual count of new filings from plaintiffs implicated in the Independent Foreclosure Review Settlement and, as described in the text, use this to investigate contagion among lenders known for the bulk processing of delinquent loans into foreclosure filings. Thirdly, we construct an indicator for the borrower being underwater on their loan: we take the principal of the mortgage at origination, conservatively assume an 80% loan-to-

value ratio to back out the value of the property, adjust the value of the property using Zillow monthly zip code housing price indices for the year of origination and the year of the foreclosure filing, and compare the adjusted value of the property to the claim made against the borrower by the plaintiff. If the claim is larger than the value of the property, then we consider the borrower to be “underwater”. For each property associated with a foreclosure filing, we then find the annual count of filings against underwater borrowers within the given radius. We have experimented with other ways to construct this underwater indicator—claim larger than loan principal at origination, claim larger than 110% of estimated value, and so on—and find little difference in our results. Finally, we construct contagion outcomes for filings from the same lender. We restrict our sample to the set of cases for which we cleaned the plaintiff name (i.e., filings with plaintiffs who appear on approximately more than 100 filings). For each case within this subsample, we identify all filings within the given radius and create an indicator for the neighboring filing having the same plaintiff. We then find the annual count of new filings from the same plaintiff and new filings from other plaintiffs.

We construct our housing market outcomes in a very similar way. For each foreclosure filing, we calculate the distance between the associated property and all residential sales in the deeds records. As with the filings, we want to exclude sales at the property associated with the completed foreclosure. However, in this case (and unlike the filings) we do not have a perfect match of address—we cannot precisely identify a sale associated with the property going through foreclosure. Instead, we drop sales within 0.01 miles. For each calendar year, we take the mean sale price of all sales occurring within that year within the given radius (0.1 miles in the baseline) of the property associated with the foreclosure filing. We then use as an outcome the log of this sale price. In the following section, we discuss how we use repeat sales to adjust this average sale price for fixed property quality. We also use as an outcome the count of residential sales that occur within the given radius in the given calendar year and an indicator for any sale occurring.

1.8.2 Adjusting Price Data for Property Quality Using Repeat Sales

In our deeds records, we first identify all repeat sales: of the 1,330,949 residential sales we observe between 1995 and 2008, there are 585,756 (44.01%) properties that transact more than once, which leaves us with 216,068 transactions during the relevant period of 2004 – 2008 (43.56% of the 496,055 residential transactions in this period). For each property, k , we assume that the sale price in year t , P_{kt} , is a function of the property’s time-invariant characteristics, δ_k , the year of sale, and whether or not there is a recent foreclosure nearby, $F_{i(k)t}$:

$$\log(P_{kt}) = \alpha_0 + \alpha\delta_k + \Psi_t + \alpha_F F_{i(k)t} + \epsilon_{kt} = \alpha_0 + \alpha\delta_k + \Psi_t + e_{kt} \quad (9)$$

where Ψ_t is a year-specific fixed effect and we denote for convenience $e_{kt} \equiv \alpha_F F_{i(k)t} + \epsilon_{kt} = \log(P_{kt}) - \alpha_0 - \alpha X_k - \Psi_t$. We want a measure of the sale price, P_{kt}^* , that removes the influence of property characteristics, but allows price to vary with foreclosure:

$$\log(P_{kt}^*) = \beta_0 + \phi_t + e_{kt}$$

To achieve this, we estimate a simple price regression that controls for property and year of sale for all repeat sales in our sample: $\log(P_{kt}) = \beta_0 + \theta_k + \Psi_t + e_{kt}$, where θ_k is a vector of property fixed effects and Ψ_t is a vector of year-of-sale fixed effects. Property fixed effects absorb the influence of the (time-invariant) property characteristics. Using the OLS parameter estimates and residuals from this model, we then estimate $P_{kt}^* = \exp(\hat{\beta}_0 + \hat{\Psi}_t + \hat{e}_{kt})$. Using these quality-adjusted sales prices, we then construct a quality-adjusted measure of sale price for each property i going through the foreclosure courts by taking the log of the average of all P_{kt}^* that transact within x miles of property i in the relevant year of observation.

1.8.3 Monotonicity of Instrument

A failure of monotonicity occurs if a higher value of the instrument means a higher probability of foreclosure for some cases, but a lower probability for others. As discussed in the main text, a failure of monotonicity may arise if judges treat different types of borrowers and lenders differently.

We examine this possibility by relating foreclosure rates for each case calendar for different subgroups to the overall value of the instrument for that case calendar. We want to check that a higher value of the instrument for the case calendar is associated with a higher foreclosure rate for the sub-groups. We first calculate the overall foreclosure rate by case calendar and filing year, and de-mean these estimates by filing year.⁵¹ We then take a given covariate (e.g., the borrower is from a predominantly white neighborhood) and calculate the foreclosure rate by case calendar, filing year, and the value of the covariate (e.g., foreclosure rate by case calendar, filing year, and whether predominantly white neighborhood), and again de-mean by filing year. We plot the de-meaned group-specific foreclosure rates against the de-meaned general case-calendar-filing-year foreclosure rate and display these plots in Figures 1.2, 1.3, and 1.4. A failure of monotonicity as we described above suggests that for certain subgroups a higher general case-calendar foreclosure rate is associated with a higher group-specific foreclosure rate, while for other subgroups a higher general case-calendar foreclosure rate is associated with a lower group-specific foreclosure rate. We construct six such plots: i) comparing foreclosure rates between properties in census tracts where the share of white residents is greater than the median to those below the median, ii) comparing properties in each quartile of median tract-level income, iii) comparing foreclosure rates among conventional (fixed-rate) mortgages vs. unconventional mortgages (adjustable rate, interest only, etc.), iv) comparing foreclosure rates by property type, v) comparing cases where the

⁵¹Recall that our estimates all include filing date fixed effects, so the relevant comparison is within filing date, although the figures are similar if we do not de-mean. We use filing year for this exercise to decrease noise in the foreclosure rate estimates; the story does not change if we use filing month, although the associated figures are noisier.

plaintiff is a large lender to those with smaller lenders, and vi) comparing foreclosure rates among cases where the lender’s attorney is a large vs. smaller attorney (as previously defined). Figures 1.2, 1.3, and 1.4 shows that there is no evidence of a failure of monotonicity. In all cases, there is a clear positive relationship between the overall case-calendar-filing-year foreclosure rate and the group-specific case-calendar-filing-year foreclosure rate—a higher value of the instrument is associated with a higher foreclosure rate in each subgroup. Thus, in terms of observables—property type, loan type, whether the plaintiff is a large bank or employs a large attorney, and census tract demographics—the monotonicity assumption appears valid.

1.8.4 Nonlinearities in Foreclosure Contagion

Our regressions are all at the foreclosure-case level, which raises two issues. Firstly, the treatment is imperfectly assigned since neighborhoods around two (or more) delinquent properties may overlap. Secondly, and relatedly, it is difficult to investigate non-linearities in the effect of a completed foreclosure. To explore nonlinearities and work with a cleaner (albeit limited) treatment, we look at the effect of the lagged foreclosures in a small neighborhood on the new filings.

We define small neighborhoods by partitioning Cook County into squares and examining foreclosure behavior within. We assign each property associated with our foreclosure cases to a 0.0625-square-mile square (0.25x0.25). Within each square, i , denote by N_{it} the count of new filings that occur in each year, t . Similarly, we count the total number of completed foreclosures in each year, F_{it} . We discard squares that have no foreclosure activity (i.e., no filings) over the period 2004–2010.

Our goal is to relate the lagged number of completed foreclosures in square i to the number of new filings in a given year. We include fixed effects for the number of ongoing foreclosure cases for properties in square i as of year t (η_{it})⁵² and for the year of observation

⁵²Since our data only covers filings from 2004 through 2010, we only consider cases filed in this period.

(ψ_t) :

$$N_{it} = \beta_0 + \beta_1 F_{it-1} + \beta_2 F_{it-2} + \eta_{it-1} + \eta_{it-2} + \psi_t + \beta X_i + e_{it} \quad (10)$$

where X_i is a vector of square-specific controls (census-tract demographics).⁵³ We instrument the lagged number of completed foreclosures using the expected number of completed foreclosures in square i in year $t - 1$ or $t - 2$ (conditional on the number of ongoing cases, $\eta_{i,t-j}$), where we take the probability of a filing ending in completed foreclosure to be the leave-one-out case-calendar/filing-month probability of foreclosure as described in Section 1.4.2.⁵⁴ In this way, we are comparing neighborhoods with the same initial foreclosure filing activity, but with different completed foreclosure outcomes owing to random assignment of these filings to different case calendars.

Aggregating in this way helps address the two issues listed above. Firstly, neighborhoods here are well defined entities that do not overlap, and so there is no mis-assignment of the treatment as defined (while this is a contained definition of the treatment, completed foreclosures on the edges of these neighborhoods may nonetheless have spillover effects to neighboring squares that we are not accounting for). Secondly, by aggregating counts of foreclosures we can more readily explore non-linearities in the effect of filings.

The results from the 2SLS estimation of Equation 10 are presented in Table 1.34 and show evidence of contagion consistent with our baseline estimates from Section 1.5.1. We split the sample by pre-crash (2006–2008) and post-crash (2009–2010) observations to reflect the differences seen for these samples in Section 1.5.1. Column 1 shows a strong positive relationship between lagged completed foreclosures and new foreclosure filings—a completed foreclosure one year prior causes 0.281 new filings, while a foreclosure two-years prior causes

⁵³We have also experimented with additional lags, although these are rarely significant and require a further reduction in sample size.

⁵⁴We treat each case filed within square i as of the given period as an independent random draw that may or may not foreclose in the given year. We construct the probability of foreclosure in the given year for each case as the leave-one-out share of cases filed in the same month and assigned to the same calendar that foreclose in the given year. Then the expected number of foreclosures in square i in a given year is the sum over these probabilities.

0.913 new filings. Off of a mean of 2.54 filings per year for this sample, this represents an increase of 11.1% to 35.9%, which is on par with the baseline estimates in Section 1.5.1. At the same time, the linear estimates for 2009–2010, in Column 4, are substantially smaller: 0.115 and 0.481 (5.2% and 21.8% relative to the mean of 2.21 filings for this sample). Columns 2 and 5 show estimates allowing for a quadratic relationship between lagged foreclosures and new filings. Interestingly, the quadratic relationship is significant—as the lagged number of foreclosures grows, the marginal effect of a completed foreclosure in the neighborhood diminishes. Again, this is consistent with the findings in Sections 1.5.1 that split the sample by pre-crisis and crisis—exposure to more foreclosure activity diminishes the contagion effect, perhaps because the marginal foreclosure conveys less information (about house prices or the foreclosure process itself).

Finally, in Columns 3 and 6 we allow for a more flexible relationship between lagged foreclosures and new filings by regressing new filings on a set of indicators for the lagged number of foreclosures. We instrument the lagged foreclosure indicators by the probability of observing that number of foreclosures in that year, conditional on the number of ongoing foreclosure cases.⁵⁵ These estimates seem to suggest that the contagion effect is strongest for the first two completed foreclosures. However, we do not put too much stock in these results as the first-stage F statistic for these regressions suggests that the instrument here is weak.

1.8.5 Lender Response to a Completed Foreclosure

For any given mortgage, divide the remaining life of the loan into three periods ($t \in 0, 1, 2$). Consider a mortgage in which the borrower owes a payment of m_t in each period, t and the mortgage is fully paid back as of $t = 2$. Suppose the borrower misses their payment in the

⁵⁵For example, if there are four ongoing foreclosure cases in a given year, then the instrument for the indicator of one completed foreclosure is the sum of the probabilities that each of the four cases ends in foreclosure, where these probabilities are given by the case-calendar-filing-month foreclosure rate in that year. We restrict the sample to observations with no more than five ongoing cases in the lagged years, since the calculation of these probabilities grows drastically with the number of outstanding cases.

current period ($t = 0$). With probability α_1 the borrower will still be delinquent the following period and the lender can foreclose on the property, in which case the lender recovers the value of the home, P , less the costs of foreclosure, λ (e.g., legal fees). However, with probability $1 - \alpha_1$, the borrower will recover in period 1 and will resume making payments. Then the value of the unmodified loan to the lender is:

$$V_u = \alpha_1 \cdot (P - \lambda) + (1 - \alpha_1) \cdot \left[m_1 + \frac{m_2}{R} \right] \quad (11)$$

where R is the discount rate. The lender may instead choose to modify the loan, which reduces subsequent loan payments to $m'_1 < m_1$ and $m'_2 < m_2$.

Lenders will be willing to modify a mortgage when modification is very effective in reducing the probability of non-payment or when the necessary reduction in the value of the loan is small. By lowering payments, modification reduces the probability of default to $\alpha'_1 < \alpha_1$.⁵⁶ Then the value of the modified loan is:

$$V_M = \alpha'_1(P - \lambda) + (1 - \alpha'_1) \left[m'_1 + \frac{m'_2}{R} \right]$$

Thus, the lender will choose modification when the value of the modified loan is larger than the unmodified loan, or the difference between the two is positive:

$$V_M - V_u = (\alpha_1 - \alpha'_1) \cdot \left[m'_1 + \frac{1}{R}m'_2 - (P - \lambda) \right] - (1 - \alpha_1) \cdot \left[m_1 + \frac{1}{R}m_2 - (m'_1 + \frac{1}{R}m'_2) \right] > 0 \quad (12)$$

Since $\alpha_1 - \alpha'_1 > 0$, $1 - \alpha_1 > 0$, and $[m_1 + \frac{1}{R}m_2 - (m'_1 + \frac{1}{R}m'_2)] > 0$, even if the probability of re-default under modification is zero (i.e., even if $\alpha'_1 = 0$), modification is still not optimal if the modified payments are too low. For example, lenders will never modify if the net-

⁵⁶Adelino et al. (2009) allow home prices to change across periods, however, for the purpose of this paper little generality is lost by assuming no price growth. Similarly, Adelino et al. (2009) operate under the assumption that modification guarantees payment in period 1 with a non-zero probability of default in period 2. Again, this does not fundamentally change the implications of the framework.

of-foreclosure-cost value of the property is greater than the present value of the modified mortgage payments (i.e., $m'_1 + \frac{1}{R}m'_2 < (P - \lambda)$). When the value of the modified payments are high enough, lenders are more inclined to modify when it is very effective in reducing the probability of redefault (i.e., the smaller is α'_1).

Assuming that the default probabilities are constant (which shuts down the borrower contagion channel), a neighboring completed home foreclosure—which lowers the value of the property under consideration—discourages new foreclosure filings by lenders. Taking the derivative of $V_M - V_u$ with respect to the value of the home, we see that $\frac{\partial V_M - V_u}{\partial P} = \alpha'_1 - \alpha_1 < 0$, relying on the assumption that modification lowers the probability of future default. A drop in the value of the home encourages mortgage modification—selling the property at auction is relatively less appealing to the lender than the modified mortgage.

We derive similar conditions for modification by mortgage servicers. Mortgage servicers are typically employed by lenders to collect mortgage payments and to manage mortgage defaults. When a mortgage is current, servicers receive a share of the interest payments that they collect. However, when a borrower is delinquent, servicers are required to forward payments to the holder of the debt while they manage the default (either by modifying the loan or seeking foreclosure). While managing a default, the servicer must incur all associated costs (e.g. legal fees). If the default ends in foreclosure, the servicer is reimbursed for all foreclosure-related expenses and fees. Thus, the value to the servicer of foreclosing on the delinquent loan (as in the above framework) is:

$$V_u^s = \alpha_1 \cdot \Pi + (1 - \alpha_1) \cdot \left[\rho \cdot m_1 + \rho \cdot \frac{m_2}{R} \right] \quad (13)$$

where Π is the total value of all foreclosure-related fees charged by the servicer and where ρ is the share of the mortgage payment that is returned to the servicer. When the loan is modified, the servicer is generally not reimbursed for any related expenses, and so the value of the modified loan to the servicer is:

$$V_M^s = \alpha'_1 \cdot \Pi + (1 - \alpha'_1) \left[\rho \cdot m'_1 + \rho \cdot \frac{m'_2}{R} \right] - C_M$$

where C_M is the cost of modification (e.g., time/labor spent in negotiations). Notice that if the modification is successful, the servicer receives a lower monthly payment for servicing the mortgage. Then the servicer will prefer to modify when:

$$V_M^s - V_u^s = (\alpha'_1 - \alpha_1) \cdot \left[\rho m'_1 + \frac{\rho}{R} m'_2 - \Pi \right] - (1 - \alpha_1) \cdot \rho \left[m_1 + \frac{1}{R} m_2 - \left(m'_1 + \frac{1}{R} m'_2 \right) \right] - C_m > 0$$

The servicer's incentives are similar to the lender's incentives, although as pointed out by Levitin and Twomey (2011), since servicers' fees have seniority over all other claims against a property, servicers are "indifferent to the amount of the [foreclosure] sale proceeds." In other words, since the payoff to a servicer of a completed foreclosure is the foreclosure fees, the servicer does not care about the value of the property (as long as it is high enough to cover their fees). Moreover, high costs of modification (which are not reimbursed) will push the servicer in favor of foreclosure. Notice, then, that if the probabilities of foreclosure (α_1 and α'_1) are invariant to a neighboring completed foreclosure, the servicer will not experience foreclosure contagion ($\partial (V_M^s - V_u^s) / \partial P = 0$).

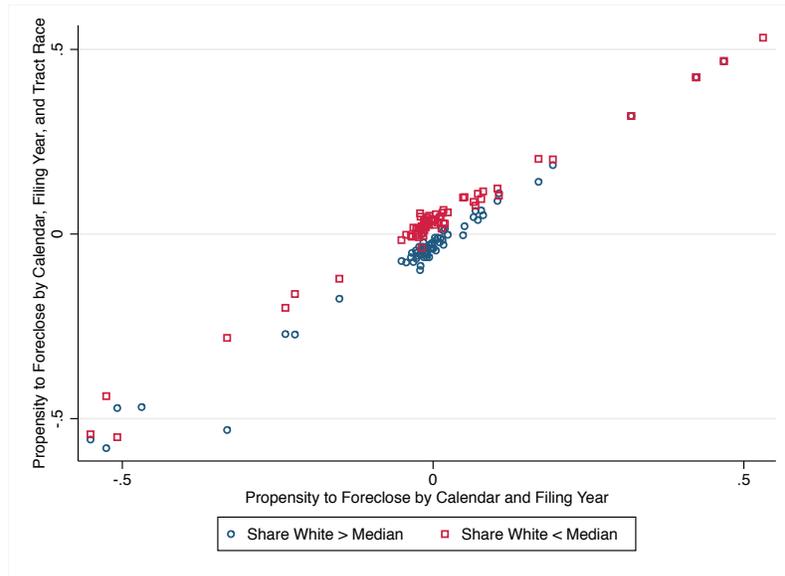
In summary, under the assumptions that a neighboring completed home foreclosure lowers the housing value P and that borrowers are unresponsive to the neighboring foreclosure so that α_1 and α'_1 are unchanged, we should observe "anti-contagion" in home foreclosures: if only lenders and servicers are responding, a completed home foreclosure should discourage neighboring foreclosure filings. Of course, if a neighboring completed foreclosure has no influence on housing values and borrowers are unresponsive, then lenders and servicers should not respond (we should see no contagion at all).

1.8.6 Estimates by Proxies for Social Connectedness

We find little systematic relationship between the extent of foreclosure contagion and several proxies for social connection. Our first attempt to proxy social connectedness is to stratify by neighborhood diversity, operating under the assumption that neighbors with similar background maintain stronger social ties. For each census tract, we calculate a Herfindahl index of neighborhood diversity using race-population shares from the 2000 Decennial Census. We then estimate the baseline treatment effect for each diversity quartile to examine whether contagion is stronger in less diverse (high-index) census tracts. We also estimate the baseline specification and interact the treatment effect with an indicator for the foreclosure taking place in a census tract where a single race makes up more than 75% of the population. Our second set of proxies draw on the notion of social capital outlined by Glaeser and Sacerdote (2000), who argue that social connections are higher when residents live in close proximity to one another. For example, studying survey data, Glaeser and Sacerdote (2000) find that there are higher levels of civic participation among residents of large condo buildings than single-family housing. We proxy neighbor proximity in two ways. First, we interact the foreclosure treatment effect with an indicator that the home undergoing foreclosure (the unit of observation) is a condominium unit. Second, we stratify census tracts by population density (again using data from the 2000 Decennial Census). In both of these cases, while we continue to find evidence of contagion, there is no systematic relationship between foreclosure contagion and our proxies for social connection. Our social-contagion estimates can be found in Table 1.31 and 1.32, and Figures 1.5 and 1.6 of the Appendix.

Figure 1.2: Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates

(a) Predominantly White vs. Non-White Tract



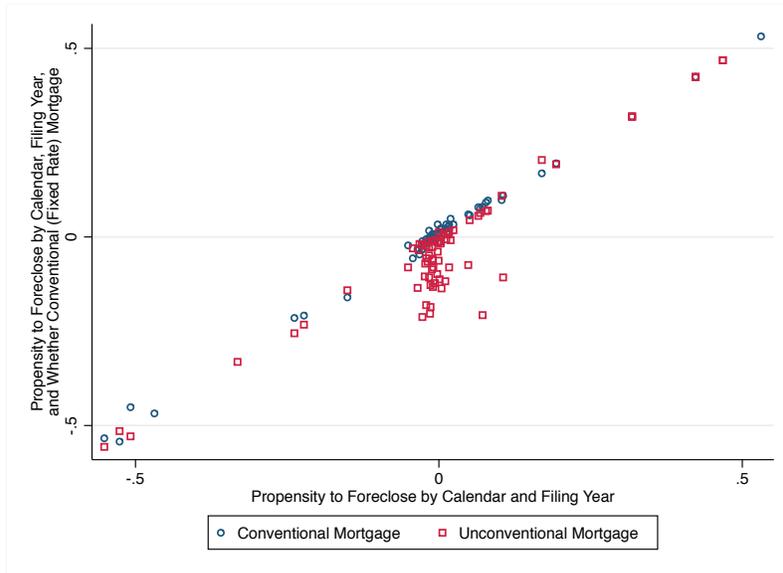
(b) Census-tract Income Quartile



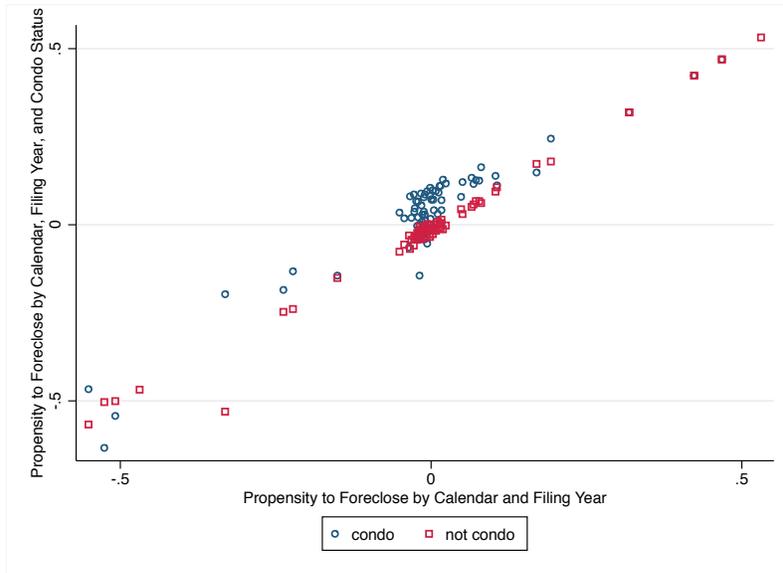
Notes: Filing-year foreclosure rates are calculated for each indicated sub-group by case calendar, demeaned by filing-year, and plotted against overall foreclosure rates for the given calendar and filing year (again, demeaned by filing year). A predominantly white census tract has an above-median share of white residents as of the 2000 Decennial Census, and income quartiles are calculated at the census tract level using median income from the 2000 Decennial Census.

Figure 1.3: Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates

(a) Conventional vs. Other Mortgages



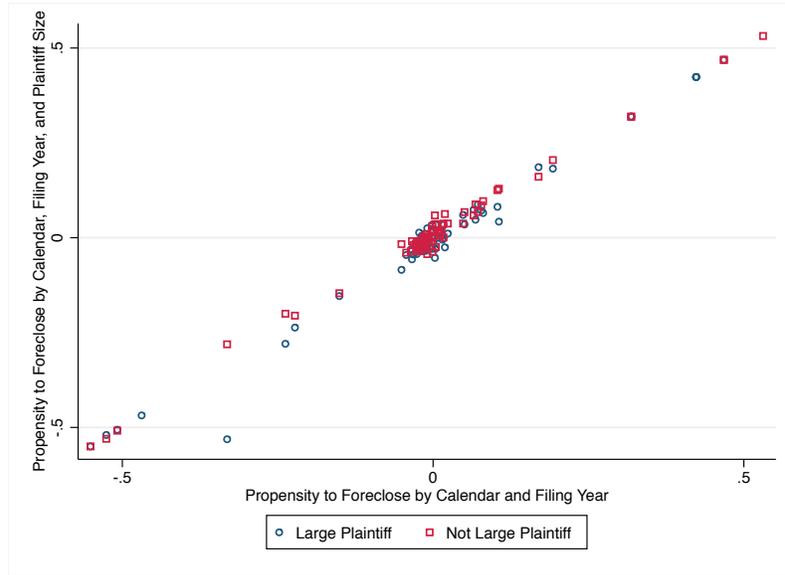
(b) Condominiums vs. Non-Condominiums



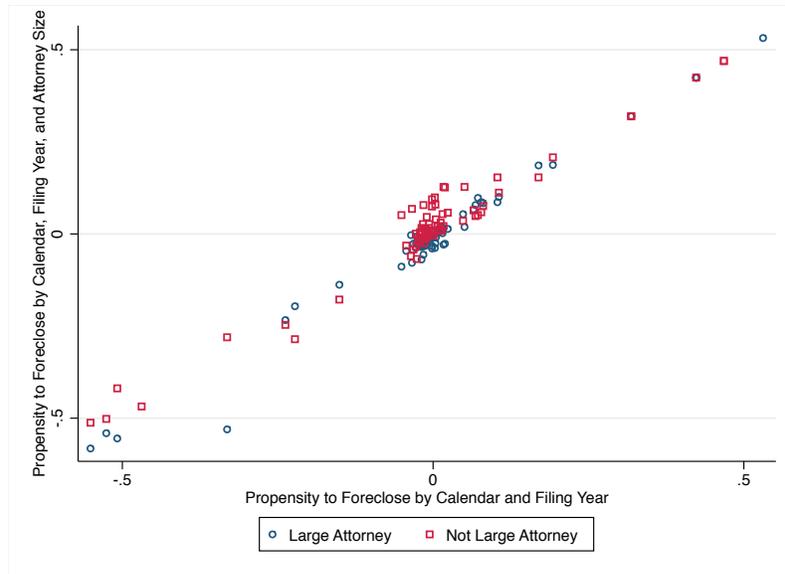
Notes: Filing-year foreclosure rates are calculated for each indicated sub-group by case calendar, demeaned by filing-year, and plotted against overall foreclosure rates for the given calendar and filing year (again, demeaned by filing year). Conventional mortgage includes all standard fixed-rate mortgages (while unconventional mortgages includes adjustable-rate mortgages, balloon-payment mortgages, reverse mortgages, and interest-only mortgages; we exclude VA mortgages). Condo status is as reported in the court documents.

Figure 1.4: Calendar-Group-Specific Foreclosure Rates vs. Calendar-Specific Rates

(a) Large Lender vs. Other Lender



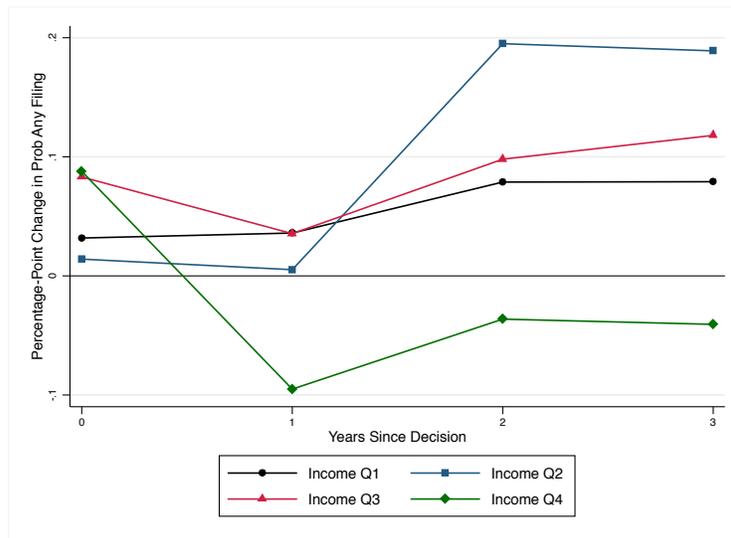
(b) Large Attorney vs. Other Attorneys



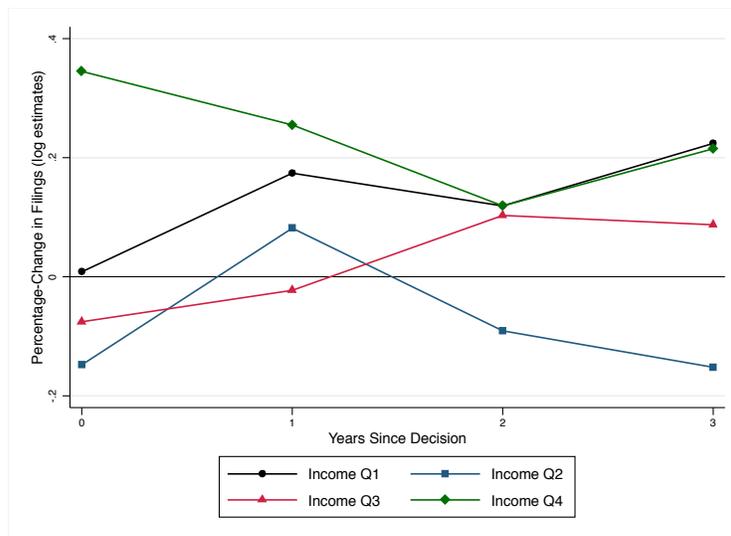
Notes: Filing-year foreclosure rates are calculated for each indicated sub-group by case calendar, demeaned by filing-year, and plotted against overall foreclosure rates for the given calendar and filing year (again, demeaned by filing year). A large lender is a plaintiff who appears on more than 7000 of the foreclosure cases filed in cook county, while a large attorney appears on greater than 10,000 cases.

Figure 1.5: Contagion Estimates by Income Quartile

(a) Effect of Foreclosure on Any New Foreclosure Filings



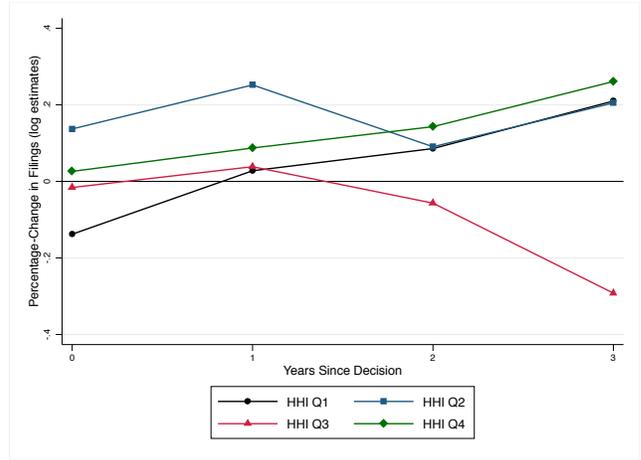
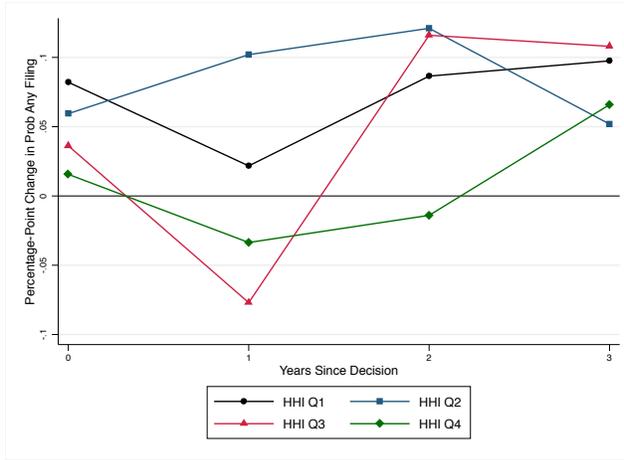
(b) Effect of Foreclosure on Log New Foreclosure Filings



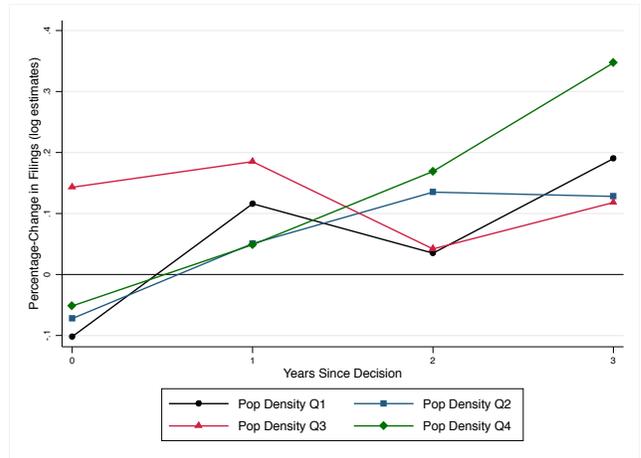
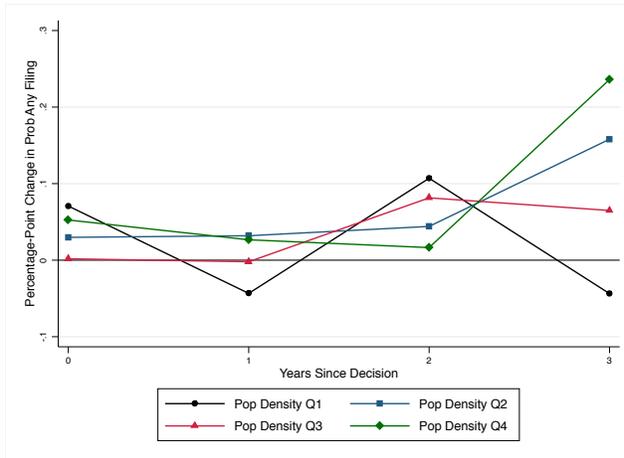
Notes: 2SLS estimates (as described in Tables 1.9 and 1.5) performed separately for each value of the given neighborhood quartile of income: income quartiles are calculated at the census tract level using median income from the 2000 Decennial Census.

Figure 1.6: Contagion Estimates by Social Connection Proxy

(a) Effect of Foreclosure on Any New Foreclosure Filings by Census Tract Diversity
 (b) Effect of Foreclosure on Log New Foreclosure Filings by Census Tract Diversity



(c) Effect of Foreclosure on Any New Foreclosure Filings by Census Tract Density
 (d) Effect of Foreclosure on Log New Foreclosure Filings by Census Tract Density



Notes: 2SLS estimates (as described in Tables 1.9 and 1.5) performed separately for each value of the given quartile. Population density is calculated as census-tract population (as of 2000) over census-tract area. Diversity is measured with a Herfindahl-Hirschman index over the shares of each 2000 Decennial Census-designated race in the tract.

Table 1.13: First Stage Regression of Foreclosure on Propensity to Foreclose

	No Property Controls	Controls [†]
Coefficient on Z_i	0.556*** (0.045)	0.554*** (0.046)
1st-Stage F	150.000	147.400
N	140667	140667

Notes: ***Indicates significance at the 1% level. Reported standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. First-stage regression of indicator for case ending in foreclosure on leave-one-out case-calendar-filing-month foreclosure rate, and filing month and property type fixed effects. [†]Unreported controls include share of tract that report race as white in 2000 decennial census, income quartile from decennial census, whether plaintiff is a “large plaintiff” (six largest plaintiffs each representing ≥ 7000 filings) or attorney is a “large attorney” (three largest attorneys each representing $\geq 10,000$ cases), whether mortgage is adjustable rate, size of initial loan, and census tract population.

Table 1.14: Contagion Estimates Measured Since Case Filing

Years Since Filing	0	1	2	3	4	5	
Any Filing	2SLS	0.057*** (0.021)	0.032 (0.021)	0.026 (0.022)	0.025 (0.026)	0.076 (0.076)	0.024 (0.087)
	1st-stage F	388.100	388.100	399	255.500	30.450	22.310
	N	140683	140683	121171	98032	59967	36685
	2SLS	0.631* (0.354)	0.867** (0.343)	0.334 (0.260)	0.276* (0.152)	0.023 (0.440)	0.456 (0.728)
Total Filings	1st-stage F	388.100	388.100	399	255.500	30.450	22.310
	N	140683	140683	121171	98032	59967	36685

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure outcome are as in Table 1.5. Outcomes (and year-of-observation fixed effects) are measured as of the year that the case is filed (rather than the year that the case is decided).

Table 1.15: Controlling for Length of Case

Years Since Decision		0	1	2	3	4	5
Control for Log Length of Case	Any	0.049	-0.006	0.107**	0.125	0.357**	0.167
	Filing	(0.042)	(0.040)	(0.043)	(0.086)	(0.177)	(0.163)
	First-stage F	149.500	133.900	114.600	24.810	10.740	8.804
	<i>N</i>	129834	118201	92804	67076	41733	23666
	Total Filings	0.728	0.782	0.669**	0.928*	2.308	0.627
	Filing	(0.553)	(0.562)	(0.299)	(0.511)	(1.522)	(1.196)
	First-stage F	149.500	133.900	114.600	24.810	10.740	8.804
	<i>N</i>	129834	118201	92804	67076	41733	23666
	log(Price)	-0.102	-0.516	-0.390			
	Filing	(0.176)	(0.341)	(0.442)			
	First-stage F	19.320	7.391	12.570			
	<i>N</i>	42880	25890	12155			
Control for Quadratic in Length of Case	Any	0.058	0.008	0.105***	0.124	0.401**	0.245
	Filing	(0.036)	(0.036)	(0.037)	(0.076)	(0.196)	(0.264)
	First-stage F	175.100	163.900	163.200	27.100	14.840	4.805
	<i>N</i>	130199	118566	93143	67379	41958	23831
	Total Filings	0.783	0.791	0.677***	0.890**	2.497	0.883
	Filing	(0.496)	(0.501)	(0.255)	(0.451)	(1.599)	(1.895)
	First-stage F	175.100	163.900	163.200	27.100	14.840	4.805
	<i>N</i>	130199	118566	93143	67379	41958	23831
	log(Price)	-0.113	-0.638	-0.511			
	Filing	(0.158)	(0.427)	(0.685)			
	First-stage F	21.270	6.922	7.338			
	<i>N</i>	43079	26047	12241			
Control for Quarterly Length Dummies	Any	0.053	0.001	0.106**	0.141	0.463**	0.214
	Filing	(0.042)	(0.040)	(0.041)	(0.096)	(0.226)	(0.225)
	First-stage F	127	123.400	104.800	24.880	14.800	8.341
	<i>N</i>	130199	118566	93143	67379	41958	23831
	log(# of Filings)	0.808	0.842	0.672**	0.990*	2.819	0.803
	Filing	(0.567)	(0.556)	(0.288)	(0.576)	(1.956)	(1.660)
	First-stage F	127	123.400	104.800	24.880	14.800	8.341
	<i>N</i>	130199	118566	93143	67379	41958	23831
	log(Price)	-0.105	-0.619	-0.462			
	Filing	(0.197)	(0.404)	(0.562)			
	First-stage F	19.450	9.194	12.400			
	<i>N</i>	43079	26047	12241			

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure outcome are as in Table 1.5, including the indicated controls for length of the case—log of the total months, a quadratic in total months, or number-of-quarter fixed effects.

Table 1.16: Robustness of Contagion Estimates By Specification

Years Since Decision	0	1	2	3	4	5	
No Covariates	Any	0.066**	0.029	0.096***	0.113**	0.253**	0.160
	Filing	(0.028)	(0.027)	(0.027)	(0.052)	(0.107)	(0.130)
	First-stage F	242.600	225	212.400	25.470	20.900	14.130
	N	130199	118566	93143	67379	41958	23831
	Total Filings	0.739*	0.741**	0.603***	0.796**	1.616*	0.621
	Filing	(0.388)	(0.375)	(0.196)	(0.320)	(0.954)	(0.972)
	First-stage F	242.600	225	212.400	25.470	20.900	14.130
	N	130199	118566	93143	67379	41958	23831
No Property-Type FEs	Any	0.071**	0.029	0.095***	0.096*	0.244**	0.150
	Filing	(0.028)	(0.027)	(0.027)	(0.053)	(0.112)	(0.133)
	First-stage F	232.700	220.200	201	24.080	18.910	13.190
	N	130199	118566	93143	67379	41958	23831
	Total Filings	1.190***	1.065***	0.876***	0.868**	1.488	0.805
	Filing	(0.453)	(0.398)	(0.197)	(0.338)	(0.939)	(0.941)
	First-stage F	232.700	220.200	201	24.080	18.910	13.190
	N	130199	118566	93143	67379	41958	23831
Drop Cases Decided in Summer	Any	0.068**	0.018	0.104***	0.089	0.171	0.085
	Filing	(0.034)	(0.032)	(0.034)	(0.058)	(0.117)	(0.137)
	First-stage F	179.200	168.800	162.300	25.500	19.180	14.600
	N	105638	95487	75004	55342	34719	19464
	Total Filings	0.914**	0.851**	0.633***	0.563	1.126	0.100
	Filing	(0.430)	(0.406)	(0.230)	(0.354)	(0.932)	(1.078)
	First-stage F	179.200	168.800	162.300	25.500	19.180	14.600
	N	105638	95487	75004	55342	34719	19464
Monotonicity-Robust IV	Any	0.083	0.118**	0.026	-0.088	-0.130	-0.417
	Filing	(0.062)	(0.057)	(0.062)	(0.109)	(0.296)	(0.662)
	First-stage F	45.180	43.800	37.200	14.260	4.733	1.432
	N	127521	115977	90795	65515	40534	22821
	Total Filings	3.471***	2.067***	1.041	-0.263	-1.259	-7.345
	Filing	(1.113)	(0.675)	(0.674)	(1.149)	(3.539)	(9.501)
	First-stage F	45.180	43.800	37.200	14.260	4.733	1.432
	N	127521	115977	90795	65515	40534	22821

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given outcome are as in Table 1.5, with the following adjustments: “No covariates” omits the case-level controls as outlined in Table 1.5; “No Property-Type FEs” omits the property-type (condo/single-family/multi-family/apartment) fixed effects; “Drop Summer Decisions” drops cases that are decided in June, July, or August (during which the courts dismiss inactive cases); “Monotonicity-Robust IV” indicates that the instrument is constructed by calendar/filing month/income quartile/large plaintiff/“white” census tract cells.

Table 1.17: Robustness of Estimates by Sample

Years Since Decision		0	1	2	3	4	5
Include Moratorium	Any	0.056**	0.009	0.083***	0.093*	0.280**	0.160
	Filing	(0.028)	(0.028)	(0.027)	(0.054)	(0.115)	(0.137)
	First-stage F	188.700	179.600	202.400	24.060	18.630	12.440
	N	140912	127802	94495	67749	42249	24048
	Total Filings	0.649	0.649*	0.553***	0.700**	1.757*	0.693
	Filing	(0.399)	(0.375)	(0.184)	(0.318)	(1.019)	(1.013)
	First-stage F	188.700	179.600	202.400	24.060	18.630	12.440
	N	140912	127802	94495	67749	42249	24048
	log(Price)	-0.122	-0.418**	-0.378			
	Filing	(0.106)	(0.210)	(0.409)			
	First-stage F	21.490	18.220	14.660			
	N	43343	26236	12343			
Full Sample	Any	0.078***	0.064***	0.071***	0.058***	0.035	0.017
	Filing	(0.022)	(0.019)	(0.018)	(0.021)	(0.025)	(0.022)
	First-stage F	202.200	201.500	225.100	144.800	125	127.600
	N	152559	139324	105628	77785	50704	31508
	Total Filings	0.402**	0.363**	0.410***	0.376**	0.360**	0.241
	Filing	(0.204)	(0.177)	(0.124)	(0.163)	(0.157)	(0.204)
	First-stage F	202.200	201.500	225.100	144.800	125	127.600
	N	152559	139324	105628	77785	50704	31508
	log(Price)	-0.159***	-0.199***	-0.222***			
	Filing	(0.038)	(0.036)	(0.036)			
	First-stage F	131.300	121.900	132.700			
	N	50879	32750	17952			
Filings < 99th Percentile	Any	0.048*	0.008	0.081***	0.089	0.251**	0.140
	Filing	(0.028)	(0.027)	(0.028)	(0.055)	(0.115)	(0.135)
	First-stage F	247.600	229.100	207.200	24.470	18.120	12.920
	N	127046	116434	91674	66454	41416	23533
	Total Filings	0.144	0.234	0.395***	0.318	0.887	0.231
	Filing	(0.194)	(0.174)	(0.128)	(0.225)	(0.561)	(0.568)
	First-stage F	247.600	229.100	207.200	24.470	18.120	12.920
	N	127046	116434	91674	66454	41416	23533

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given outcome are as in Table 1.5, with the following sample adjustments (recall that the baseline sample omits cases filed during the foreclosure moratorium of 2009, omits filings against VA loans, and omits cases filed during the first year of a case-calendar's existence): "Include Moratorium" maintains the baseline sample, but includes cases filed during the 2009 foreclosure moratorium; "Full Sample" relies on the set of all matched foreclosure cases in Cook County for 2004–2010; "Filings < 99th Percentile" drops observations where the number of new foreclosure filings near a case in a given year is greater than the 99th percentile.

Table 1.18: Contagion in Any New Filing, Omitting Each Filing Year

Years Since Decision	0	1	2	3	4	5
No 2004	0.044 (0.028)	0.014 (0.027)	0.078*** (0.027)	0.076 (0.054)	0.233* (0.130)	0.167 (0.174)
First-stage F	221.700	208.300	190.800	21.330	15.180	10.010
<i>N</i>	124342	112717	87309	61592	36265	18391
No 2005	0.046* (0.028)	0.009 (0.027)	0.072** (0.028)	0.109* (0.063)	0.349** (0.171)	-0.735 (1.183)
First-stage F	258.600	248	247.400	22.270	8.089	0.937
<i>N</i>	116682	105064	79676	54020	28854	11860
No 2006	0.053* (0.027)	0.006 (0.026)	0.078*** (0.027)	0.060 (0.047)	0.189* (0.103)	0.233* (0.132)
First-stage F	267.700	255	249	28.580	26.780	13.700
<i>N</i>	112963	101380	76056	50592	26403	17411
No 2007	0.061** (0.029)	0.015 (0.028)	0.073*** (0.028)	0.072 (0.052)	0.316* (0.185)	0.140 (0.133)
First-stage F	176.900	167.300	172.700	21.770	8.266	12.970
<i>N</i>	107066	95565	70614	46462	34353	23831
No 2008	0.079 (0.073)	0.003 (0.072)	0.224*** (0.081)	0.242* (0.126)	0.247** (0.112)	0.140 (0.133)
First-stage F	62.950	55.030	34.530	16.530	19.060	12.970
<i>N</i>	93508	83050	62869	56850	41957	23831
No 2009	0.055* (0.028)	0.018 (0.028)	0.078*** (0.029)	0.090* (0.053)	0.247** (0.112)	0.140 (0.133)
First-stage F	232.200	218	173.100	24.620	19.060	12.970
<i>N</i>	110228	101633	89191	67379	41958	23831
No 2010	0.043 (0.028)	0.012 (0.027)	0.082*** (0.027)	0.090* (0.053)	0.247** (0.112)	0.140 (0.133)
First-stage F	245.500	224.400	205	24.620	19.060	12.970
<i>N</i>	116405	111987	93143	67379	41958	23831

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for any new neighboring filing are as in Table 1.5, although omitting cases filed in the given year.

Table 1.19: Contagion in Total New Filings, Omitting Each Filing Year

Years Since Decision	0	1	2	3	4	5
No 2004	0.658 (0.403)	0.642* (0.376)	0.462** (0.182)	0.531* (0.318)	1.464 (1.181)	-0.253 (0.253)
First-stage F	221.700	208.300	190.800	21.330	15.180	20.300
<i>N</i>	124342	112717	87309	61592	36265	13557
No 2005	0.667 (0.407)	0.653* (0.386)	0.531*** (0.189)	0.840** (0.341)	1.609 (1.157)	0.571 (1.206)
First-stage F	258.600	248	247.400	22.270	8.089	1.423
<i>N</i>	116682	105064	79676	54020	28854	8716
No 2006	0.695* (0.396)	0.650* (0.373)	0.501*** (0.175)	0.597** (0.286)	1.524 (0.946)	-0.003 (0.231)
First-stage F	267.700	255	249	28.580	26.780	24.660
<i>N</i>	112963	101380	76056	50592	26403	12983
No 2007	0.659 (0.427)	0.686* (0.405)	0.528*** (0.187)	0.529* (0.314)	1.882 (1.594)	-0.062 (0.217)
First-stage F	176.900	167.300	172.700	21.770	8.266	24.500
<i>N</i>	107066	95565	70614	46462	34353	17628
No 2008	1.118* (0.654)	0.635 (0.517)	1.192** (0.573)	1.178 (1.012)	1.551 (0.983)	-0.062 (0.217)
First-stage F	62.950	55.030	34.530	16.530	19.060	24.500
<i>N</i>	93508	83050	62869	56850	41957	17628
No 2009	0.749* (0.394)	0.768** (0.378)	0.524*** (0.199)	0.657** (0.319)	1.551 (0.983)	-0.062 (0.217)
First-stage F	232.200	218	173.100	24.620	19.060	24.500
<i>N</i>	110228	101633	89191	67379	41958	17628
No 2010	0.627* (0.375)	0.667* (0.367)	0.536*** (0.183)	0.657** (0.319)	1.551 (0.983)	-0.062 (0.217)
First-stage F	245.500	224.400	205	24.620	19.060	24.500
<i>N</i>	116405	111987	93143	67379	41958	17628

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for count of filings per year are as in Table 1.5, although omitting cases filed in the given year.

Table 1.20: Baseline Estimates for 0.25 Mile Radius

Years Since Decision		0	1	2	3	4	5
Any Filing per Year	2SLS	0.017 (0.012)	-0.003 (0.011)	0.016 (0.011)	0.026 (0.016)	0.032 (0.041)	-0.064 (0.052)
	1st-stage F	238.300	224.200	205	24.620	19.060	12.970
	N	130199	118566	93143	67379	41958	23831
	2SLS	2.187*** (0.777)	2.358 *** (0.785)	1.693 *** (0.520)	1.219 *** (0.764)	2.665 (1.792)	0.762 (2.217)
Total Filings per Year	1st-stage F	238.300	224.200	205	24.620	19.060	12.970
	N	130199	118566	93143	67379	41958	23831
	2SLS	-0.161** (0.067)	-0.269 * (0.144)	-0.410 * (0.220)			
	1st-stage F	24.760	17.870	16.860			
log(price)	N	63639	39508	21860			

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for any new neighboring filing are as in Table 1.5, although defining outcomes using a 0.25 mile radius around delinquent property.

Table 1.21: Estimates for Loans with Multiple Claimants

Years Since Decision	0	1	2	3	4	5
Any Filing	0.076*	0.007	0.087**	0.071	0.112	0.058
	(0.039)	(0.038)	(0.039)	(0.075)	(0.143)	(0.164)
First-stage F	210.100	213.600	175.800	23.360	16.810	12.410
<i>N</i>	65366	59917	48271	34865	21252	11377
95% CI Lower Bound	0.154	0.081	0.163	0.217	0.392	0.379
Total Filings per Year	0.820	0.669	0.454	0.470	2.026	0.707
	(0.575)	(0.511)	(0.286)	(0.495)	(1.634)	(1.267)
First-stage F	210.100	213.600	175.800	23.360	16.810	12.410
<i>N</i>	65366	59917	48271	34865	21252	11377
95% CI Upper Bound	1.947	1.671	1.015	1.440	5.229	3.190
log(price)	0.035	-0.172	-0.198			
	(0.147)	(0.231)	(0.346)			
First-stage F	19.580	24.230	18.530			
<i>N</i>	22348	13423	6071			
95% CI Upper Bound	-0.254	-0.625	-0.876			

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for any new neighboring filing are as in Table 1.5. Sample is restricted to cases where multiple lien holders are listed in the foreclosure filing.

Table 1.22: Contagion in Completed Foreclosures/Dismissals Filed After Decision

Years After End of Case	0	1	2	3	4	5
Contagion in Completed Foreclosures Filed After Decision						
Any Completed Foreclosure	-0.015 (0.011)	-0.150*** (0.032)	0.054 (0.035)	0.186*** (0.055)	-0.023 (0.115)	0.119 (0.154)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
Total Completed Foreclosures	-0.028* (0.014)	-0.285** (0.129)	0.323** (0.160)	0.705*** (0.211)	-0.044 (0.389)	0.180 (0.539)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
Contagion in Dismissals Filed After Decision						
Any Dismissal	-0.018** (0.008)	-0.029 (0.023)	-0.011 (0.017)	0.034 (0.030)	0.047 (0.072)	0.077 (0.095)
First-stage F	238.300	224.200	205	24.620	19.060	12.970
<i>N</i>	130199	118566	93143	67379	41958	23831
Total Dismissals	-0.026** (0.011)	-0.050 (0.036)	-0.026 (0.028)	0.049 (0.059)	0.003 (0.127)	0.216 (0.176)
First-stage F	232.700	234.500	172.700	25.030	15.020	8.676
<i>N</i>	94866	84906	66291	46383	28310	15230

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure or dismissal outcome are as in Table 1.5. Dismissal outcomes are counted in the same way as foreclosure outcomes: number of dismissals or indicator for any dismissal (among cases filed after observation case is decided).

Table 1.23: Robustness of Price Estimates By Specification

Years Since Decision		0	1	2
No Covariates	log(Price)	-0.153 (0.113)	-0.439** (0.214)	-0.374 (0.394)
	First-stage F	23.290	20.600	15.210
	<i>N</i>	43079	26047	12241
No Property-Type FEs	log(Price)	-0.113 (0.106)	-0.416** (0.204)	-0.361 (0.397)
	First-stage F	22.140	19.140	14.580
	<i>N</i>	43079	26047	12241
Drop Summer Decisions	log(Price)	-0.127 (0.122)	-0.482** (0.240)	-0.549 (0.361)
	First-stage F	22.860	18.330	9.793
	<i>N</i>	31977	19176	8837
Monotonicity-Robust IV	log(Price)	-0.200 (0.256)	-0.929* (0.502)	-0.337 (0.678)
	First-stage F	12.080	5.523	3.730
	<i>N</i>	41593	24902	11517

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given outcome are as in Table 1.5, with the following adjustments: “No covariates” omits the case-level controls as outlined in Table 1.5; “No Property-Type FEs” omits the property-type (condo/single-family/multi-family/apartment) fixed effects; “Drop Summer Decisions” drops cases that are decided in June, July, or August (during which the courts dismiss inactive cases); “Monotonicity-Robust IV” indicates that the instrument is constructed by calendar/filing month/income quartile/large plaintiff/“white” census tract cells.

Table 1.24: Price Effects, Omitting Each Filing Year

Years Since Decision	0	1	2
No 2004	-0.068 (0.111)	-0.418* (0.252)	-0.207 (0.464)
First-stage F	17.390	12.490	12.150
<i>N</i>	38223	21575	8701
No 2005	-0.140 (0.127)	-0.282 (0.405)	-2.818 (2.370)
First-stage F	15.530	5.944	1.824
<i>N</i>	32114	16558	6011
No 2006	-0.143 (0.101)	-0.393** (0.182)	-0.156 (0.341)
First-stage F	28.770	29.790	13.780
<i>N</i>	30601	17345	9770
No 2007	-0.113 (0.112)	-0.531* (0.292)	-0.358 (0.395)
First-stage F	17.230	9.746	14.390
<i>N</i>	32369	22664	12241
No 2008	-0.206 (0.206)	-0.411** (0.203)	-0.358 (0.395)
First-stage F	13.530	18.830	14.390
<i>N</i>	39009	26046	12241

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for log of neighboring residential real estate sale prices are as in Table 1.5, although omitting cases filed in the given year.

Table 1.25: Sub-Sample Price Effects

Years Since Decision	0	1	2
Repeat Sample, Unadjusted	-0.073 (0.115)	-0.395 (0.247)	-0.322 (0.440)
log(price) 1st-stage <i>F</i>	22.070	11.470	13.980
<i>N</i>	30482	17916	7904
Non-Repeat Sample	-0.235 (0.146)	-0.291 (0.274)	-0.705 (0.518)
log(price) 1st-stage <i>F</i>	24.040	16.870	9.938
<i>N</i>	30522	17925	8143

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates for given completed foreclosure or dismissal outcome are as in Table 1.5, although using the given sample. “Repeat sample” includes all properties for which we observe at least one other sale; “non-repeat” is all other sales.

Table 1.26: Contagion Among Loans with Lenders Implicated in Independent Foreclosure Review Settlement

	Years Since Decision					
	0	1	2	3	4	5
Any Filing per Year	0.024 (0.030)	0.038 (0.030)	0.068 (0.051)	0.091 (0.111)	0.059 (0.153)	0.009 (0.680)
First-stage F	269.300	161.600	25.060	19.060	12.970	2.765
N	94735	83564	67028	41958	23831	9579
Total Filings per Year	0.121 (0.103)	0.115 (0.091)	0.139 (0.086)	0.033 (0.224)	0.226 (0.283)	0.357 (1.240)
First-stage F	269.300	161.600	25.060	19.060	12.970	2.765
N	94735	83564	67028	41958	23831	9579
Log of Filings per Year	0.094 (0.068)	0.089 (0.068)	0.064 (0.084)	0.034 (0.281)	0.174 (0.225)	0.024 (0.396)
First-stage F	312.200	139.700	18.160	8.790	10.440	5.390
N	27570	23426	17780	11119	6159	2551
Log of Cumulative Filings since Decision	0.083 (0.079)	0.155** (0.070)	0.258*** (0.083)	0.252 (0.228)	0.680** (0.319)	1.911 (1.454)
First-stage F	221.700	167.300	22.240	20.410	15.760	3.238
N	27929	36395	34539	23998	14714	6317

Notes: **Indicates significance at the 1% level, *5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS specification as in Table 1.5. Outcomes for “Automating Lenders” are measured based on new foreclosure filings by lenders implicated in the Office of the Comptroller of the Currency and the Board of Governors of the Federal Reserve System Independent Foreclosure Review Settlement, while “Baseline” includes new foreclosure filings among all lenders. Both samples are restricted to cases decided between 2007 and 2010 (during which automation of foreclosure filings is thought to be most common).

Table 1.27: Contagion among Borrowers in Positive vs. Negative Equity

	Years Since Decision					
	0	1	2	3	4	5
Any Filing From Non-Underwater Borrower	0.044 (0.028)	0.037 (0.028)	0.105*** (0.029)	0.131** (0.056)	0.175 (0.111)	0.191 (0.145)
Any Filing From Underwater Borrower	0.043 (0.030)	-0.003 (0.029)	0.010 (0.027)	0.073 (0.052)	0.189 (0.153)	-0.251 (0.155)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.985	0.187	0.000	0.286	0.917	0.010
<i>N</i>	260398	237132	186286	134758	83916	47662
Total Filings From Non-Underwater Borrowers	0.785** (0.352)	0.736** (0.319)	0.619*** (0.179)	0.650** (0.268)	0.921 (0.748)	0.836 (0.822)
Total Filings From Underwater Borrowers	-0.094 (0.089)	-0.066 (0.084)	-0.082 (0.058)	0.007 (0.109)	0.629* (0.362)	-0.298 (0.329)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.006	0.005	0.000	0.006	0.605	0.077
<i>N</i>	260398	237132	186286	134758	83916	47662
Log of Filings From Non-Underwater Borrowers	0.059 (0.067)	0.149** (0.060)	0.106** (0.053)	0.119 (0.089)	0.097 (0.217)	0.086 (0.228)
Log of Filings From Underwater Borrowers	-0.078 (0.053)	0.025 (0.050)	0.030 (0.052)	-0.004 (0.073)	0.546** (0.272)	-0.020 (0.199)
First-stage F	127.300	107.400	83.690	15.110	4.309	14.370
<i>p</i> val for difference between groups	0.080	0.018	0.281	0.217	0.103	0.641
<i>N</i>	127218	116391	91160	65799	41166	23303

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5. Outcomes are for two separate counts of new foreclosure filings by either borrowers in negative equity or borrowers in positive equity (as defined in the text). P-value tests the significance between the two responses (positive vs. negative equity)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

Table 1.28: Interaction with Positive Price Growth

	Years Since Decision					
	0	1	2	3	4	5
Any Filing	0.046 (0.028)	0.017 (0.028)	0.068** (0.027)	0.076 (0.051)	0.227** (0.112)	0.121 (0.134)
Interaction of Foreclosure with Positive Price Growth	-0.042 (0.065)	-0.052 (0.045)	0.528* (0.308)	0.094 (0.080)	-0.125 (0.141)	-0.087 (0.109)
First-stage F	75.470	98.680	3.115	12.140	9.705	6.768
N	127980	116456	91358	65987	41063	23280
Total Filings	0.697 (0.549)	0.605* (0.364)	0.255 (0.232)	0.609* (0.324)	1.566* (0.946)	0.385 (0.999)
Interaction of Foreclosure with Positive Price Growth	-0.342 (2.215)	1.648 (1.268)	17.110* (10.130)	2.593 (2.196)	-7.903* (4.072)	-1.292 (4.832)
First-stage F	75.470	98.680	3.115	12.140	9.705	6.768
N	127980	116456	91358	65987	41063	23280
Log of Filings in Zip Codes/Years with Negative Price Growth	-0.003 (0.074)	0.135** (0.061)	0.069 (0.054)	0.176* (0.091)	0.202 (0.225)	-0.093 (0.217)
Interaction of Foreclosure with Positive Price Growth	0.156 (0.281)	0.084 (0.150)	1.805 (1.182)	0.328 (0.238)	-0.879 (0.606)	-0.021 (0.386)
First-stage F	51.130	96.730	2.502	14.250	8.214	13.970
N	93583	85496	67258	48491	30449	17273

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, but allowing for the interaction of the foreclosure treatment with an indicator for positive price growth in that zip code in that year (from Zillow price indices).

Table 1.29: Response Among Borrowers with Conventional Mortgages

	Years Since Decision					
	0	1	2	3	4	5
Any Filing Among Borrowers with Non-Conv. Mortgages	0.022 (0.028)	0.050* (0.030)	0.034 (0.030)	0.087 (0.054)	0.302** (0.133)	0.060 (0.158)
Any Filing Among Borrowers with Conventional Mortgages	0.019 (0.033)	0.023 (0.032)	0.094*** (0.031)	0.074 (0.048)	0.131 (0.117)	0.189 (0.187)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.952	0.528	0.088*	0.821	0.209	0.529
<i>N</i>	260398	237132	186286	134758	83916	47662
Total Filings Among Borrowers with Non-Conv. Mortgages	0.226 (0.212)	0.391** (0.181)	0.145 (0.093)	0.365* (0.193)	0.894 (0.598)	0.223 (0.491)
Total Filings Among Borrowers with Conventional Mortgages	0.465** (0.210)	0.278 (0.222)	0.392*** (0.120)	0.292* (0.173)	0.656 (0.434)	0.315 (0.586)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.099*	0.503	0.017**	0.623	0.448	0.831
<i>N</i>	260398	237132	186286	134758	83916	47662
log(# of Filings) Among Borrowers with Non-Conv. Mortgages	-0.032 (0.054)	0.119** (0.049)	0.075 (0.051)	0.119 (0.081)	0.242 (0.251)	-0.084 (0.210)
log(# of Filings) Among Borrowers with Conventional Mortgages	0.125** (0.058)	0.065 (0.066)	0.089* (0.052)	0.134 (0.083)	0.231 (0.201)	-0.060 (0.187)
First-stage F	117.800	96.760	72.820	14.810	8.370	13.850
<i>p</i> val for difference between groups	0.001***	0.357	0.791	0.838	0.959	0.885
<i>N</i>	136995	126760	100375	73225	45688	25633

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5. Outcomes are for two separate counts of either new foreclosure filings among borrowers with conventional mortgages or filings among borrowers with unconventional mortgages. P-value tests the significance between the two responses (conventional vs. non-conventional)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

Table 1.30: Contagion among Filings with -10% - 10% Equity versus < -10% or > 10%

	Years Since Decision					
	0	1	2	3	4	5
Any Filing with Equity < -10% or > 10%	0.024 (0.025)	-0.019 (0.027)	0.033 (0.028)	-0.019 (0.044)	0.150 (0.103)	-0.076 (0.151)
Any Filing with Equity between -10% and 10%	0.042 (0.028)	0.038 (0.027)	0.083*** (0.032)	0.156*** (0.053)	0.188 (0.115)	0.161 (0.150)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.582	0.077	0.187	0.001	0.727	0.267
<i>N</i>	260398	237132	186286	134758	83916	47662
Total Filings with Equity < -10% or > 10%	-0.233*** (0.068)	-0.212*** (0.069)	-0.085 (0.062)	0.028 (0.091)	0.228 (0.273)	-0.178 (0.332)
Total Filings with Equity between -10% and 10%	0.924** (0.391)	0.881*** (0.337)	0.621*** (0.178)	0.629** (0.270)	1.323* (0.789)	0.716 (0.853)
First-stage F	119.200	112.100	102.500	12.310	9.534	6.486
<i>p</i> val for difference between groups	0.004	0.001	0	0.014	0.085	0.251
<i>N</i>	260398	237132	186286	134758	83916	47662
Log of Filings with Equity < -10% or > 10%	-0.168** (0.079)	-0.012 (0.052)	-0.031 (0.051)	0.167** (0.081)	0.037 (0.237)	0.105 (0.209)
Log of Filings with Equity between -10% and 10%	0.070 (0.062)	0.155** (0.064)	0.117** (0.053)	0.107 (0.089)	0.264 (0.228)	-0.028 (0.198)
First-stage F	58.860	101.600	90.690	14.320	3.189	10.910
<i>p</i> val for difference between groups	0.015	0.013	0.008	0.539	0.290	0.602
<i>N</i>	116814	108614	86569	63587	40557	23479

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5. Outcomes are for two separate counts of new foreclosure filings by either borrowers where the difference between outstanding debt and loan principal is within 10% of initial loan principal and all other borrowers (see definition in the text). P-value tests the significance between the two responses—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

Table 1.31: Contagion by Proxy for Neighborhood Homogeneity

	Years Since Decision					
	0	1	2	3	4	5
Main Effect of Completed Foreclosure on Any Filing	0.037 (0.040)	0.031 (0.038)	0.090** (0.038)	0.105 (0.066)	0.241* (0.141)	0.185 (0.185)
Interaction with Single Race > 75% of Tract	0.025 (0.054)	-0.038 (0.052)	-0.020 (0.053)	-0.034 (0.093)	-0.008 (0.212)	-0.095 (0.260)
First-stage F	108.600	99.770	84.300	9.773	9.330	5.559
N	130199	118566	93143	67379	41958	23831
Main Effect of Completed Foreclosure on Total Filings	0.404 (0.558)	0.834 (0.662)	0.570* (0.296)	1.346** (0.529)	1.677 (1.047)	-0.096 (1.512)
Interaction with Single Race > 75% of Tract	0.534 (0.813)	-0.338 (0.650)	-0.132 (0.369)	-1.385** (0.630)	-0.394 (1.416)	1.216 (1.810)
First-stage F	108.600	99.770	84.300	9.773	9.330	5.559
N	130199	118566	93143	67379	41958	23831
Main Effect of Completed Foreclosure on Log of Filings	0.011 (0.094)	0.172* (0.088)	0.090 (0.082)	0.276** (0.121)	0.252 (0.273)	-0.277 (0.345)
Interaction with Single Race > 75% of Tract	0.009 (0.110)	-0.082 (0.108)	-0.009 (0.104)	-0.257 (0.166)	-0.119 (0.391)	0.380 (0.423)
First-stage F	95.590	96.070	72.730	11.620	7.097	10.310
N	95155	86961	68482	49422	31016	17628

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, but allowing for the interaction of the foreclosure treatment with indicator for property being in a census tract where one race makes up greater than 75% of the population.

Table 1.32: Contagion by Condo Status

	Years Since Decision					
	0	1	2	3	4	5
Main Effect of Completed Foreclosure on Any Filing	0.060** (0.029)	0.014 (0.028)	0.074** (0.029)	0.074 (0.057)	0.250** (0.116)	0.155 (0.137)
Interaction with Property Being a Condo	-0.045 (0.041)	-0.013 (0.035)	0.042 (0.044)	0.092 (0.066)	-0.017 (0.109)	-0.074 (0.085)
First-stage F	120	112.400	102.400	12.400	9.603	6.414
N	130199	118566	93143	67379	41958	23831
Main Effect of Completed Foreclosure on Total Filings	0.563 (0.392)	0.261 (0.296)	0.079 (0.192)	0.549* (0.331)	1.529* (0.916)	0.248 (1.032)
Interaction with Property Being a Condo	0.704 (1.656)	2.236** (0.986)	2.446*** (0.869)	0.615 (1.326)	0.117 (2.005)	1.453 (1.776)
First-stage F	120	112.400	102.400	12.400	9.603	6.414
N	130199	118566	93143	67379	41958	23831
Main Effect of Completed Foreclosure on Log of Filings	0.082 (0.075)	0.125** (0.063)	0.024 (0.057)	0.106 (0.094)	0.112 (0.227)	-0.051 (0.231)
Interaction with Property Being a Condo	-0.321 (0.196)	0.036 (0.146)	0.352** (0.161)	0.307 (0.219)	0.501 (0.456)	-0.061 (0.240)
First-stage F	139.100	107.200	97.200	14.660	8.312	12.390
N	95155	86961	68482	49422	31016	17628
N	95155	86961	68482	49422	31016	17628

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5, but allowing for the interaction of the foreclosure treatment with the given indicator that the property is a condo.

Table 1.33: Contagion Among Borrowers with the Same Lender

Years Since Decision	0	1	2	3	4	5
Any Filing (Different Lender)	0.062** (0.029)	0.037 (0.027)	0.073*** (0.028)	0.092* (0.054)	0.287** (0.119)	0.070 (0.163)
Any Filing (Same Lender)	0.001 (0.021)	0.014 (0.021)	0.020 (0.019)	0.013 (0.029)	0.048 (0.074)	0.029 (0.094)
First-stage F	124.700	112.400	112.700	13.060	8.506	4.524
<i>p</i> val for diff. = same	0.104	0.488	0.126	0.146	0.040	0.794
<i>N</i>	231920	211018	164628	117636	70910	38210
Total Filings (Different)	0.816** (0.362)	0.887*** (0.343)	0.683*** (0.182)	0.742** (0.347)	1.112 (0.700)	-0.081 (1.108)
Total Filings (Same)	-0.194*** (0.061)	-0.185*** (0.062)	-0.160*** (0.050)	-0.116 (0.083)	0.024 (0.190)	0.063 (0.254)
First-stage F	124.700	112.400	112.700	13.060	8.506	4.524
<i>p</i> val for diff. = same	0.007	0.002	0.000	0.012	0.018	0.860
<i>N</i>	231920	211018	164628	117636	70910	38210
log(# of Filings) (Different)	0.023 (0.067)	0.127** (0.054)	0.121** (0.053)	0.182** (0.090)	0.145 (0.226)	-0.053 (0.246)
log(# of Filings) (Same)	0.046 (0.087)	0.001 (0.080)	0.033 (0.072)	-0.154 (0.103)	-0.425 (0.394)	0.952 (3.326)
First-stage F	55.050	107.600	41.270	15.800	5.857	0.056
<i>p</i> val for diff. = same	0.838	0.162	0.157	0.000	0.121	0.750
<i>N</i>	98694	88191	67543	47654	28798	15282

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5. Outcomes are for two separate counts of new foreclosure filings by either different lenders than that in the observed case or the same lender. P-value tests the significance between the two responses (same lender vs. other lender)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

Table 1.34: Neighborhood-Level Aggregate Contagion Regressions

	2006–2008			2009–2010		
	(1)	(2)	(3)	(4)	(5)	(6)
F_{it-1}	0.281** (0.123)	-0.051 (0.225)		0.115 (0.075)	0.725*** (0.148)	
F_{it-1}^2		0.066 (0.045)			-0.078*** (0.020)	
F_{it-2}	0.913*** (0.108)	1.177*** (0.163)		0.481*** (0.064)	0.742*** (0.096)	
F_{it-2}^2		-0.062* (0.035)			-0.026*** (0.009)	
$F_{it-1} = 1$			1.793 (2.381)			-0.192 (0.288)
$F_{it-1} = 2$			3.895 (4.288)			0.096 (0.470)
$F_{it-1} = 3$			5.209 (9.861)			0.844 (1.025)
$F_{it-1} = 4$			7.004 (13.671)			-0.841 (2.557)
$F_{it-2} = 1$			1.296*** (0.414)			0.259 (0.238)
$F_{it-2} = 2$			10.262 (8.890)			1.062** (0.503)
$F_{it-2} = 3$			-2.192 (5.702)			1.687 (1.613)
$F_{it-2} = 4$			-2.160 (11.965)			0.006 (2.853)
N	30291	30291	25317	19807	19807	12967
First-Stage F Stat	351.578	205.429	0.060	244.386	145.119	0.071

Notes: ***Indicates significance at the 1% level, **5%, and *10%. Standard errors (and corresponding F statistics) are the maximum of the standard error (minimum of F-stat) clustered on the census tract, clustered on filing month, and multi-way clustered on tract and filing month. 2SLS estimates as in Table 1.5. Outcomes are for two separate counts of new foreclosure filings by either different lenders than that in the observed case or the same lender. P-value tests the significance between the two responses (same lender vs. other lender)—estimates are performed simultaneously (pooling both outcomes and allowing for differing filing month and year of observation fixed effects).

2 Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market (with Wojciech Kopczuk⁵⁷)

⁵⁷Department of Economics and School of International and Public Affairs, Columbia University, NBER and CEPR. wojciech.kopczuk@columbia.edu.

Abstract

Houses and apartments sold in New York and New Jersey at prices above \$1 million are subject to the so-called 1% “mansion tax” imposed on the full value of the transaction. This policy generates a discontinuity (a “notch”) in the overall tax liability. We rely on this and other discontinuities to analyze implications of transfer taxes in the real estate market. Using administrative records of property sales, we find robust evidence of substantial bunching and show that the incidence of this tax for transactions local to the discontinuity falls on sellers, may exceed the value of the tax, and is not explained by tax evasion (although supply-side quality adjustments may play a role). Above the notch, the volume of missing transactions exceeds those bunching below the notch. Interpreting our results in the context of an equilibrium bargaining model, we conclude that the market unravels in the neighborhood of the notch: its presence provides strong incentive for buyers and sellers in the proximity of the threshold not to transact. This effect is on top of the standard extensive margin response. Finally, we show that the presence of the tax affects how the market operates *away* from the threshold—taxation increases price reductions during the search process and in the bargaining stage and weakens the relationship between listing and sale prices. We interpret these results as demonstrating that taxation affects the ultimate allocation in this search market.

2.1 Introduction

Purchasing real estate is a time-consuming and complicated process with large financial stakes and potentially important frictions. Beyond the price, a typical transaction involves many associated costs, including broker’s fees, inspection costs, legal fees, title insurance, mortgage application and insurance fees, and moving costs. In this paper, we rely on a particular type of cost—transfer taxes that are imposed on the value of real estate transactions—to understand how frictions affect the functioning of this market.

Our objective is fourfold. Real estate transfer taxation is common and given the importance of this market it is of interest to understand the empirical implications of such taxes. Second, we take advantage of variation in tax incentives and data on both transactions and listings in order to gain better understanding of the importance of search and matching frictions in this market. Third, we use this context to develop a framework for understanding tax incidence and efficiency costs of transaction taxes in search and matching markets more generally. Other contexts where similar issues arise are labor markets and financial transaction taxes. Fourth, our theory and empirics allows for studying the impact of discontinuous incentives on the *existence* of the market itself.

Our empirical approach relies on variation generated by the discontinuous nature of the taxes imposed in New York and New Jersey, which are levied as a function of the appropriately defined purchase price. A prominent example is the so-called “mansion tax” in New York state (since 1989) and New Jersey (since 2004) that applies to residential transactions of \$1 million or more. The tax rate is 1% and is imposed on the *full* value of the transaction so that a \$1 million sale is subject to a \$10,000 tax liability, while a \$999,999 transaction is not subject to the tax at all. In New York City, all real estate transactions are also subject to the real property transfer tax (RPTT) and in New Jersey they are subject to the Realty Transfer Fee (RTF)—both of these schedules happen to have (smaller) discontinuities as well, as we discuss in Section 2.2. Hence, all of these taxes create tax notches (see Slemrod, 2010), while the introduction of the tax in New Jersey also creates

a time discontinuity.⁵⁸ Furthermore, the statutory incidence is different for the mansion tax (which is the responsibility of the buyer) than for the New Jersey RTF and New York City RPTT (which are the responsibility of sellers, with the exception of sales of new constructions in NYC). Interestingly, such discontinuities are not uncommon—for example, they are also present in the UK (Besley et al., 2013; Best and Kleven, 2013) and D.C. (Slemrod et al., 2012).

Our results allow us to reach three sets of conclusions. First, and perhaps least surprisingly, we find that the tax distorts the price distribution resulting in significant bunching just below \$1 million. This bunching is evident in the distribution of sales in New York displayed in Figures 2.1 and 2.2.⁵⁹ A similar pattern appears in New Jersey after the introduction of the tax. Figures 2.3 and 2.4 demonstrate that the onset of this effect is immediate. The bunching we observe is substantial: our estimates robustly indicate that about \$20,000 worth of transactions shift to the threshold in response to the \$10,000 tax. The strength of this effect does not significantly vary with our proxies for tax evasion and we show, using listings data, that a distortion of similar magnitude is already present when properties are first advertised by sellers, which we interpret as inconsistent with tax evasion. We find some evidence that the effect is weaker (but still very strong) for newly built properties that sell when already finished, suggesting that real adjustments to the characteristics of a property may be part of the effect. Still, we conclude that real responses do not fully account for the extent of bunching and the tax near the threshold imposes a substantial burden on sellers. Results from smaller discontinuities that shift statutory incidence are consistent with this conclusion.

Second, we build a theoretical framework to illustrate and test for unraveling of the market above the threshold—the possibility that the tax locally destroys trades of matches

⁵⁸There are also *geographic* discontinuities that we do not exploit: the RPTT changes at the New York City border, and both RTF and the pre-2004 mansion tax change at the New Jersey-New York border.

⁵⁹Figure 2.1 corresponds to the whole state, while Figure 2.2 is for New York city itself. We chose to present the two figures with different binning and overlaying the fit on just one of them in order to present both the visual evidence of the effects and illustrate salient features of the data (round number bunching) that we control for in the empirical analysis.

with remaining positive surplus. Figures 2.1, 2.2 and 2.3 show that the distribution of prices features a large gap above the threshold. If all transactions with positive surplus when taxed continue to transact in the presence of the discontinuous tax, the gap is expected to reflect observations shifting from just above the notch (in absence of the tax) to bunch below the threshold. In particular, the gap above the notch is expected *not to be bigger* than the extent of bunching below. This is a testable prediction. If it fails, it implies that some of the observations are not occurring over and beyond the standard extensive margin response—the phenomenon that we refer to as “unraveling”: transactions with positive surplus that could otherwise occur not far from the notch are not taking place at all. Moreover, we show that the difference between the size of the gap and the extent of bunching is a lower bound for the number of missing transactions. One explanation for such unraveling is that sellers, who face a large burden from these sales, may instead opt out or continue waiting, or buyers may prefer to continue searching in order to benefit from locally depressed prices.

The implementation of our test for unraveling is straightforward and illustrated in Figure 2.5 (that we explain in more detail later in the paper). Conceptually, we estimate bunching at the threshold by constructing the counterfactual distribution based on the data to the left of the threshold and comparing this to the observed bunching at the notch. We estimate the gap by constructing the counterfactual distribution based on the data well to the right of the threshold (i.e., affected by the tax and, thus, accounting for the standard extensive-margin response) and comparing this to the observed gap in sales above the notch.⁶⁰ Intuitively, the after-tax counterfactual used to estimate the size of the gap is, by construction, already net of the standard extensive margin response (matches that have surplus lower than the tax) and thus the gap only reflects sales that have shifted to the threshold or matches with continued positive surplus in the presence of the tax that do not sell. We argue that these missing transactions would have sold if they were far from the threshold, but are discouraged

⁶⁰In practice, we usually simply allow for a shift in the distribution at the threshold to parsimoniously capture the two different counterfactual distributions. We also report results that rely on separate estimation on both side.

by the incentives presented by the notch to continue searching.

We indeed find that more transactions are missing from the gap than we can observe bunching at the threshold, indicating that the market unravels locally. This effect is large: our baseline estimates indicate that over \$40,000 worth of transactions (i.e., equivalent to the mass of transactions that would sell between \$1,000,000 to \$1,040,000 in absence of the tax) that would still yield positive surplus even with the tax do not take place. This corresponds to 2,800 missing transactions in New York City, out of 380,000 that occurred over the whole period. Hence, by our estimates, this one percent tax, applying at a relatively large threshold, managed to eliminate 0.7% of transactions due to the unraveling effect. To reiterate, our interpretation of this response is conceptually different from the standard demand response that is due to higher taxes discouraging transactions with low surplus: unraveling corresponds to eliminating transactions with positive surplus in the presence of the tax. This additional extensive-margin response indicates that the substantial friction introduced by the transaction tax hampers functioning of the market in some region above the threshold. Given our sources of variation, estimating the standard response would require making strong assumptions about comparability of distributions with and without taxes, and we do not pursue it in this paper.

Recognition and empirical identification of the type of extensive margin response that we focus on is a novel contribution. There are, of course, known examples of frictions eliminating particular markets—most prominently, asymmetric information affecting the existence of insurance markets. We find that the notched design of the tax can destroy a market for housing with values close to the notch, which has not been previously recognized. We argue that such a response is present because of the search frictions in the housing market and may apply in any situation where search is present and the population affected by distortionary incentives is not fixed.⁶¹

Unraveling has important implications for empirical work that relies on notches. Much

⁶¹More speculatively, we briefly comment in Section 2.6 on figures from our online appendix that indicate possible changes in patterns of behavior during the search process.

of this literature focuses on contexts where only the intensive margin is of interest, such as income taxation, and hence this point has not been recognized before. In particular, our results indicate that in situations where the volume of taxable units is endogenous (as in housing, but perhaps not under the income tax far from the filing threshold where non-filing may be negligible), exits around kinks/notches cannot be assumed to be the standard extensive margin and hence such responses cannot be generalized as reflecting the effect on behavior elsewhere. One must be careful about separating unraveling from the standard effect—response to the tax in general may differ substantially from behavior close to the notch.⁶² Similarly, estimation of the extent of optimization frictions as in Kleven and Waseem (2013) relies on measuring the size of the gap, which would partially reflect unraveling if present. Moreover, the possibility of unraveling complicates estimation of a general (i.e., not specific to the notch) intensive-margin response to a notched policy, such as the approach described by Kleven and Waseem (2013). We demonstrate how a general response of price to the transfer tax, as determined by the relative bargaining power of buyers and sellers, could be estimated by adjusting bunching by the size of the gap. This procedure eliminates the part of the response that is driven by the discrete impact of the notch, leaving only the effect of a continuous tax. However, such an estimate will be biased (perhaps substantially) in the presence of unraveling.

Our third set of conclusions finds evidence indicating that the impact of the tax is not limited to the proximity of the threshold, but extends much further. Both price reductions while properties are listed and discounts (the difference between final advertised and sale price) increase permanently above the threshold, indicating that the search and matching process is affected everywhere by the tax. Furthermore, we find that in the presence of the tax listing prices are a weaker signal of the final sale price of the property. Relying on our theoretical arguments, we interpret this greater dispersion of sale price conditional on asking price as corresponding to increased deviation from the efficiency-maximizing matching

⁶²In particular, extensive margin responses in other papers exploiting discontinuities in transaction taxes (Slemrod et al., 2012; Best and Kleven, 2013; Besley et al., 2013) are subject to this critique.

equilibrium and conclude that a general transaction tax increases inefficiency in the search process.

A small literature focuses on the effect of transfer taxes on the functioning of the real estate market (Benjamin et al., 1993; Van Ommeren and Van Leuvensteijn, 2005; Dachis et al., 2012). Contemporaneously, three other papers (Slemrod et al., 2012; Best and Kleven, 2013; Besley et al., 2013) look at similar distortions to the distribution of final prices (but not listings) in the United Kingdom and Washington, D.C. These studies focus on the standard extensive margin response (the general effect on sales) to policy changes, rather than incidence, listings and search frictions as we do. We are also unique in showing evidence of unraveling—that the extensive margin effect of the tax goes beyond eliminating transactions with negative net-of-tax surplus. Beyond offering the first, to our knowledge, evidence of this type of an effect, these results also cast doubt on generalizing from responses around notches and kinks (where market can unravel) to elsewhere (where only standard extensive margin response should be present) in the presence of matching frictions.

Another strand of literature to which we contribute analyzes the search and matching process in the real estate market. Several studies focus on the role of listing prices and bargaining in determining the final sale outcome (c.f. Han and Strange, 2012, 2013; Merlo and Ortalo-Magne, 2004; Haurin et al., 2010), while a few apply more general search models to real estate data (c.f. Carrillo (2012); Genesove and Han (2012)). However, these studies do not explicitly identify the effect of transaction costs, such as transfer taxes, on outcomes. A related line of study focuses on the role of real estate agents, attempting to unbundle the effect of cost from information provision (Levitt and Syverson, 2008; Jia and Pathak, 2010; Bernheim and Meer, 2013). Finally, a number of empirical papers incorporate information available in real estate listings data to the study of seller behavior in the housing market (Genesove and Mayer, 2001; Carrillo and Pope, 2012).

Our paper is also related to the broader body of work on behavioral responses to taxation. As in the research on responses to income taxation, we are interested in separating real,

timing, avoidance, and evasion responses (Slemrod, 1990; Saez et al., 2012). Contrary to that strand of work, our context requires considering both sides of the market. There has been a recent revival of interest in estimating the incidence of specific taxes/transfers (e.g., Doyle Jr. and Samphantharak, 2008; Mishra et al., 2008; Hastings and Washington, 2010; Marion and Muehlegger, 2011). Real estate tax is more complicated due to non-homogeneity of goods traded, and the closest analogue in the literature is work on incidence of income/payroll taxes or credits (Rothstein, 2010; Saez et al., 2011).

The structure of the paper is as follows. In the next two sections, we discuss the institutional and policy context and our data. In Section 2.4 we present our theoretical framework. We start by introducing a bargaining framework that illustrates the effect of the tax for a particular match, followed by discussion of frictionless equilibrium and predictions regarding the effect of the tax on the price distribution. We then derive simple testable implications of the presence of frictions to matching. In Section 2.5, we present empirical results about the distribution of prices, both graphical evidence and local incidence estimates for various types of taxes, relying on price and listings data. We also show evidence for various subsamples in order to shed a light on the role of evasion and real adjustments. In Section 2.6, we focus on distortions to the matching process near the threshold and present our results about the extent of unraveling. In Section 2.7 we demonstrate the global effect of the tax on discounts and informational content of listings. Conclusions are in the final section.

2.2 Policy

Real estate transfer taxes are common across the U.S. These taxes are applied to the sale price of real property, and range from as low as no tax in Texas to 2% in Delaware. In New York and New Jersey, the tax rates change discontinuously with total consideration, creating corresponding notches in total tax liability. Table 2.1 contains details of the relevant tax schedules. One notch arising in both states is due to the mansion tax: a 1% tax on the total consideration for homes costing \$1,000,000 or more. Under the mansion taxes of both

New York and New Jersey buyers' total tax liability jumps by \$10,000 when the sales price moves from \$999,999 (where the tax does not apply) to \$1,000,000 (where the tax comes into effect). In New Jersey, the mansion tax was introduced on August 1st, 2004, and covers all residential real-estate transactions. In New York state, the mansion tax was introduced in 1989 and applies to the sale of individual coop and condo units, and one-, two-, and three-family homes, with few exceptions.⁶³

Real estate sales in New York City and New Jersey are subject to additional taxes with discontinuous average rates. In New York City, the Real Property Transfer Tax (RPTT) applies to residential sales (as defined for the New York mansion tax) with a rate of 1% if the total consideration is \$500,000 or less, and 1.425% above \$500,000. Commercial sales are also subject to the RPTT at a rate of 1.425% below \$500,000 and 2.625% above. Unlike the mansion tax the statutory incidence of the RPTT falls on the seller by law, however it is customary for the buyer to pay the tax when purchasing directly from a sponsor (i.e., purchasing a newly developed condo or a newly offered coop). Thus, the RPTT is a unique tax in that there is variation in the statutory incidence.

Residential sales in New Jersey are subject to the Realty Transfer Fee. This transfer fee (or tax) has a non-linear schedule (see Table 2.1) that shifts when total consideration is greater than \$350,000. The marginal tax rate for consideration above \$200,000 is 0.78% if the total price is less than or equal to \$350,000, while this tax rate jumps to 0.96% when the total price is greater than \$350,000. Moving from a price of \$350,000 to \$350,001 increases the buyer's tax liability by \$630.

⁶³Exceptions are as follows. If a residential unit is partially used for commerce, only the residential share of the total consideration is subject to the tax (although the entire consideration is still used to determine if the tax applies). Similarly, multiple parcels sold in the same transaction are taxed as one unit unless the parcels are evidently not used in conjunction with one another. Vacant lots are exempt from the mansion tax, and, finally, any personal effects sold with the home are deducted from the total consideration for tax purposes (but are subject to state sales taxes).

2.3 Data

We study administrative records on real estate transactions in New York state and New Jersey as well as historical real estate listings in New York City. Sales records, which cover the universe of recorded real estate transactions in the given geography and time period, come from three sources: the New York City Department of Finance’s (NYCDOF) Annualized Rolling Sales files for 2003–2011 (covering all of NYC), real property transfer reports compiled by the New York state Office of Real Property Services (NYSORPS) for 2002–2006 and 2008–2010 (all of NY State, excluding the five counties of NYC), and SR1A property transfer forms collected in the New Jersey Treasury’s SR1A file for 1996–2011 (all of NJ). These records contain details of each transaction, including the date of sale, the total consideration paid by the buyer, the address of the property, the property type (e.g., one-, two-, or three-family home, residential coop or condo, etc.), the year of construction of the building (in NYC and NJ) or whether the property is newly constructed (in New York state), and whether the sale is arms-length (NYSORPS only).⁶⁴ We also use deeds records for 1996–2008 for NYC collected by an anonymous private data provider from the county records, which indicate whether the buyer relied on a mortgage (see Appendix 2.9.3 for more information about data sources).⁶⁵

We identify sales that are subject to each of the transfer taxes. Misclassification in taxable status, if any, will introduce a bias against finding effects, however it is unlikely to be substantial: our information comes from administrative records that contain sufficient information to classify, even though an explicit taxability flag is not provided. In New Jersey, we consider all “residential” sales to be taxable (mansion tax and Transfer Fee). For New York state, we define all single-parcel residential sales of one-, two-, or three-family homes, condos, and seasonal properties as subject to the mansion tax. Finally, in New York City

⁶⁴This definition excludes sales between current or former relatives, between related companies or partners in business, sales where one of the buyers is also a seller, or sales with “other unusual factors affecting sale price.” See Appendix 2.9.3 for more details.

⁶⁵We are grateful to Chris Mayer and the Paul Milstein Center for Real Estate for access to this data.

we define all single-unit (non-commercial) sales of one-, two-, or three-family homes, coops, and condos as taxable.⁶⁶

We match the Rolling Sales data for Manhattan to a subset of historical real estate listings in order to get a broader picture of the effect of the tax on the real-estate search process. Our listing data comes from the Real Estate Board of New York’s (REBNY) electronic listing service and covers all closed or off-market listings between 2003 and 2010. REBNY is a trade association of about 300 realty firms operating in New York City and represents a substantial share of listings in Manhattan and Brooklyn. Members are required to post all listings and updates to the electronic listing service. We limit attention to the more complete Manhattan listings, which accounts for approximately 45% of Manhattan sales in the Rolling Sales files.

From the REBNY listing service we observe details of all (REBNY-listed) closed or off market listings since 2003. These data include initial asking price and date of each listing, all subsequent price updates, and the date the property sells or is taken off the market. To acquire the final sale price, we match these listings by precise address (including apartment number) and/or tax lot to the NYCDOF data for Manhattan. We obtain approximately a 90% match rate for listings identified as “closed” in the REBNY data.⁶⁷

Table 2.2 presents descriptive statistics from the three sources of sales data. Overall, we have records for 3,256,597 taxable sales (with non-zero sale price) spanning 1996–2011. The distribution is skewed: mean sale price is higher than the median. Unsurprisingly, prices are highest in New York City. Although median (and mean) sale price is well below the \$1,000,000 threshold of the mansion tax, there are still several thousand sales per geography within \$50,000 of the mansion-tax cutoff. Table 2.3 presents the count of taxable residential sales and median prices over time for the three regions. The growth of housing sales and prices throughout the early 2000s is evident here, as is the subsequent drop in total sales

⁶⁶While multi-parcel sales in New York state are typically subject to the mansion tax, such a sale may be split for tax purposes if structures on adjacent parcels are not used in conjunction with or clearly related to one-another. Since we cannot identify such cases in the New York City and New York state data, we err on the side of caution and exclude all multi-parcel and partially-commercial sales from taxable status.

⁶⁷The match-rate is continuous across the tax thresholds. See Appendix 2.9.3 for details on the matching process.

and median price at the onset of the recession in 2007/2008.

Table 2.4 presents statistics for the matched REBNY listings data. The sales covered by REBNY have much higher prices, on average, than the general NYC rolling sales data (\$1.24 million versus \$658,000). This is due to the REBNY data only covering sales in Manhattan, which is considerably more expensive than the outer boroughs. About 64% of the homes in our sample end up closing (rather than being taken off the market). At the same time, 8% of listings do not close but have a corresponding sale in the rolling sales data, which we interpret as corresponding to either direct sales by owner or sales with a non-REBNY agent. Homes that do not close spend more time on the market (200.5 days vs. 146 for sold properties). These statistics suggest that search frictions are non-negligible in the housing market—the process of finding a buyer for a home is lengthy and sellers are often unable to find a match. There also appears to be bargaining between buyers and sellers: for properties that close and are matched to the rolling sales data, the average discount between the initial asking price and the sale price ($\frac{\text{initial} - \text{sale}}{\text{initial}}$) is 5.9%: 3.2% discount from initial asking to final asking price and a 2.8% discount from final asking price to sale price. However, the median listing in our sample has no price updates between the initial and final asking prices.

2.4 Theoretical Framework

To interpret our empirical findings, we present a simple model of real-estate transactions. We first discuss implications of taxation in a bargaining framework given a match between a buyer and a seller. We then characterize the equilibrium and its responsiveness to taxation absent search frictions. The equilibrium in this situation corresponds to assortative matching. We follow with a discussion of how the price distribution might respond when individuals search and need not transact conditional on matching. Finally, we elaborate on how the equilibrium price dispersion (which is present when there are matching frictions) may respond to taxation, and how these effects can be empirically characterized by relying on observable information (including listings data).

2.4.1 Bargaining

We start with a bargaining model that clarifies the nature of distortions to the price distribution around the notch and relates tax incidence observed at the notch to the bargaining power that determines incidence elsewhere. For now, we abstract from equilibrium considerations and instead characterize pricing behavior given a match. This is a building block of the equilibrium analysis that we come back to below. The Nash bargaining model itself is formally presented in the Appendix 2.9.1. Here, we introduce the intuition underlying the model and illustrate key results on Figure 2.6. The figure corresponds to a lump-sum tax, but the main insights apply as well to the proportional tax that we discuss in more details in the Appendix.

We assume that transaction prices are determined by Nash bargaining between the buyer and seller. Consider a single match (b, s) between a buyer with a reservation value of b and a seller with reservation value s . Given the price that a buyer and a seller negotiate, p , and a lump-sum tax T imposed on the buyer (as in the case of the mansion tax), the parties end up with surpluses of $S^B = b - T - p$ and $S^S = p - s$, respectively. We assume Nash bargaining with seller's weight β so that the price maximizes $\beta \ln(p - s) + (1 - \beta) \ln(b - T - p)$ and is, thus, set to $p(b, s) - \beta T$, where $p(b, s) \equiv \beta b + (1 - \beta)s$ is the price absent taxes. Consequently, the surplus of each side is equal to $S^B = (1 - \beta)(b - T - s)$ and $S^S = \beta(b - T - s)$, which implies that the parameter β determines how the total surplus $b - T - s$ is split between the two parties.

Necessarily, the incidence of the tax is determined by the seller's relative bargaining power, β , as well. Although it follows automatically from this framework, it is worth highlighting that the party with *lower* bargaining power bears a *lower* share of the tax. This party has a lower claim to the surplus in the first place and, symmetrically, experiences a lower reduction in surplus when the tax is imposed. At the extreme, when $\beta = 0$ or $\beta = 1$, one of the parties has no bargaining power and no surplus, and thus is completely inelastic so that it cannot bear any burden of the tax.

Transaction taxes may also discourage sales. Transactions take place when the surplus is non-negative ($b - T - s \geq 0$). All matches (b, s) that satisfy $\beta b + (1 - \beta)s = p$, for some p , sell at exactly the same price (equal to $p - \beta T$) or do not sell at all if total surplus is negative. By reducing the surplus, the uniform lump-sum levy discourages some sales (note that this is the “standard” extensive margin response; we discuss unraveling in Section 2.4.3).

However, the lump-sum tax does not lead to re-ranking of transactions. All prices simply adjust by βT , so that transactions that were occurring at the same price absent the tax continue to sell at equal (although different from the original) prices. This lack of re-ranking is not general: it does not survive considering proportional rather than lump-sum taxation, but it simplifies the following discussion and provides a natural benchmark (and, as discussed in the appendix, the key qualitative results generalize to the proportional tax case).

We illustrate our formal results regarding the price and sales responses to the tax graphically on Figure 2.6. The figure shows reservation values of buyers and sellers on the two axes, and also allows for tracing prices. The contract line (the relationship between prices of buyers and sellers) in absence of the tax requires that the prices of buyers and sellers have to be the same: $p^B = p^S$; while in the presence of the tax above the notch it is given by $p^B = p^S + T \cdot (p^S \geq H)$, where H is the notch ($H = \$1,000,000$ for the mansion tax). The solid black line represents this pricing/budget constraint below the notch, the dashed black line represents the no-tax situation above the notch, and the solid blue line represents the contract line in the presence of the tax (above the notch). The straight green line and other lines parallel to it show the locus of matches with the same sale price in absence of the tax—i.e., a constant value of $p(b, s)$ (corresponding to the intersection of the given parallel line with the black no-tax contract line). In the presence of the lump-sum tax all matches on any of these lines sell at the same price, given by $p(b, s) - \beta T$, which corresponds to a point where a given constant price line intersects with the contract line. Transactions in the gray shaded area, marked “Z”, have positive surplus and would sell without the tax, but they do not sell when the tax is present because the net-of-tax surplus turns negative. Matches in

this region reflect the standard extensive margin response.

Price adjustments are affected by the presence of the notch, resulting in some transactions moving from above to just below the notch, and can be broken-down into four cases that depend on the initial (absent-tax) price. Case 1 are transactions that originally occur below the price of H and are not affected by the tax (buyer-seller matches in region “D”). Case 2 are transactions that, in absence of the tax, sell at a price between H and $H + \beta T$ and would sell below H if the tax was uniform (i.e., not notched), but in the presence of the notch occur there. These transactions are illustrated in Figure 2.6 by the yellow area marked “A” that is bounded by the $p(b, s)$ schedules corresponding to $p(b, s) = H$ and $p(b, s) = H + \beta T$. Given the assumption of maximizing the Nash-bargaining objective function, transactions that originally sell at a higher price than $H + \beta T$ may sell at the notch depending on whether total surplus at the notch or with the tax is higher. We show that for some intermediate range of original (absent-tax) prices above the threshold (Case 3), some transactions will bunch at the notch (region “B”), while others will sell at a new price above the notch (region “C”), depending on the relative size of the seller’s and buyer’s reservation values for the property. However, when the original price is high enough (greater than some *finite* pre-tax price corresponding to the solid green line) no transaction will bunch, constituting our final case (region “E”). In the appendix we establish that the qualitative characterization of the effect of taxation on prices described by Figure 2.6 is general.

The formal model delivers an additional result that is not a priori obvious: there exists a pre-tax price above which the presence of the notch does not affect any transactions. Firstly, for any β there exists a pre-tax price above which transactions are not affected by the notch—this boundary is determined by considering matches involving a seller with reservation value of zero. Furthermore, and less intuitively, there exists a single *finite* bound for such maximum pre-tax prices above which transactions do not bunch that applies uniformly for any value of β . In the lump-sum tax case, the tight bound (applying as $\beta \rightarrow 0$) is the solution to $\ln(H) - \ln(p) - \frac{H-T}{p} + 1 = 0$ and corresponds to $p \approx \$1,144,717$ when $H = \$1,000,00$ and

$T = \$10,000$. In the case of the proportional tax, the bound solves $\ln(H) - \ln(p) - \frac{H-p\ln(1+t)}{p} + 1 = 0$ (where t is the marginal tax rate, 1% in the case of mansion tax) corresponding to $p \approx \$1,155,422$. While this precise characterization is the consequence of functional form assumptions in the case of Nash bargaining, it does provide a theoretical justification for the assumption that we make in our empirical analysis that only matches in some finite omitted region around the notch might bunch. In our empirical analysis, we use this theoretical bound for defining the omitted region in our baseline specification (but of course we explore the sensitivity to this choice).

2.4.2 Equilibrium

The framework that we have introduced so far focuses on price determination given a match between a buyer and seller. This is a component of the equilibrium description—given matches that lead to sales, we assume Nash bargaining as the approach for determining the price. The missing component of the model is a description of how matches form. Providing a complete search framework is beyond the scope of this paper and, in fact, we are not aware of a framework in the literature that would incorporate two-sided search in the real estate context (Carrillo, 2012, makes a step in this direction by setting up, but not explicitly solving, a model of this kind). We make two arguments that we then investigate empirically. First, we consider what the distribution of prices reveals about the distribution of the underlying matches and efficiency of the equilibrium. Second, we consider which of the matches are likely to be “stable” in the presence of the tax.

The simplest way to introduce equilibrium consideration in this framework is to assume random matching followed by very large frictions (search costs) preventing both parties from further search so that, conditional on a match, the only decision to make is whether to transact. Under this assumption, there is a matching technology $M(b, s)$ that results in some (smooth) distribution of matches over (b, s) and Figure 2.6 reflects which of those matches correspond to transactions.

An alternative is to consider a situation with no search frictions. Suppose that we have an equal number of buyers and sellers. Absent taxation, the overall surplus from a match in our Nash bargaining model is strictly supermodular (it is given by $\ln(b - s) + \text{constant}$). Hence, the equilibrium and, simultaneously, the *efficient allocation* that maximizes the overall surplus involves positive assortative matching. In the presence of taxation, the surplus for transactions not subject to the tax and not at the threshold remains $\ln(b - s) + \text{constant}$, while the surplus for transactions subject to the tax is $\ln(b - s - T) + \text{constant}$ (which is, naturally, also supermodular). Hence, *within* each of these groups maximization of the overall surplus involves assortative matching.

The efficient allocation absent taxation corresponds to an increasing profile of matches (b, s) . This profile is illustrated on Figure 2.6 using a wiggly solid gray line. If these matches were to remain when the tax is introduced, a match corresponding to point X on the figure would be the highest priced one that is not subject to taxation, the match marked by Y would be the lowest priced one that does not shift to the notch, and matches between X and Y would move to the notch.

Introduction of the tax may affect which matches take place in the equilibrium. The efficient allocation will retain the main features visible on Figure 2.6, although the actual equilibrium profile and points X and Y need not coincide with the allocation absent tax distortions. As argued above, the equilibrium matches subject to the tax will continue to be assortatively matched—that is, matches above point Y will lie on an increasing profile. Similarly, matches below point X will form an upward-sloping profile. For matches that are priced at the threshold, the price is fixed at H so that any permutation of residences between buyers would deliver exactly the same surplus (a feature that is arguably peculiar to this model), so that the precise identity of matches between X and Y is indeterminate.⁶⁸

The equilibrium allocation in the presence of frictions will not be efficient, although

⁶⁸The location of X and Y may change because fewer transactions may take place in the presence of a tax. Figure 2.6 appears to preclude this possibility by using the efficient matching schedule that does not involve matches that are crowded out by the tax, but it need not be so in general.

efficiency serves as a natural reference point. Trivially, our theoretical framework implies that under the efficient allocation variance of the price (or, equivalently, buyer’s type) conditional on seller’s type (and vice versa) is zero, because the efficient allocation corresponds to an upward sloping line.⁶⁹ In the presence of frictions, matches would occur not just on the efficient allocation line as in Figure 2.6, but rather could be spread over the rest of the region corresponding to surplus from transacting. Intuitively, one might expect that an increase in dispersion of prices for a given type of home corresponds to an allocation that is further from the efficient one. While we do not prove this result, it has to be so at least when the frictions are small. To examine whether transaction taxes affect the efficiency of the housing market allocation, we test whether there is an increased dispersion of prices in the presence of the tax. Because matches are indeterminate in the bunching region, this exercise is of interest for transactions that are not local to the threshold—i.e., those far enough from the threshold in either direction. We discuss this further in Section 2.4.4.

2.4.3 Measuring the Impact on the Price Distribution

We expect that the tax affects the distribution of home sales by inducing both bunching of sales at the notch and by creating a “gap” in the distribution above the notch, and we estimate both. Our objectives are twofold. First, descriptively, these estimates allow us to quantify the magnitude of distortions to the price distribution. Second, we use these estimates to back out values analogous to the mass in regions A and B of Figure 2.6. Intuitively, extra transactions bunching at the threshold correspond to regions A and B , while transactions that are missing from the distribution above the threshold (relative to the distribution further to the right) reflect region B . In principle then, these values may be used to recover the mass in region A that is tightly linked to the bargaining parameter, β .

However, in what follows, we argue that the tax notch provides a strong incentive for neighboring “productive” matches—those close to the boundary between regions B and C ,

⁶⁹While the efficient allocation is not unique in the bunching region A and B on Figure 2.6, the price is constant and equal to the threshold level in that region.

which have positive surplus in the presence of the tax—to break. Consequently, there may be more transactions missing from the distribution above the notch than locate at the notch itself. We show that our estimates can be used to test for and bound this local extensive-margin response: if such a response is not present, the excess number of transactions at the notch must exceed the gap in the distribution. We view this test as one of the central contributions of our paper, as it corresponds to testing for unraveling of the market due to the presence of the notch. Equivalently, this is a test of whether the extensive margin response is standard (the gray region in Figure 2.6)—a condition that is necessary to generalize from estimates based on a notch to behavioral responses to a general tax (as done by Best and Kleven (2013) and Slemrod et al. (2012)). Previewing our results, we find that overwhelmingly the answer is that it is not. The rest of this section serves to define quantities that we estimate and to formalize the test for the local unraveling of the market.

Observed and counterfactual distributions. We first consider how the distribution of sales is distorted by the tax notch. The discussion is graphically illustrated on Figure 2.5, which is a distribution analogue of Figure 2.6 (with corresponding region labels and coloring). We denote by $F(p)$ the “true” (population) price distribution in the presence of the actual (notched) tax from which our observations are drawn. We denote by $F_T(\cdot)$ the observed cumulative population price distribution—a draw from $F(p)$. In order to characterize and interpret distortions to the distribution, we presume that there is a set of potential matches (from some matching technology) that may result in transactions. We leave the origin of the set of matches unspecified, and simply assume that there is some matching technology that we take as given. Matches are indexed by i and have associated with them three prices $(p_i, \check{p}_i, \tilde{p}_i)$: the actual price p_i in the presence of a “notched” tax, the shadow price \check{p}_i that would prevail if the tax did not apply anywhere, and the shadow price \tilde{p}_i that would prevail for the same transaction if the tax was proportional everywhere (i.e., involved no threshold). In the context of our model as illustrated on Figure 2.6, $p_i = \check{p}_i$ for transactions taking place

in region D , $p_i = \tilde{p}_i$ in regions C and E , and $\check{p}_i \geq H > p_i$ in regions A and B . Analogously, this respectively corresponds to the transactions far below the notch, transactions above the notch, and the bunching region on Figure 2.5. We assume that $\check{p}_i > \tilde{p}_i$, which excludes the polar case of incidence fully borne by buyers, but allows for simplifying notation. We also allow for either of the prices to be infinite, corresponding to the transaction not taking place in a given regime. We do not rule out in general that the notched tax affects the equilibrium distribution everywhere (even below the threshold): $F(p)$ corresponds to the actual equilibrium outcome; however, \check{p}_i and \tilde{p}_i are prices specific to matches that form in the observed equilibrium given a notched tax, so that their marginal distributions do not reflect any changes regarding which matches would form if the the tax was removed or replaced by one that is proportional everywhere.

We rely on two counterfactual distributions for our estimates: one that corresponds to the non-taxable regime and another corresponding to the taxable one. $F_0(p) = P(\check{p}_i < p)$ is the counterfactual distribution corresponding to the non-taxable regime. Below the taxable threshold, H , F_0 is the true distribution net of transactions that are affected by the presence of the tax ($\check{p}_i \geq H > p_i$). We define a counterfactual distribution $F_1(p) = P(\tilde{p}_i < p)$ corresponding to the region subject to the tax: it is the distribution under a proportional tax with no notch (in particular it is accounting for the “standard” extensive margin response but, of course, without allowing for equilibrium adjustment to the set of matches).

In the presence of the tax, the actual distribution will display an excess mass bunching below the threshold (relative to F_0) and a gap above the threshold (relative to F_1). At some abuse of the notation, we use the region descriptors from Figure 2.6 (e.g., A) to denote the set of matches and the mass (probability) of *transactions* in the corresponding region. Transactions that are distorted by the threshold have $\check{p}_i \geq H > p_i$ and come from a number of different sources: $\tilde{p}_i < H$ (region A except for A'); $\tilde{p}_i = \infty$ (region A'); and $\tilde{p}_i \geq H$ (region B). We note that $A+B = P(\check{p}_i \geq H > p_i)$ and $B = P(p_i < H \leq \check{p}_i \wedge \tilde{p}_i \geq H) = P(p_i < H \leq \tilde{p}_i)$. As discussed in Section 2.4.1, these transactions move from above the threshold (in absence

of the tax) to bunch just below the notch. This movement of sales from above the threshold (and any additional extensive margin response beyond the standard one already embedded in the adjustment to distribution F_1) leaves a gap in the observed distribution just above the notch.⁷⁰

Crowd out of productive matches. One concern is that not every productive match corresponds to a transaction: i.e., that there may be transactions for which $\tilde{p}_i < \check{p}_i < \infty = p_i$ —those with sufficient surplus to survive the tax, but that do not occur in the presence of the notch. To see why, recall that our basic framework assumed that equilibrium matches in the neighborhood of the notch form in a way similar to those away from it, with only the outcome of the bargaining process affected. However, proximity to the threshold provides strong incentive for some buyers and sellers in matches near the notch to continue or delay searching (and perhaps not transact in the end at all). Firstly, consider buyer-seller pairs who would move to the threshold if a sale occurs (region B of Figure 2.6). Sellers in this region—who face a substantial reduction in sale price in moving to the notch—may prefer searching for a buyer with slightly higher reservation value who is, thus, willing to buy above the notch. Secondly, buyers in the buyer-seller pairs that would transact in the gap region above the notch (region C) may have an incentive to return to the market to find a seller with slightly lower reservation value.

Whether this type of local extensive-margin response is present is of interest in its own right, corresponding to both an efficiency loss due to a notched tax in markets with search frictions (productive matches that do not transact) and the importance of search in the housing market. We assume that such exits do not occur for transactions below the threshold ($\check{p}_i < H$, region D) and for transactions that have sufficiently high prices (region E , $\tilde{p}_i > \bar{P}$ for some sufficiently large price \bar{P}). We denote the mass of such exits that comes from matches that could otherwise sell above the threshold in the presence of proportional taxation

⁷⁰Transactions to the left of the threshold (region D) have $p_i = \check{p}_i$, and those with $p_i = \tilde{p}_i$ correspond to regions C and E . The gray region—transactions with surplus low enough that the tax crowds them out—have $\check{p}_i < \infty = p_i = \tilde{p}_i$.

(regions B and C) by $M \equiv P(H \leq \tilde{p}_i \leq \bar{P} < \infty = p_i)$.

Local incidence. Using an estimate of the excess mass bunching at the threshold (I), we can estimate a measure of the incidence of the tax. We define

$$I = F(H) - F_0(H) = P(\check{p}_i \geq H > p_i) = A + B$$

as the number of observations that shift below the threshold due to the tax. Given the observed distribution F_T and an empirical estimate of the counterfactual distribution \hat{F}_0 , we can construct an empirical estimate of the volume of responsive sales making up I as $\hat{I} = F_T(H) - \hat{F}_0(H)$. In practice, we construct the counterfactual price distribution of sales \hat{F}_0 by relying on the actual distribution F_T to the left of the notch, but omitting sales near it—the specifics are in Section 2.4.5. The estimate of bunching, \hat{I} , is represented by the red/yellow area in the first panel of Figure 2.5.

Given an estimate of the counterfactual distribution F_0 to the right of the threshold, we can also define a dollar measure \hat{h} as

$$F_0(H + \hat{h}) - F_0(H) = \hat{I} \tag{14}$$

to represent the magnitude of the shift to the threshold. In other words, \hat{h} is obtained by finding the dollar value such that the integral under the counterfactual to the right of the notch is equal to the excess mass (represented as the green area of the first panel of Figure 2.5). We refer to \hat{h} as “local” or “reduced-form” incidence of the tax. Our preferred interpretation is that, as with the “kinked” budget-set methodology outlined by Saez (2010), \hat{h} represents the average amount of money that is lost (relative to their corresponding non-taxed sale price, \check{p}_i) by sellers participating in the marginal transactions affected by the presence of the threshold. However, the value of \hat{h} does not, on its own, inform us about the underlying bargaining power of the two sides of the market and hence does not reflect incidence of the tax away from the threshold.

We construct and estimate \hat{h} throughout, but our interpretation of \hat{h} as local incidence depends on assumptions about the nature of the counterfactual distribution. With the exception of our data for New Jersey prior to implementation of the tax, we do not observe \check{p}_i for values greater than H at all. When we use data below the notch to project F_0 above the notch, the interpretation of \hat{h} requires additional assumptions. If $F_\tau(p)$ below the notch coincides with the distribution absent taxation $F_0(p)$ (i.e., $p_i = \check{p}_i$ for $\check{p}_i < H$), then the projected counterfactual above the notch corresponds to the distribution absent taxation as well, so that \hat{h} can be thought of as a reduced form dollar estimate of local incidence (this interpretation also applies when we build our counterfactual using the non-distorted distribution in New Jersey before the tax was introduced). However, if the untaxed part of the distribution (below the notch) is affected by the tax via general equilibrium effects (and so $F_\tau(p) \neq F_0(p)$), the projection of the counterfactual above the notch, and the dollar value \hat{h} that relies upon this projection, do not have clear interpretations (although \hat{h} remains a convenient way of standardizing mass of sales bunching below the notch). On the other hand, even if $F_\tau(p) \neq F_0(p)$, our estimate of the excess mass \hat{I} (and unraveling estimates that depend on it) remains valid.

Gap. We also construct a measure of the gap to the right of the notch by comparing an estimated counterfactual above the threshold to the observed distribution. We presume that there is a known value of \bar{P} such that F_1 and F coincide for prices greater than $H + \bar{P}$ (by Theorems 3 and 4, this has to be so in our framework; also recall that we assumed away exit of productive matches for high enough prices) and define the gap in the distribution as:

$$G = [F_1(\bar{P}) - F_1(H)] - [F(\bar{P}) - F(H)]$$

i.e., the difference between the number of transactions taking place in the presence of taxation with and without the notched implementation of the tax. For the estimation, we replace F by F_T , and F_1 by its empirical estimate \hat{F}_1 . The estimate of $\hat{F}_1(\bar{P}) - \hat{F}_1(H)$ reflects the

expected number of observations in regions B and C , while $F_T(\bar{P}) - F_T(H)$ is the actual number of observations in region C . Using our definitions, we can show that

$$G = P(\bar{P} \geq \tilde{p}_i \geq H \wedge (p_i < H \vee p_i = \infty)) = P(p_i < H \leq \tilde{p}_i) + P(H \leq \tilde{p}_i \leq \bar{P} < \infty = p_i) = B + M$$

Recall that M represents the mass of transactions that would have taken place under a proportional tax at prices exceeding the threshold, but do not take place in the presence of the notch. Thus, the gap reflects two effects: exit from the market of productive matches and shift to the threshold.

To reiterate, the gap reflects transactions that are missing from the distribution to the right of the threshold relative to the counterfactual *with taxes*. In particular, it does not include the standard extensive margin response—matches with surplus small enough that they are no longer economically viable in the presence of the tax (the gray region Z in Figures 2.5 and 2.6).

Testing for market unraveling. The gap, G , and behavioral response, I , are related and can be used to test for the presence of market unraveling and, in the absence of unraveling, to estimate β . Both G and I partially reflect transactions in region B —those that would sell at prices higher than the threshold in the presence of a continuous tax, but sell at the threshold when it is discontinuous. Clearly,

$$G - I = M - A$$

We report an estimate of $\hat{G} - \hat{I}$ converted to a dollar figure: $\hat{Z} = \hat{h} \cdot \left(\frac{\hat{G}}{\hat{I}} - 1\right)$.⁷¹ If the market does not unravel in the neighborhood of the notch ($M = 0$), then $G - I = -A \leq 0$. Intuitively, if all buyer-seller matches continue to transact in the presence of the notched

⁷¹Alternative definitions would be to define $F_1(H + \hat{Z}) - F_1(H) = \hat{G} - \hat{I}$ or $F_0(H + \hat{Z}) - F_0(H) = \hat{G} - \hat{I}$. The choice we implement has two advantages. First, it is in terms of the distribution F_0 so that it is directly comparable to \hat{h} . Second, knowing \hat{h} (which we report as well) allows for directly recovering an alternative metric of the gap and behavioral response, $\frac{\hat{G}}{\hat{I}}$.

tax, then the mass bunching at the threshold should always be at least as large as the gap. Hence, given estimates of I and G , we can then test whether the tax destroys productive matches.

Remark 1. Rejecting a testable hypothesis $\hat{Z} \leq 0$ implies market unraveling ($M > 0$).

If the hypothesis of $M = 0$ cannot be rejected, one could construct a straightforward estimate of β . With no missing sales, $I - G = A$, so that $\hat{\beta}$ would solve $F_0(H + \hat{\beta} \cdot T) - F_0(H) = \hat{I} - \hat{G}$.⁷²

Previewing our results, however, we find that $\hat{Z} \leq 0$ is rejected or, put differently, we find that the size of the gap is larger than the number of transactions that bunch. We conclude that there are transactions that do not take place because of the proximity to the threshold so that the market (partially) unravels in its neighborhood.⁷³

One can bound local exit from the market (M) by considering how much missing mass is required to explain our estimates assuming different values of β . In particular, consider the two extreme cases of $\beta = 0$ (buyer captures all surplus) and $\beta = 1$ (seller captures all surplus). In the first case, $A = 0$, while in the second case A corresponds to the mass in the interval of prices $(H, H + T)$. Noting that \hat{Z} is expressed in dollar terms, the dollar-valued mass in the second case is, thus, T . Hence, the implied missing mass when $\hat{Z} > 0$ is between \hat{Z} (when $\beta = 0$) and $\hat{Z} + T$ (when $\beta = 1$). In our discussion of the results, we will refer to the lower bound \hat{Z} .

Finally, note that this discussion provides three qualifications of general interest when relying on notches and kinks in tax schedules for identification. First, for a clean interpretation of our incidence parameter \hat{h} (and, analogously, for estimating elasticities or other measures of behavioral response based on bunching), the counterfactual distribution F_0 needs to correspond to the situation absent the tax—this is a strong assumption that is violated if there

⁷²We treat T as a lump-sum tax here for simplicity of exposition; the effect of adjusting for the marginal tax of 1% is negligible for the purpose of this exercise.

⁷³Naturally, unraveling occurs here because the tax reduces incentives to transact, but in other contexts the incentives may go the other way. Studying a time-notch affecting marriages in Sweden, Persson (2013) compares bunching at the notch and the gap above the notch and finds that in that context discontinuous incentives may encourage transactions at the (time) notch.

are spillover effects from the notch/kink to the non-taxable region. However, our estimates of the number of observations bunching \hat{I} and missing mass \hat{Z} do not require such an assumption. Second, the presence of a notch may provide incentives to exit, corresponding to local unraveling of the market. In this case, gap estimates partially reflect such an exit and can be used to test for its presence when combined with the magnitude of the shift to the notch. Third, as a consequence, when there is such market unraveling, extensive-margin responses estimated by studying local effects of notches/kinks do not generalize to extensive response elsewhere.

2.4.4 Implications for Efficiency of the Equilibrium Allocation

In order to shed a light on how taxation interacts with search frictions away from the threshold in this market, we proceed as follows. Recall that the equilibrium price for a given match (b, s) is equal to $\beta b + (1 - \beta)s - \beta T$. Conditional on the seller's type, $\text{var}[p|s] = \beta^2 \text{var}[b|s]$. If we could directly observe s , the comparison of the variance of prices conditional on s with and without the tax would constitute a test of the hypothesis that taxation affects price dispersion. Evidence of this kind would suggest that the tax increases deviation from efficiency.

In practice, we are unable to observe the seller's type and instead rely on a set of indicators, X , that proxy for it. In that case,

$$\text{var}[p|X] = \beta^2 \text{var}[b|X] + (1 - \beta)^2 \text{var}[s|X] + 2\beta(1 - \beta) \text{cov}[b, s|X]$$

When X contains s , the second and third term are zero. In order to understand how $\beta^2 \text{var}[b|s]$ varies with and without taxation, we consider expanding the set of indicators X —as they become more informative about s , the influence of the last two terms declines and the first term should tend towards $\text{var}[b|s]$. We test whether there is a difference between $\text{var}[p|X]$ with and without taxes for a large set of indicators X correlated with seller's type.

One of the primary indicators of seller's type that we consider is the seller's asking price.

It is natural to think that this price is correlated with the seller’s type, but it may also be endogenous to taxation. In that case, the alternative interpretation of the effect of taxation on $\text{var}[p|X]$ is as a test of whether taxation changes informativeness of this important signal available to buyers.

2.4.5 Econometric Implementation

We estimate the price distribution of sales by maximum likelihood as follows. We specify a parametric distribution of prices absent the tax:

$$\ln f_0(p) = g(p) + \alpha D(p) \tag{15}$$

and the distribution in the presence of the tax:

$$\ln f(p) = \ln f_0(p) + \gamma \cdot I(p > H) \tag{16}$$

where the left-hand side is the log of the probability distribution function at price p , $g(\cdot)$ is a parametric function (a polynomial—third degree in our baseline specification) and D is a set of controls for round numbers. We allow for discontinuity of the density at the threshold (γ) to account for the shift in sale price and global extensive-margin response (gray region Z of Figures 2.5 and 2.6) among transactions subject to the tax. We estimate this model using data that excludes some region around the threshold ($H - \underline{P}, H + \bar{P}$) to ensure that our estimates are not biased by the distortions to the distribution near the notch.⁷⁴ Given the observed distribution of prices outside of the omitted region, we estimate the distribution given by equation (16) by maximum likelihood.⁷⁵ This procedure yields $\hat{f}_0(p)$, our estimate of the counterfactual distribution function of prices absent the tax and $\hat{f}_1 = e^{\hat{\gamma}} \hat{f}_0$, the

⁷⁴Our theoretical framework establishes that there is a value of \bar{P} above which the notch (although not the tax) is irrelevant for the distribution.

⁷⁵Formula 16 is already the log-likelihood and implementation only requires imposing conditions guaranteeing that $f(p)$ is a probability distribution function, i.e., that it integrates to one over the considered interval.

counterfactual in the presence of the tax.⁷⁶ Using these counterfactual distributions, we estimate bunching and gap, local incidence, and bounds on attrition near the threshold as outlined above. Specifically, we estimate the excess mass as the difference between the observed mass in the region $(H - \underline{P}, H)$ and the predicted mass $(\int_{H-\bar{P}}^H \hat{f}_0(p)dp)$, and estimate the gap as the difference between the predicted mass in the region $(H, H + P)$ allowing for a discontinuity at H $(\int_H^{H+\bar{P}} \hat{f}_1(p)dp)$ and the observed mass in $(H, H + \bar{P})$.

While it is common in existing literature exploiting kinks and notches for identification to include round-number dummies to control for bunching at these points, our implementation of the round-number effects, $D(p)$, is more involved. Our baseline approach is to rely on the maximum likelihood estimation and hence specify the density at any point. In order to parsimoniously capture various forms of bunching (in particular, there is bunching in listings data just under round numbers — e.g., at \$899,000), we introduce “bunching” regions for each round number R that extend from $R - b$ to R . Within the bunching regions, the distribution is specified as $g(p) + D_R + D_R \cdot p$ where D_R are the relevant round-number dummies, while it is $g(p)$ otherwise. Interacting the round-number dummies with price allows the extent of rounding to vary with price (perhaps \$1.2 million is not equally as salient as \$600,000). In practice, we allow for rounding at multiples of \$25,000, but allow for separate effects for multiples of \$25,000 and \$50,000. We set $b = \$1000$, which makes the bunching region extend from, for example, \$899,000 to \$900,000 allowing both for bunching at the \$900K level and just below it (e.g., \$899,999). Since our objective is to estimate the counterfactual in the omitted region (in particular, at \$1 million), this approach amounts to assuming that bunching at other round prices is a valid counterfactual for the magnitude of bunching at the tax notch—this is not a directly testable assumption but, as discussed before, data in New Jersey before the introduction of the mansion tax (Figure 2.3) provides support for this assumption. Except for a gain in information by allowing for continuous prices, this approach

⁷⁶We show as a robustness check an estimate of \hat{h} resulting from estimating f_0 using only the data to the left of the threshold. Moreover, we combine this with an estimate of f_1 using only data on the right of the omitted region to attain an estimate of \hat{Z} .

is very close to binning the data in \$1000 bins (and we show the more restrictive “binned” specifications as one of our robustness checks). Beyond that, using maximum likelihood instead of fitting a polynomial to binned data replaces the arbitrary zero-mean restriction for the error terms by a natural restriction that the estimated specification represents a distribution function—arguably, a much more natural assumption than OLS.

All reported standard errors are obtained by bootstrapping the whole procedure 999 times. Note that the estimates of incidence and gap may fall into one of the round-number bunching regions. When this is the case, the estimates are not very sensitive—small changes in parameters correspond to staying in the bunching region. In such cases reported standard errors for estimates of gap and incidence are small, even though standard errors for parameters of the parametric density are not.

2.5 Distortion to the Price Distribution

We begin by demonstrating graphically that the tax has a causal effect on the distribution of prices and timing of transactions. Response to the tax notch is evident in Figure 2.1, which shows the empirical distribution of taxable sales in New York with sales grouped into \$5,000 bins. There is clear bunching in the sale price just below \$1 million and a drop in the volume of sales just above \$1 million. These features, especially the gap above the notch, are obvious when looking at the distribution in logarithmic scale (and \$25,000 bins), as in Figure 2.2. Figure 2.11 in the Appendix shows analogous patterns at the smaller (0.425%) RPTT notch at \$500,000, which also demonstrates some evidence of a response.⁷⁷

We can also verify explicitly that the tax induces bunching by comparing sales in New Jersey before and after the introduction of the tax. Figure 2.3 presents plots of the (log-

⁷⁷There is also significant bunching at other round price levels (at every \$50,000 and, to a lesser extent, remaining \$25,000 multiples), which may confound our bunching analysis. Unlike the bunching at \$1 million, this round-number bunching occurs in the bin above rather than the one below the round number. A priori it is possible that, although this observed bunching below \$1 million is consistent with theoretical predictions, it may simply reflect adjustments to the tax by very small amounts. However, aggregating the data to larger bins in Figure 2.12 mostly eliminates such round-number bunching, while continuing to indicate that the response covers more than just the immediate neighborhood of the threshold.

scaled) histogram of sales in NJ before and after the tax is introduced in \$25,000 bins, with the pre-tax distribution adjusted to account for sales growth and inflation, as discussed in the Figure note (and corresponding to our later empirical implementation that compares NJ before and after the tax).⁷⁸ We see pronounced bunching after the tax is introduced in 2004 and minimal bunching prior to 2004. Data prior to 2004 also shows that \$1 million is no more salient than other multiples of \$50,000 before the tax arrives. Moreover, this figure displays clear visual evidence of the gap above the notch in the post-tax distribution, a feature that is not shared with the pre-tax distribution.

As Figure 2.4 illustrates, the number of sales in NJ just below the threshold clearly increases precisely at the time of the introduction of the tax and the number of sales above the threshold falls. Focusing on the region within \$10,000 of the threshold it is evident that the increase in the mass below the threshold is larger than the shift from just above \$1 million. This difference provides the first clear indication that the local effect of the tax may extend beyond 100% of its value (\$10,000). Figure 2.4 in the Appendix demonstrates that retiming is strong for sales well above the threshold, but that it is unlikely to have lasted for an extended period of time—the excess sales just before the introduction of the tax do not correspond to more than a couple months worth of sales. On the other hand, the pattern of sales in the combined \$900,000–1,100,000 range suggests that, despite pricing effects post-introduction of the tax, overall retiming of transactions in the neighborhood of the threshold is not particularly strong.⁷⁹

⁷⁸We choose this larger bin size to smooth out the bunching at other multiples of \$50,000. The conclusions are the same using \$5,000 bins—in the presence of the tax there is excess bunching just below \$1,000,000 and a gap just above that is not present when there is no tax. However, the bunching at multiples of \$50,000 makes the figure difficult to read when both distributions (pre and post tax) are overlaid.

⁷⁹Appendix Figure 2.14, which shows the monthly distribution of sales in New Jersey between \$900,000 to \$1,100,000, further demonstrates that responses extend well beyond the \$10,000 value of the tax. There is no evidence of the distribution being distorted below \$975,000, and clear evidence of a shift of the mass of sales to just under \$1 million from as far above as the \$1,025,000 to \$1,050,000 range. Additionally, Figures 2.4 and 2.14 show patterns that may be consistent with anticipation effects—there is a spike in sales with prices over \$1 million just before the introduction of the tax. This is not surprising: the tax had been announced prior to coming into effect, and the lengthy process of closing a real estate transaction may allow for the possibility to speed up the timing of final sale. This effect is not long lasting and of no relevance for New York where the tax was introduced long before our data starts.

Our baseline estimates, presented in Table 2.5, demonstrate that the observed bunching translates to a local incidence on sellers that exceeds the magnitude of the tax. The first row reports estimates for New York City and corresponds to the specification shown on Figure 2.2. For our baseline, we use a 3rd order polynomial, while omitting data in the \$990,000-\$1,155,422 region (the upper bound is the theoretical limit discussed in Section 2.4.1), and allowing for an additional constant shift above the tax threshold. Our estimate of the local incidence parameter in the baseline specification is \$21,542.098: bunching at the threshold is equivalent to all transactions over the following \$21,000 shifting to the threshold. These estimates are consistent with impressions from the graphical evidence presented above. Taken literally as an incidence estimate this corresponds to over 200% incidence of the tax on sellers for the marginal transaction. The fourth column presents the corresponding estimates of \hat{Z} : the positive value indicates that there has to be substantial unraveling due to the presence of the threshold or, in other words, there is no β that can rationalize behavior if no such extensive-margin response is present.

In Appendix 2.9.4 we discuss robustness of these results to reasonable modifications of our specification, including choice of polynomials, changes in the omitted region and estimating incidence by OLS relying on binned data. We find that our baseline local incidence results are very robust. We also consider “placebo” treatments at other round numbers that, as expected, show no effect.

Our estimate of about 200% reduced-form incidence local to the threshold is consistent across geographies and data sets. In the top panel of Table 2.5, we report results for New York State (excluding New York City) and New Jersey. Estimates for these regions are remarkably similar—within \$2,000—to those for New York City. In contrast, estimates for New Jersey before the introduction of the tax show no evidence of bunching.

As an alternative to our maximum likelihood approach, we estimate incidence in New Jersey using the pre-tax period as a counterfactual for the post-tax period and find similar results. We implement this pre/post comparison as follows. We omit transactions within 90

days of the policy change (to avoid the retiming response) and focus on the following year (Oct. 30, 2004–Oct. 29, 2005). We rescale the period before the tax (May 3, 2003 to May 2, 2004) to account for sales growth over time. Specifically, we construct a counterfactual growth factor by taking the ratio of the count of sales within \$2500 of each price from May 3, 2002 to May 2, 2003 to the count of sales from Nov. 5, 2000 through Nov. 4, 2001 (omitting sales between Nov. 2001 and May 2002 to mimic the 180 day gap around the introduction of the tax in August, 2004). Figure 2.3 shows the corresponding distributions. We find excess sales at the mansion tax threshold as the difference between total post-period sales and adjusted pre-period sales in the region \$990,000–\$999,999, and estimate the incidence as the price, p^* , at which the number of sales in the pre-period between the threshold and p^* is equal to the excess. We estimate the missing mass in the gap in the same way by taking the difference between total sales in the range \$1,000,000–\$1,155,422 in the pre and post periods. We find standard errors for these estimates by bootstrapping this procedure (including the growth factor for the pre-period) 999 times. Our incidence estimate using the pre/post comparison is slightly larger than the baseline estimate (\$25,000 vs. \$21,542).⁸⁰

The similarity of our estimate of \hat{h} using cross-sectional data (our baseline specification) and the pre/post comparison in New Jersey is reassuring, given that these sources of identification rely on different assumptions about the counterfactual. In the cross-sectional case, the counterfactual distribution is potentially distorted due to general equilibrium effects. In the pre/post comparison, the counterfactual is the distribution absent taxation. Similarity of the estimates of \hat{h} suggests that general equilibrium spillover to the distribution outside of the omitted region is not very important.

We find some heterogeneity by property vintage (years since construction) in estimates

⁸⁰We find little change in local incidence over time. Figure 2.15 in the appendix displays monthly incidence estimates for NJ. Prior to the introduction of the mansion tax in August 2004 estimates show no response to the threshold. Once the tax is introduced, prices quickly respond—incidence estimates reach the \$20,000 level within four months. We find no evidence that the response to the tax is changing with the housing boom and bust. We also see no obvious relationship between incidence and the real estate market in New York. Table 2.10 presents incidence estimates for all three geographies over time. While the estimates vary somewhat from year to year and region to region, we see no clear pattern over time and they all hover around our baseline estimates of \$20,000.

shown in Table 2.5, suggesting that some of the local response to the tax may be due to supply-side quality adjustments. We expect that negotiating a purchase of property before construction is finished allows for significant response in terms of the level of finish, appliances and other amenities, allowing for price reductions driven by adjustments in property characteristics. Similarly, older properties may require renovation and hence allow for quality to more readily respond to the tax. In contrast, original sales of apartments or houses after they have been constructed and finished may have less flexibility. Our data for New York City and New Jersey contain information about year of construction of the property. In particular, in New York City, which is dominated by large apartment buildings, there is a non-trivial number of sales that occur before construction is finished. For New York (but not New Jersey), we find that bunching is very large for sales before construction is complete and for sales that occur three or more years after construction. In contrast, sales that occur soon after construction—presumably original sales of already fully constructed and equipped properties—show smaller, but still significant (exceeding \$10,000 incidence estimate), bunching. Recall though that the introduction of the tax in New Jersey induced bunching immediately (Figures 2.4 and 2.14), so that investment-related adjustments are unlikely to explain the bulk of the response. We interpret these results as evidence that supply-side response along the quality/finish dimension is important: quality adjustments may perhaps explain around half of the price shift, still leaving local incidence of over 100%. While we do not pursue full welfare analysis, note that such a tax-motivated response corresponds to welfare loss by the same logic that implies that taxable income response reflects the efficiency cost of income tax (Saez et al., 2012).

Our estimates for the NYC RPTT and NJ RTF, reported in the first and second panels of Table 2.6, are smaller, but still consistent, with the mansion tax results. We find no evidence of response to the small (\$600) New Jersey RTF threshold. For the \$2125 RPTT notch that applies only in New York City, we find a response for new sales, but not for old sales (and we find no effect in the rest of New York State where the tax does not apply). The variability of

these results coincides with shifts in statutory incidence. Like the mansion tax, the RPTT on new sales is the responsibility of the buyer and we find evidence of a response that is consistent with the mansion tax estimates, albeit somewhat smaller: \$1758.225 represents an 82.7% local incidence on sellers of the \$2125 increase in tax liability. The RPTT on old sales and the NJ RTF schedule are the responsibility of the seller, and in none of these cases do we find any evidence of response. Both of these results are consistent with our reduced-form incidence estimates based on the mansion tax: a 100% burden of the tax on sellers should correspond to no price change for the sellers when they are the party with statutory incidence. Alternatively, switching incidence can correspond to changes in the salience of the tax—perhaps a tax imposed on sellers is less salient than the one imposed on the buyers.⁸¹

One concern is that the response to the tax may be driven by tax evasion, which we investigate by examining proxies for the availability of evasion opportunities. In the third panel of Table 2.6, we split the NYC sample by coop vs. condo status. Coop transactions have to be approved by coop boards that have the power to veto sales. In particular, there is anecdotal evidence that coop boards disapprove transactions that occur below the expected (or, perhaps, desired) market price. One might expect that if underreporting of the price is the important margin of response to the tax, the extent of bunching in coop apartments should be smaller than otherwise. This is indeed what we find, although the margin is small: \$16,292.354 for coops vs. \$23,292.602 for non-coops.

We also investigate a more direct proxy for evasion—the nature of the transaction—presented at the bottom of Table 2.6. All-cash transactions involve fewer parties (in particular, no financing) and more liquidity, which may increase the likelihood of side payments. We find the opposite: incidence of \$16,018 for cash transactions versus \$20,676.666 when the sale has an associated mortgage. In the context of real estate transactions, tax evasion

⁸¹For both the RTF and the RPTT there is a small dominated region where sellers would be better off accepting a lower sale price below the threshold, which is offset by a lower tax bill. Thus, there should be a small amount of bunching and a small gap at the RTF and RPTT thresholds when sellers remit. That we find no evidence of this is potential support of these taxes' limited salience. The dollar value of these discontinuities is small, however, so we refrain from drawing sharp conclusions on this point.

is certainly possible, but one might expect that it is not completely straightforward: both parties have to agree and money has to change hands at some point during the long closing process. Evasion in this context likely requires an aspect of trust between the two parties. Our New York State data contains a dummy for whether a transaction is “arms-length” (i.e., between related parties). We find no evidence that arms-length transactions involve more bunching (\$23,169 for arms-length sales, \$23,786 for non-arms-length).

Finally, examining listings data for Manhattan, we find comparable bunching in seller listing prices, suggesting that evasion is not a driving force of the observed sale-price response. Figure 2.7 shows the smoothed distribution of listing prices around the mansion tax threshold for properties sold and matched to the tax data.⁸² There are three prices shown for sold listings: the initial asking price, the final price in the listing data, and the sale price. Among properties that sell, bunching appears most prominent for the final asking price, followed by the sale price, and the initial asking price. These visual perceptions are confirmed by our estimates in Table 2.7 that find substantial bunching for both initial and final listing prices that actually exceed the response at the sale stage.⁸³ The response of the listing prices indicates that sellers internalize the presence of the tax (which is the responsibility of the buyer) even before meeting the buyer. Since these listings responses occur before the seller identifies a buyer who would be willing to engage in tax evasion, we find it unlikely that the ultimate sale price response is driven by such cheating.

The evidence so far shows clearly that notched transaction taxes distort the distribution of sale prices. We find some evidence of supply-side response in quality adjustments, as

⁸²Since bunching is more prominent in listings data, we adjust the distributions on the graphs to remove the common round number bunching. Specifically, we regress the log of the per-\$25,000-bin count on a cubic in price and dummy variables for multiples of \$50k and \$100k interacted with the price. We then subtract the predicted bunching effect from the actual counts. The remaining peaks in the data are the result of noise and do not necessarily correspond with salient round numbers. Figure 2.16 and Figure 2.17 in the appendix show the unadjusted distribution and the distribution of listing prices for all listed properties in Manhattan, respectively.

⁸³We do not find conclusive evidence of a similar response of listings prices to the NYC RPTT. This is consistent with the results that we discussed before: we find a response to the RPTT only for new sales where the tax applies to buyers. However, the number of new sales in the listing data is very small and we run into power issues.

well as differences in estimates based on the side of the market responsible for the tax. Our tests of tax evasion are weak, but do not suggest that this is the main force. These results reflect local reduced-form incidence estimates—the adjustment of prices in response to the threshold. By themselves, they do not reveal the strength of bargaining power and are not informative about the incidence of the tax away from the threshold. As discussed in our theoretical section, understanding the bargaining power and, relatedly, the possibility of unraveling in the market requires investigating the size of the gap in the distribution as well (reflected by \hat{Z}).

2.6 Unraveling: Market Distortions Local to the Threshold

Our estimates of \hat{Z} imply that there is significant unraveling of the market local to the tax threshold. In general, we find the number of sales bunching at the threshold to be smaller than the number of sales missing from the gap, translating into a positive value of \hat{Z} , as reported in Table 2.5. As discussed in Section 2.4.3, a positive sign on \hat{Z} cannot be reconciled with positive values of β (seller’s bargaining power) and instead indicates unraveling of the market in the proximity of the threshold. Hence, our results show that the threshold design of the tax discourages transactions that would have taken place if the tax rate was the same but discontinuity was not present—even after controlling for the usual extensive margin response (sales with positive surplus in absence of the tax, but negative surplus when taxed) we find that there are sales that do not occur. Moreover, the presence of unraveling suggests that the matching process is an important part of real estate sales and that this process may be disrupted by the tax.

To be specific, our baseline estimate of $\hat{Z} = \$43,861.766$ implies that the tax eliminates transactions that correspond to a range of original prices with a width between $\$43,861.766$ and $\$43,861.766 + \text{tax} \approx \$54,000$ above the threshold. For NYC, this corresponds to about 2,800 sales over the 2003-2011 period. Altogether, we find that the tax induces a \$20,000 price-range worth of transactions from above the threshold to bunch at the threshold, and

discourages another \$50,000 or so of sales from occurring at prices just above the threshold. The finding that \hat{Z} is substantial and positive is robust. Firstly, as can be seen in Table 2.5, \hat{Z} ranges from 37,409.87 to 43,861.766 across the three geographic regions that we consider (NYC, NYS, and NJ). Secondly, \hat{Z} is small and economically insignificant in NJ before the tax is introduced. Thirdly, as discussed in Appendix 2.9.4, the estimates of \hat{Z} are robust to the specification choices. Moreover, using data for NJ and constructing our counterfactual distribution using data from before the tax is introduced gives an estimate of \hat{Z} somewhat smaller than the baseline maximum likelihood estimate and with large standard errors, but with a point estimate that is still economically significant.

This extensive-margin response highlights an important margin of efficiency loss due to the transaction tax notch. As discussed above, a positive estimate of \hat{Z} suggests a very specific extensive-margin response, which we refer to as unraveling: some buyer-seller pairs who have a positive joint surplus under the tax (regions *B* and *C* of Figure 2.6) are exiting the market. This does not imply that these parties do not trade at all—buyers and sellers can continue to search and some of them may ultimately transact at different prices away from the bunching/gap region—but it provides evidence that the market in the region just above the threshold is unraveling. Note that this is different from the usual extensive margin response in which buyer-seller pairs who would transact in absence of the tax find that the tax reduces their joint surplus below zero and so the sale does not occur. Our estimation procedure explicitly controls for this traditional extensive-margin response by allowing for a level shift (discontinuity) in the distribution above the notch.

Examining real-estate listings data for New York City, we find suggestive evidence that the tax further disrupts the search process in a region above the notch. We interpret the presence of substantial bunching in the listings price (discussed above) as evidence that the tax influences seller search behavior. We also find that those who list just above the notch (between \$1M and \$1.075M) are still very likely to sell below one million (see the relationship between listing and sale price in Figure 2.18 in the Appendix), and are more likely to sell

than those who list below the notch or much higher above the notch (Figure 2.19) despite spending more time on the market (Figure 2.20). Interestingly, those who list just above the notch are more likely to leave their original REBNY realtor and sell with another realtor or on their own (Figure 2.21), suggesting an additional margin of adjustment to the tax—saving on realtor fees to compensate for a lower price.

The evidence thus far points to extensive distortions to price, unraveling of the market, and some disruption of the matching process local to the transfer tax notch. However, many of these responses occur because of how the tax is implemented—prices can adjust below the notch to avoid the tax and the notch creates specific local incentives for buyers and sellers to break matches. In what follows, we examine how the tax affects home sales more generally, away from the notch.

2.7 Global Market Distortions

In this section, we show evidence indicating that the transfer tax may distort the matching process everywhere above the notch. Conditional on initial listing price, sellers are taking larger discounts in the presence of the tax. While this higher discount could be explained by a shift in bargaining power or by endogenous listing prices, we find evidence that the efficiency of matches themselves is distorted by the tax. In particular, the variance of sale price conditional on property characteristics increases with the tax. While these results are largely descriptive—we must rely on observations below the threshold to form a counterfactual above—they do show sharp effects and, as discussed in Section 2.4.4, we interpret this higher variance in selling price as a decrease in the efficiency of the matching process.

Price discounts, which we define as the percent drop from listing to sale price, increase under the transfer tax. Figure 2.8 shows that the discount from the initial price to the final advertised price (i.e., before a buyer is identified) and to the final sale price increase as the initial listing price moves above \$1,000,000. We present the median and 75th percentiles of the distribution of discounts (many listings are not revised) in the figures. The effect is not

immediate at the \$1,000,000 threshold, because the tax applies to the sale price and not the initial price, and it is the latter that constitutes the running variable here: close to \$1 million a small discount is sufficient to bring the sale price below the notch. Interestingly, the increase is persistent well above the threshold—beyond \$1.1M larger discounts persist, even though these discounts are generally not large enough to move the transaction below the notch. We find analogous evidence for the discount from final to sale price—see Appendix Figure 2.22—suggesting that the price response is slowly revealed and reinforced throughout the search process by distorting the initial prices, subsequent revisions and, finally, during the bargaining stage.⁸⁴

We investigate the relationship between transfer taxes and price discounts more formally and find that the increase above the notch is significant and persistent. We regress the discount from initial asking price to sale price on a linear spline in initial asking price, with nodes at every multiple of \$100,000 (restricting the sample to listings with initial prices between \$500,000 and \$1,500,000). We follow the same procedure for the discount from initial asking price to final asking price. We plot the difference between the predicted discount (first price and final price) at each node and the predicted discount at \$1,000,000 in Figure 2.9. These estimates show a significant jump in the price discounts at the notch that shows no signs of reversing before reaching \$1.5 million.⁸⁵

There are several explanations for the increase in discounts. Firstly, it could be that sale prices above the notch are not changing, but the response is driven entirely by a change in asking prices (and we know from Section 2.5 that asking prices respond to the tax). In this case, since the final outcomes would be unchanged there could be little efficiency loss due to the tax. Secondly, it could be that the tax increases buyers' bargaining power. This may not entail any welfare loss due to the tax, but rather a redistribution of surplus from sellers to buyers. These two explanations still reveal that the tax affects behavior, although not the

⁸⁴This can also be seen in Figure 2.23, in which we focus on the mean discounts from the initial price that allow us to decompose the response (but blur the response that is predominantly present at high quantiles). Roughly half of the response is due to price revisions and half due to discounts at the bargaining stage.

⁸⁵ The corresponding point estimates (and related slope estimates) are in Appendix Table 2.11.

final outcomes. Alternatively, as we argue next, it could be that the tax disrupts the search process and reduces match quality with associated efficiency losses.

We find evidence consistent with transfer taxes disrupting the buyer-seller match process: conditional on seller characteristics, the tax increases the variance of sale prices. We investigate the relationship between asking prices and the variance of sale-prices with a two-stage spline estimation procedure. We first regress sale price on a linear spline in asking price (nodes at every \$100,000 between \$500,000 and \$1,500,000). We then estimate by median regression the relationship between the squared residuals from this first stage and a linear spline in asking price. We estimate standard errors using a clustered bootstrap procedure.⁸⁶

We find a significant increase in the variance of sale price as (initial or final) asking price crosses the \$1,000,000 notch. In Figure 2.10 we plot the predicted dispersion at the given node (from the two-stage spline procedure) relative to that at the notch.⁸⁷ This increased dispersion is pronounced and persistent well above the notch. The estimates show that, in general, the predicted variance below the notch is very close to the predicted variance at \$1,000,000, while the variance of sale price conditional on asking price is significantly higher above the notch than below the notch. Even as asking price rises to \$1,500,000, the variance of sale price does not return to pre-notch levels. This is inconsistent with asking prices and discounts simply scaling up without changes to other aspects of the matching process. As with the price discount, these estimates suggest that the real estate market is affected by the transfer tax even far above the notch—in the presence of the tax, asking price is a noisier signal of final sale price.⁸⁸

⁸⁶We use median regression in the second stage, because squaring the residuals makes these specifications sensitive to outliers. The bootstrap is as follows: for each observation, we resample first-stage residuals from the 50 nearest observations (by listing price). We use these residuals and the first-stage predicted values to construct a bootstrap sample and re-estimate the two-step process. We iterate this process 999 to acquire the distribution of estimates.

⁸⁷We present the numerical estimates of the difference between the predicted value at each node and the predicted value at \$1,000,000 in Table 2.12. The estimates are fairly consistent for both asking prices and are insensitive to the choice of quantile or mean regression in the first stage—see also Figure 2.10.

⁸⁸This increase in dispersion is also confirmed in the raw data. In Figures 2.24 and 2.25 we plot the variance of sale price conditional on initial asking price and final asking price, respectively. In both cases, the variance of sale price increases above the notch.

In Section 2.4.4 we argued that an increase in the dispersion in sale price conditional on seller characteristics implies a movement further away from the optimal allocation of assortative matching. An increase in the variance of sale price conditional on asking price is suggestive of such an efficiency loss, but it is not conclusive: the ideal measure is the variance of price conditional on seller type. Asking price is an imperfect proxy for seller type—if asking price is endogenous to the tax, the increase in variance may be driven by a change in the composition of sellers who list at each price.

We repeat the two-step estimates of sale-price variance versus asking price including controls for property characteristics, in order to better approximate the variance of sale price conditional on seller type. In both stages we control for year-of-sale fixed effects, zip code fixed effects, building type (single-family home, multi-family home, apartment in walkup, apartment in elevator building, etc.), whether the sale is of a new unit, and the log of years since construction (plus an indicator for missing years since construction). We plot the results from this procedure in Figure 2.10. Including these controls somewhat reduces the effect of the tax on the variance of sale price, suggesting that endogeneity of listing prices may have influenced the previous estimates, but the difference is not large and generally not statistically significant. In particular, there is still a significant increase in the dispersion of sale price above the notch. We interpret this increase as evidence that the transfer tax reduces efficiency of the housing market by disrupting the matching process.

2.8 Conclusions

Our empirical analysis demonstrates how a notched transaction tax affects the distribution of housing sales by creating bunching below the notch, a large gap in the distribution above the notch, and distorts the relationship between asking price and sale price. The price responses that we identify suggest that sellers local to the threshold take large price cuts—greater than the cost of the tax—although this may partially be driven by quality adjustments. The finding of a gap above the notch that is larger than the number of sales bunching at the

threshold—even after controlling for the usual extensive-margin effect—suggests that there are productive transactions that do not occur because of the notched tax. We show that there is bunching in the distribution of residential listing prices and an effect of the tax on discounts, both of which indicate that not just the final price, but also some aspects of the search process, are affected by the tax. Finally, we find that in the presence of transaction taxes the relationship between seller asking price and final sale price weakens, and that this persists when we control for property characteristics. This effect extends far above the threshold. We interpret this increase in price dispersion as a movement away from the efficient allocation of positive assortative matching in the housing market and as revealing the distortion due to taxation globally. Hence, we conclude that our analysis of real-estate transaction tax notches reveals substantial price response local to tax thresholds, that a notched tax crowds out productive transactions, and that transaction taxes may increase search-related inefficiencies.

That we find substantial distortions due to the design of the transaction taxes in NY and NJ (0.7% of all transactions are eliminated due to the threshold) raises the question of why notched taxes are present in the first place. One of the few general results of the optimal income tax literature is that the marginal tax rate should be lower than 100% (no notches). Indeed, our bunching estimates show that the transfer tax threshold leads to a substantial response and loss of revenue. Moreover, unraveling adds efficiency and revenue costs. Phasing in the mansion tax, which is likely administratively feasible, would eliminate these inefficiencies. Still, phasing in would not eliminate all distortions: an extensive margin response may be present and one of our findings is that search is affected even far from the threshold. One might ask, then, why have a progressive transfer tax in the first place. Public finance theory, going back to the Atkinson-Stiglitz result, suggests that such an instrument is useful if it can push redistribution beyond what can be achieved via an income tax (a tool which New York State, New York City, and New Jersey all have in place). The real estate market raises an additional issue though: some investors may be non-residents,

so that resources invested in real estate need not be reachable by income taxes. Still, while a transfer tax might indeed hit non-residents, the same would be true about general property taxes, which avoid many of the inefficiencies we identify herein. One redeeming quality of the transaction tax is that it will disproportionately affect more frequently traded properties. This feature is appealing if frequency of trading varies for different groups of purchasers.

Our evidence of exit from markets near the tax threshold raises important issues for implementation and interpretation of studies relying on bunching at notches for identification. To date, most of this literature has assumed one-sided markets and abstracted from extensive-margin responses. Our framework highlights that there are two different types of exit from markets. One is standard—transactions with low surplus do not take place. In the income tax context, this is akin to the labor force participation decision.⁸⁹ The second type of exit, which has not been previously recognized in this literature, is unraveling of the market near the notch that corresponds to destroying productive matches. Both of these responses imply that bunching does not fully characterize the consequences of a notch. While in some contexts (e.g., an income tax notch at a high value) it is reasonable to ignore the first type of extensive-margin response, the second type of response is likely to be intrinsic to any matching context in which parties have an option to continue search (e.g., firms may continue searching for a worker willing to accept a wage below the notch; employees may make different occupational choices). The possibility of this second type of response undermines the assumption that the excess mass bunching at the threshold is identical to the missing mass in the gap (as in Kleven and Waseem, 2013). Our empirical framework relaxes this assumption and allows us to explicitly test for exit local to the notch.

More generally, our results suggest that taxes may introduce inefficiencies into other search markets. Labor markets are, perhaps, the most obvious example of a market with non-trivial search costs where matches are subject to taxation. Many labor-market regulations follow notched designs, such as notched income and wage taxation or requirements to provide

⁸⁹See Marx, 2013 for explicit modeling of this decision to exit the market.

health insurance or comply with Value Added Tax if the number of employees crosses a given threshold. That firms might face such discrete costs to increasing scale may lead not simply to supply- or demand-side adjustments, but perhaps to the destruction of equilibrium opportunities that require costly search—in the same way that productive real-estate matches are discouraged near the transaction tax notch. Moreover, our finding that transaction taxes increase search-related inefficiencies well above the threshold suggests that even non-notched policies may lead to less efficient worker-firm matches. Of course, the housing market differs from the labor market in ways that may make the mechanisms we study particularly pronounced. Firstly, the availability of alternative margins of adjustment affects the ease of moving below the tax threshold. Secondly, the housing market is a “spot” market where the current price determines tax treatment; this is not necessarily the case in labor markets, where contracts may be long lasting and as earnings adjust over time they move further away from the notch. Nonetheless, our results underscore the importance of considering how taxation affects search, especially in the context of policies that follow a threshold design.

Table 2.1: Real Estate Transfer Tax Schedules

Tax	Threshold (\$)	Rate Below	Rate Above	Jump	Statutory Incidence
Mansion Tax (NY & NJ)	\$1,000,000	0%	1%	\$10,000	Buyer
RPTT (NYC, Residential)	\$500,000	1%	1.425%	\$2,125	Seller*
RPTT (NYC, Commercial)	\$500,000	1.425%	2.625%	\$6,000	Seller
RTF (NJ)	\$350,000	0.78%**	0.96%	\$630	Seller

Notes: *Buyer remits tax if the sale is of newly developed property (otherwise Seller remits). **NJ RTF schedule features non-linear tax schedule below \$350,000, all of which changes when the sale price crosses the notch; 0.78% and 0.96% are simply the marginal rates faced above and below the notch.

Table 2.2: Sample Statistics for Taxable Sales

	NYC (2003 – 2011)	NYS (2002 – 2010)	NJ (1996 – 2011)
Number of Sales	380629	1172708	1703260
Sales ∈ (\$950K, \$1.05M)	7932	6242	7556
Median Price	405600	159900	200000
Mean Price	660719	258363	262122

Notes: NYC data is from the Department of Finance Rolling Sales file for 2003-2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos). Data for NYS from Office of Real Property Service deeds records for 2002-2006 and 2008-2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes). Data for NJ from the State Treasury SR1A file for 1996-2011 (taxable defined as any residential sale).

Table 2.3: Median Price of Taxable Sales Over Time

Year	NYC		NYS		NJ	
	<i>n</i>	Price	<i>n</i>	Price	<i>n</i>	Price
1996111759127000	
1997115470130000	
1998131485137500	
1999139167143000	
2000136891151000	
2001136733169000	
2002	.	.163491132000	145718197000			
2003	47679293000	167709149500	148906235000			
2004	53342340000	175766165000	159220270155			
2005	52310395460	175873184640	155340315000			
2006	47973445000	152220170000	127630327000			
2007	48552480000	5934287000	108790321050			
2008	40354475000	117000162500	86151288000			
2009	31368420000	110408155500	83407257500			
2010	27132463000	104307160666	80944250000			
2011	131919456000	.	.28766250000			

Notes: NYC data is from the Department of Finance Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos). Data for NYS from the Office of Real Property Services deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes). NYS observations in 2007 are from sales made in 2007, but recorded in 2008–2011 and omits sales recorded in 2007. Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale).

Table 2.4: REBNY Listings Sample Statistics

All Listings						
	Sold but not Matched,	Days on Market	Initial Asking Price	Final Asking Price	Discount (First to Final)	
Mean	0.671	0.103	197.859	1604547	1602670	0.019
Median	1.000	0.000	110.000	899000	875000	0.000
n	71875	71875	67550	71875	71875	71875

Closed and Matched Listings						
	Days on Market	Initial Asking Price	Final Asking Price	Sale Price	Discount (First to Final)	Discount (Final to Sale)
Mean	146.107	1384028	1435747	1241209	0.032	0.059
Median	80.000	825000	799000	784052	0.000	0.043
n	40680	44320	44320	44193	43506	43309

Notes: Data from the Real Estate Board of New York's listing service; represents all REBNY listings in Manhattan between 2003 and 2010 that are closed or off market. Sold is an indicator equal to one if the final status of the listing is "Closed." Days on the market is calculated as the number of days between the initial active listing and the final status of "in contract" (if the property sells with REBNY) or "permanently off market" (otherwise). Initial asking price is the asking price on the listing when first active; final asking price is the price listed immediately prior to the listing being "in contract" or being taken off the market (if unsold). Sale price is the price reported in the NYC DOF data, and is available only for REBNY listings that have a match in the DOF data (sale price of 0 is considered missing). Discount is defined as $1 - \frac{\text{final price}}{\text{first price}}$ and is winsorized at the 1st and 99th percentiles. 5,940 listings have invalid listing and off-market dates (missing or obviously misreported), and are omitted from days on market calculations. "Matched, but not closed" is an indicator that a listing has a match in the NYC DOF data, but is never reported as "Closed" by the REBNY agent.

Table 2.5: Response to Mansion Tax, by Region and Years Since Construction

Sample	Incidence	Std. Error	\hat{Z}	Std. Error	n	
NYC	21542.098	1150.878	43861.766	4142.953	102493	
NYS (excl. NYC)	23227.515	1084.482	41610.588	4334.170	108462	
NJ Post Tax	21477.388	1474.300	37409.873	4310.896	111936	
NJ Pre Tax	-784.065	38.892	2958.261	285.872	57836	
NJ Pre/Post Comparison [†]	25000.000	8515.132	14223.330	11628.94	2020	
	<0	37329.701	16009.929	-13709.572	42837.004	559
	0	13759.258	7671.441	-2451.076	27109.852	1048
NYC	1	11309.339	3457.124	45053.550	15619.853	3422
(Yrs. Since Contr.)	2	14118.294	3145.743	47311.896	15259.170	4388
	3	24467.069	5927.603	36654.826	26001.895	2253
	4-6	25586.634	6045.138	82880.894	29444.354	2433
	7+	22780.508	1275.592	48877.574	5193.370	72128
NYS (excl. NYC)	Old	22677.619	1283.443	40945.774	4264.702	104576
	New	34254.287	8350.604	62173.984	25965.831	3886
	0	24949.365	10542.512	19598.806	33121.546	988
	1	24684.039	7485.882	55501.171	26732.423	1896
NJ Post Tax	2	23730.361	7183.484	54344.928	25413.421	1773
(Yrs. Since Constr.)	3	24950.530	8043.393	25353.913	25917.361	1882
	4-6	19718.802	4392.694	41103.750	17066.774	6148
	7+	19967.636	1633.977	40047.784	4989.520	85551
	0	-351.083	4149.632	1619.879	26316.982	723
	1	-142.337	6362.861	4103.325	71219.647	852
NJ Pre Tax	2	-798.644	1743.317	3214.845	13783.929	1082
(Yrs. Since Constr.)	3	-2121.267	2519.545	8725.490	10998.073	1249
	4-6	-661.898	537.723	9250.812	18873.573	3335
	7+	-778.278	46.481	3226.236	380.842	36530

Notes: Estimates from baseline procedure for given data source/sample (3rd-order polynomial, omit \$990k-\$1,155,422).

[†]Estimates for NJ Pre/Post comparison using NJ taxable sales omitting 90 days around the implementation of the policy: from Oct. 30, 2004 to Oct. 29, 2005 (post-period) and May 3, 2003 to May 2, 2004 (pre-period). Incidence estimate is the price at which the number of sales in the pre-period to the right of the threshold equal the difference between the number of sales in the bunching region (\$990,000-\$999,999) in the post- and adjusted pre-periods—pre-period distribution is adjusted as described in text. \hat{Z} calculated as in the text, where the excess number of sales bunching at the gap is the difference between the post- and adjusted pre-period distributions in the bunching region, while the gap is calculated as the difference between the distributions above the notch (\$1M-\$1,155,422). Number of observations listed for the pre/post comparison is the total count of taxable sales between \$990,000 and \$1,155,422 for these dates.

Table 2.6: Heterogeneity in Response by Notch and Sub-Sample

Geog.	Sample	Incidence	Std. Error	\hat{Z}	Std. Error	n
RPTT						
NYC	New	1758.225	751.923	3071.659	1975.290	21683
	Old	-390.461	28.339	488.460	69.80225	59840
NYS	New	-679.032	510.129	7.829	332.599	19147
	Old	-889.542	18.945	1105.190	56.53168	7807
RTF						
NJ	Post Aug. 2004	-699.552	15.333	-268.101	44.13154	6882
	Pre Aug. 2004	-591.400	20.099	-300.348	68.34483	6832
Mansion Tax						
NYC	All Coops	16292.354	2184.555	58245.269	8590.430	26950
	All Non-Coops	23292.602	1027.253	38817.535	4612.629	75543
	Old Non-Coops	24196.113	535.699	37340.124	5067.660	63971
NYC (Deeds)	Cash Only	16018.361	2356.852	74976.653	8814.378	28339
	Mortgage	20676.666	2223.237	119838.215	11273.628	49421
NYS	Arms-Length	23168.557	1163.640	43274.495	4606.535	97936
	Non-Arms-Length	23786.484	3149.902	24496.898	12693.600	10526

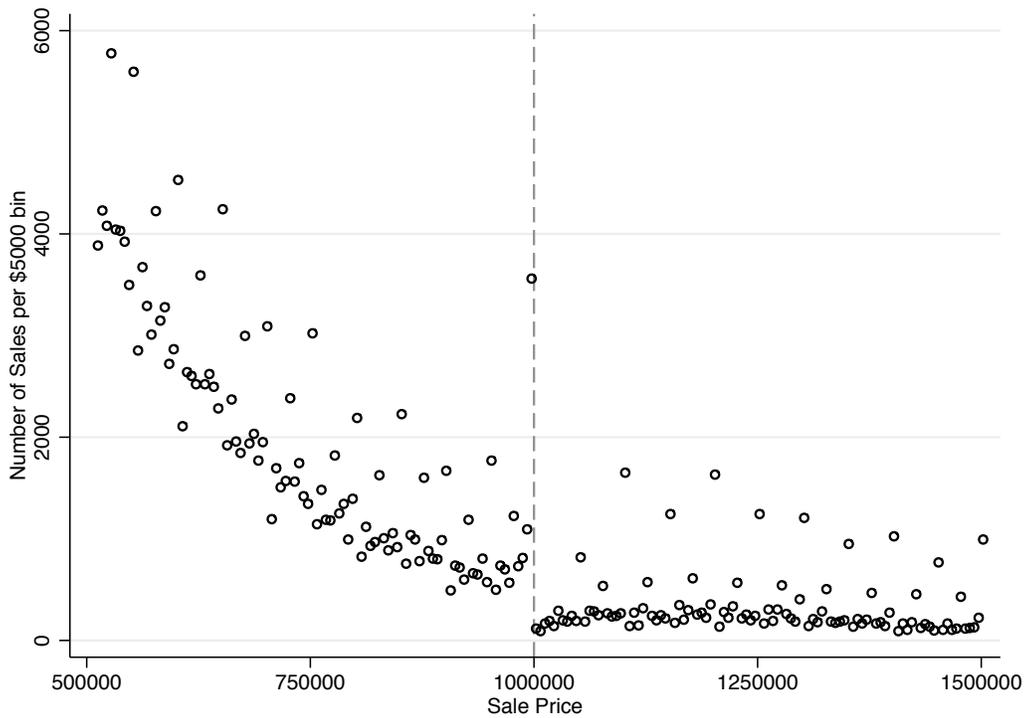
Notes: RPTT and RTF estimates based on 5th-order polynomial, omitting sales between \$490,000 and \$550,000 (for NYC and NYS) and between \$340,000 and \$400,000 (for NJ), Mansion tax estimates as in the baseline specification. NYC data is from the Department of Finance Rolling Sales file for 2003–2011. Data for NYS from the Office of Real Property Services deeds records for 2002–2006 and 2008–2010 (excluding NYC) restricted to all single-parcel residential sales of one-, two-, or three-family homes. Sales in NYC are defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos. New sales are defined as any sale occurring within three years of unit’s construction (in NYC) or any sale flagged as new construction (in NYS, excluding NYC). Data for NJ from the State Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). Coops are identified in the rolling sales data as sales with associated building codes equal to “Coops - Walkup Apartments” or “Coops - Elevator Apartments.” NYC Deeds Records data from deeds records collected by private data provider (taxable defined as any residential sale). Non-arms-length sales in NYS defined by the Office of Real Property Services as a sale of real property between relatives or former relatives, related companies or partners in business, where one of the buyers is also a seller, or “other unusual factors affecting sale price” (ex. divorce or bankruptcy).

Table 2.7: Mansion Tax: Listings

Sample	Price	Incidence	Std. Error	\hat{Z}	Std. Error	n
All	First	24443.666	166.150	101702.509	10589.033	36232
	Final	34992.894	3380.599	76606.311	9122.157	35714
Sold	First	24363.369	353.901	95113.750	12794.903	25112
	Final	38445.445	4289.510	68828.242	10640.681	24755
Unsold	Sale	19148.096	2089.630	53521.404	8077.065	24474
	First	24700.339	2638.516	113962.040	21381.710	7612
	Final	32926.816	6670.229	94839.590	19962.747	7539

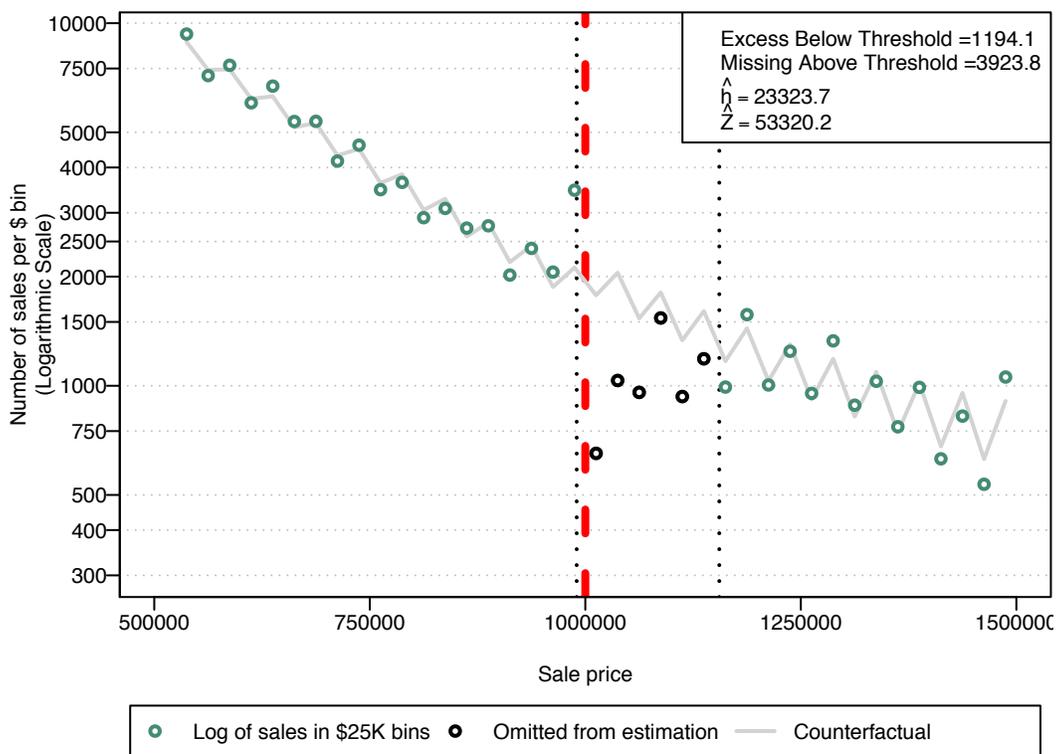
Notes: Data from the Real Estate Board of New York's listing service; represents all REBNY listings between 2003 and 2010 that are closed or off market. Unsold sample defined as all listings with final status not equal to "closed." Sold sample defined as all listings that match to a NYC Department of Finance sale record with final status equal to "closed." First price is the initial price posted on the listing. Final price is the last price posted while the listing is active (prior to status being changed to "in contract" or "off market"). Sale price is the recorded price from the NYC Department of Finance.

Figure 2.1: Distribution of Taxable Sales in New York State



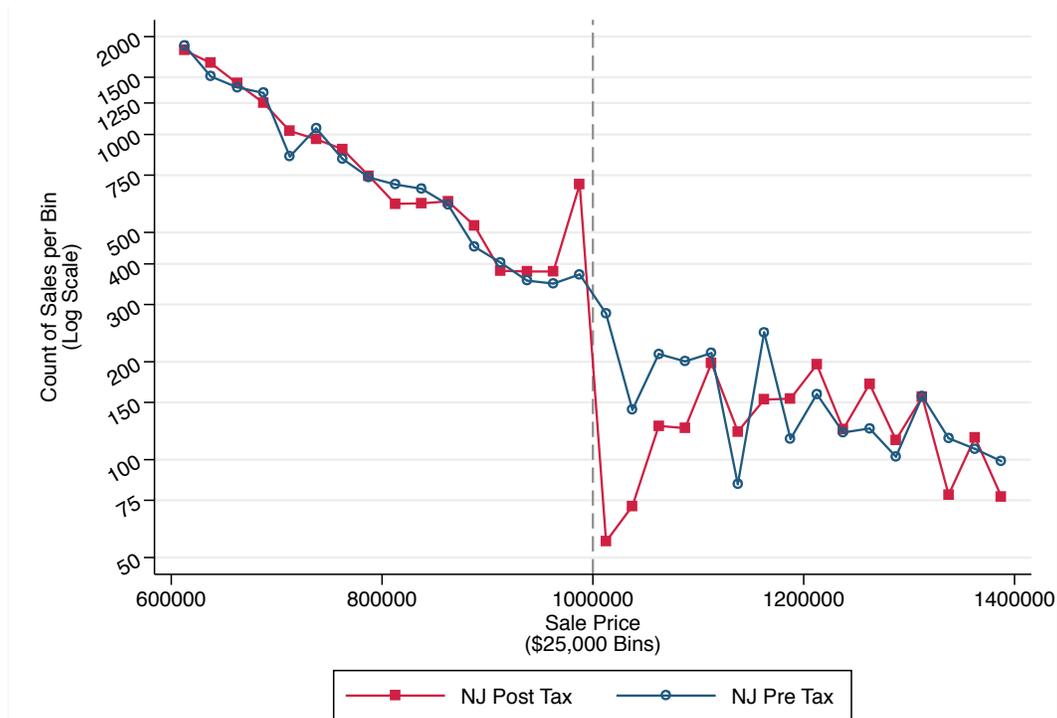
Notes: Plot of the number of mansion-tax eligible sales in each \$5,000 price bin between \$510,000 and \$1,500,000. Data from the NYC Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family properties) and from the N.Y. State Office of Real Property Service deeds records for 2002–2006 and 2008–2010 (taxable defined as single-parcel residential sales of one-, two-, or three-family homes).

Figure 2.2: Distribution of Taxable Sales in New York City



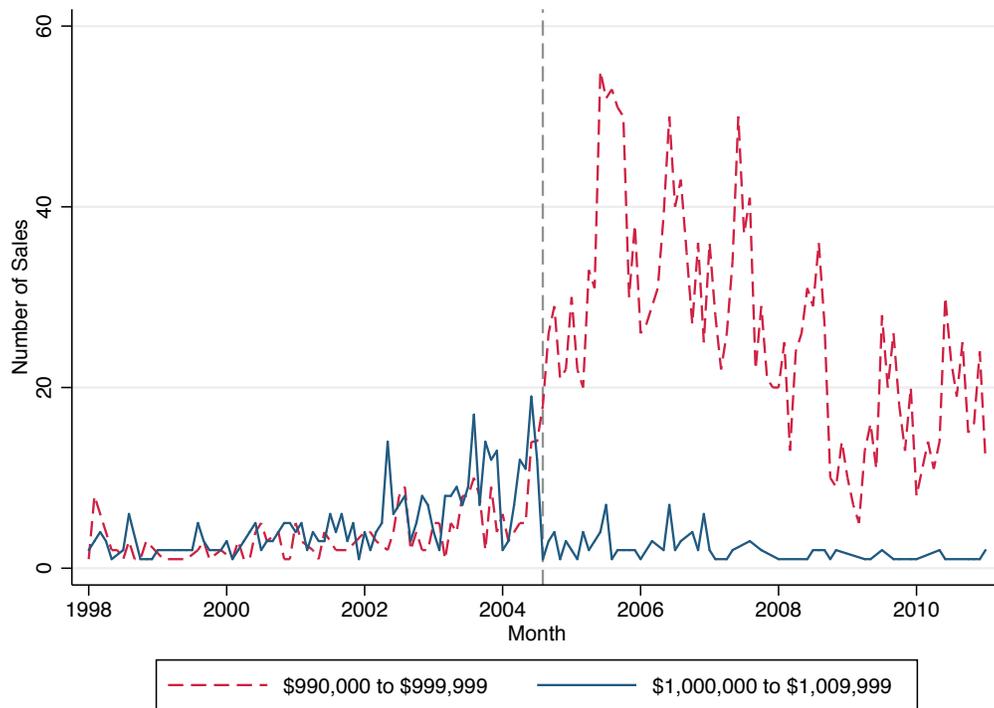
Notes: Plot of the number of mansion-tax eligible sales in each \$25,000 price bin between \$510,000 and \$1,500,000. Data from the NYC Rolling Sales file for 2003–2011. Fit corresponds to the baseline specification in Table 2.5.

Figure 2.3: Distribution of NJ Sales Pre- and Post-Mansion Tax



Notes: Plot of the number of mansion-tax eligible sales in each \$25,000 price bin between \$510,000 and \$1,500,000 before and after the introduction of the tax. Data from NJ Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). We implement this pre/post comparison as follows. We omit transactions within 90 days of the policy change (to avoid the retiming response) and focus on the following year (Oct. 30, 2004–Oct. 29, 2005). We rescale the period before the tax (May 3, 2003 to May 2, 2004) to account for sales growth over time. Specifically, we construct a counterfactual growth factor by taking the ratio of the count of sales within \$2500 of each price from May 3, 2002 to May 2, 2003 to the count of sales from Nov. 5, 2000 through Nov. 4, 2001 (omitting sales between Nov. 2001 and May 2002 to mimic the 180 day gap around the introduction of the tax in August, 2004).

Figure 2.4: New Jersey Monthly Sales Above \$990,000



Notes: Total taxable NJ sales in given price range by month. Data from NJ Treasury SR1A file for 1998–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (denoted by gray dashed line).

Figure 2.5: Incidence and Gap Concepts

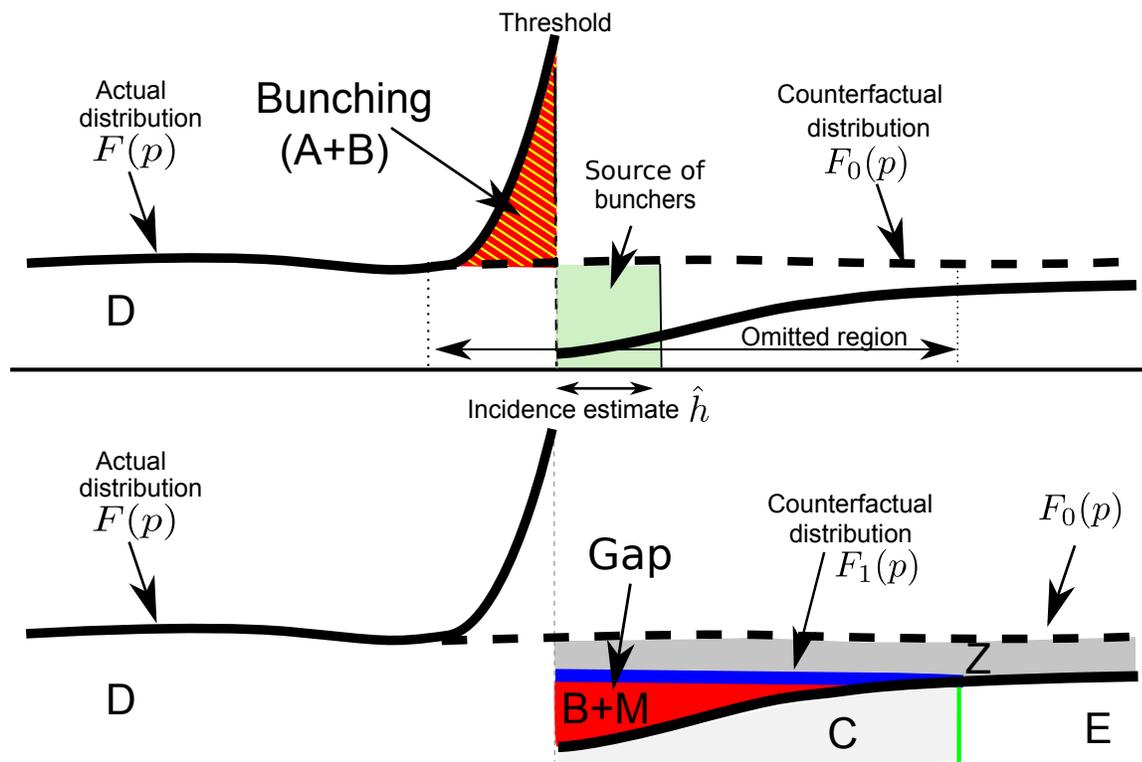


Figure 2.6: Bunching at the notch and efficient allocation

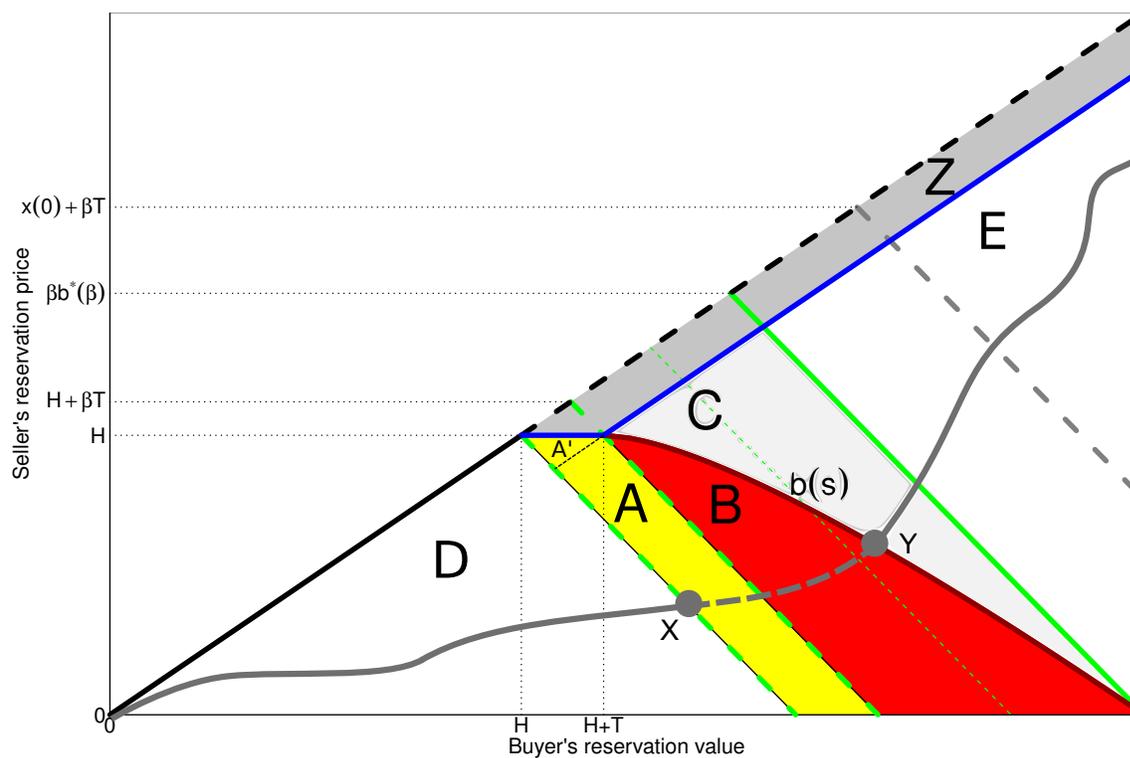
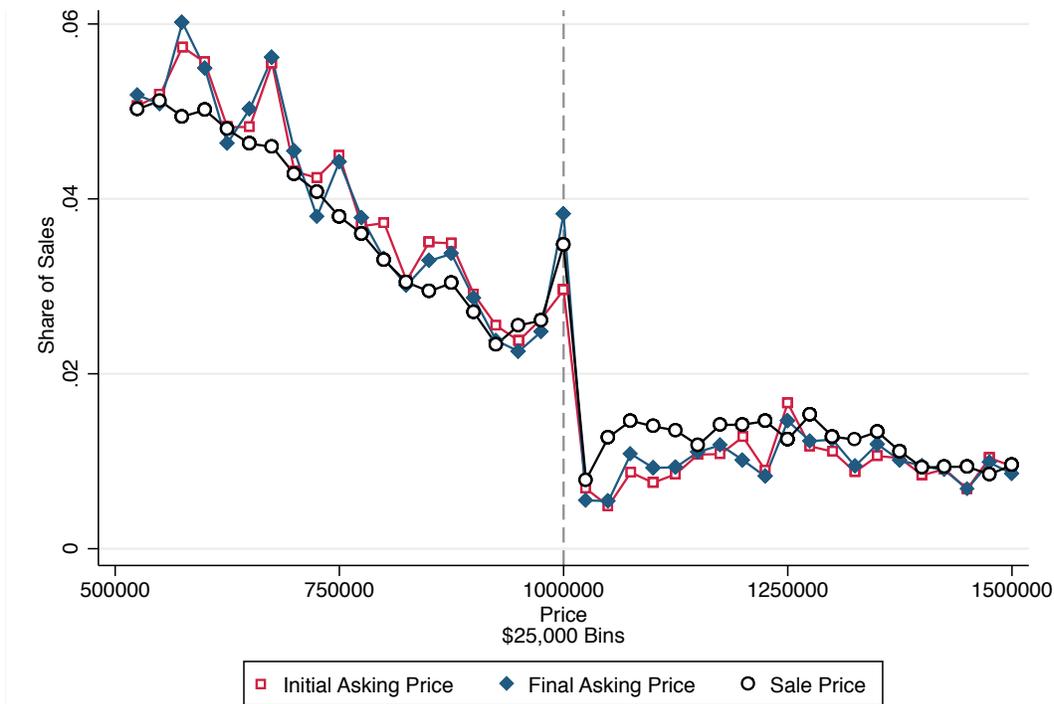
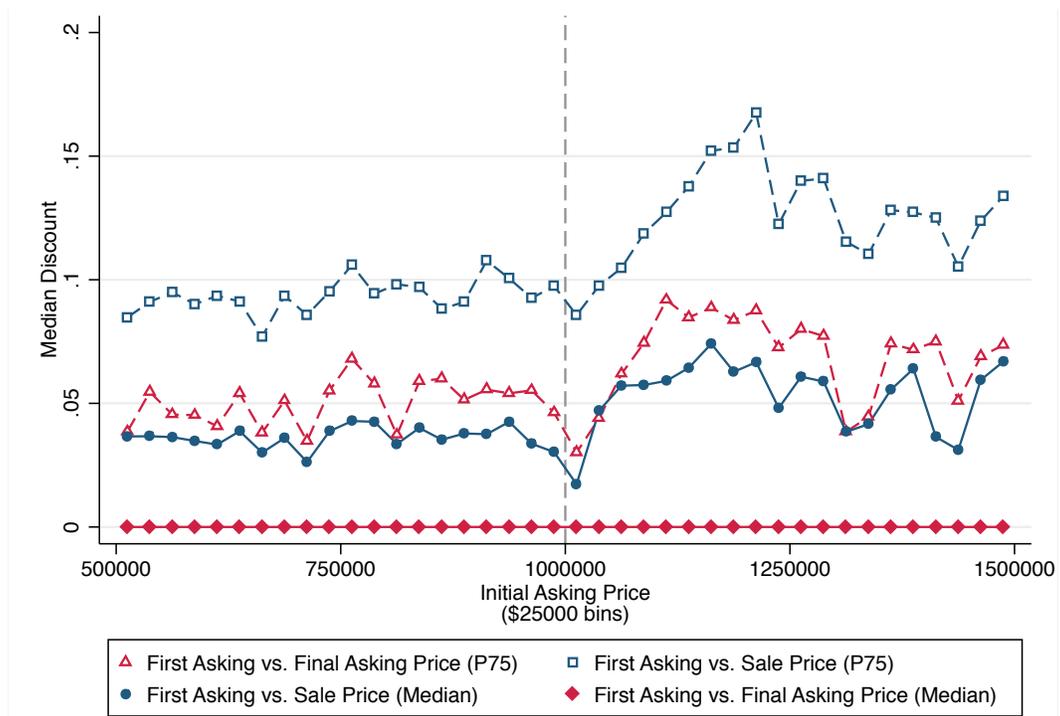


Figure 2.7: Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only)



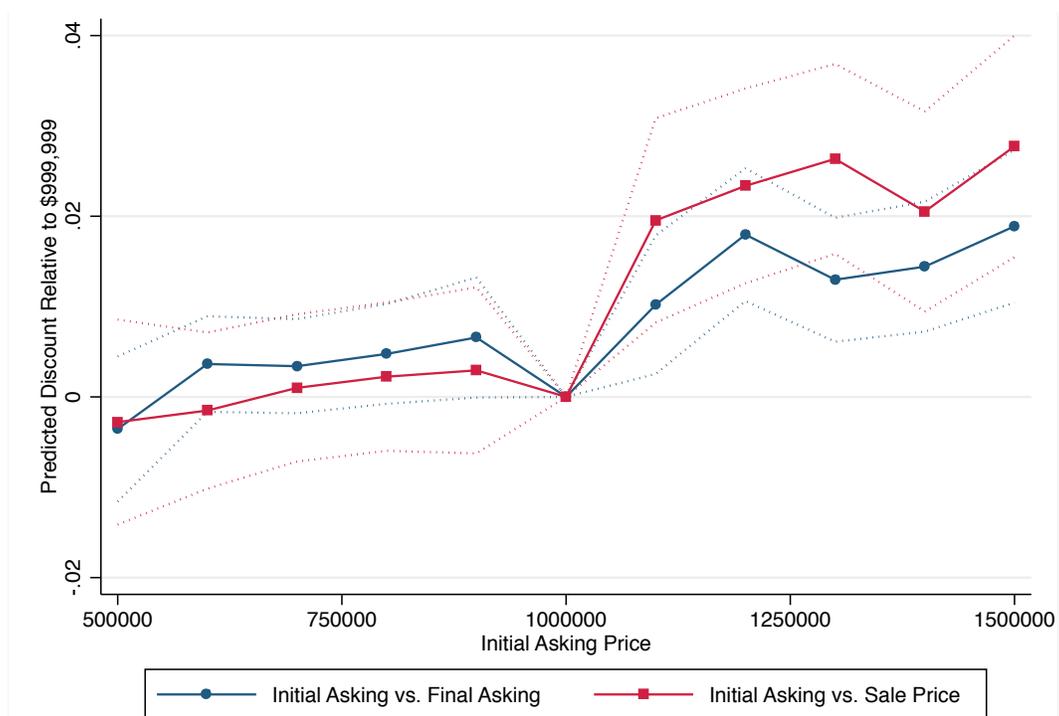
Notes: Data from REBNY listings matched to NYC Department of Finance sales records. Sample restricted to “sold” listings: last listing status is “closed” and property can be matched to NYC sales data. Smoothed plot of the distribution that accounts for round-number bunching. The log of the per-\$25,000-bin counts are regressed on a cubic in price and dummy variables for multiples of \$50,000 and \$100,000 interacted with the price. Predicted bunching for round-number bins are then subtracted from the corresponding counts.

Figure 2.8: Median & 75th Percentile Price Discounts by Initial Asking Price



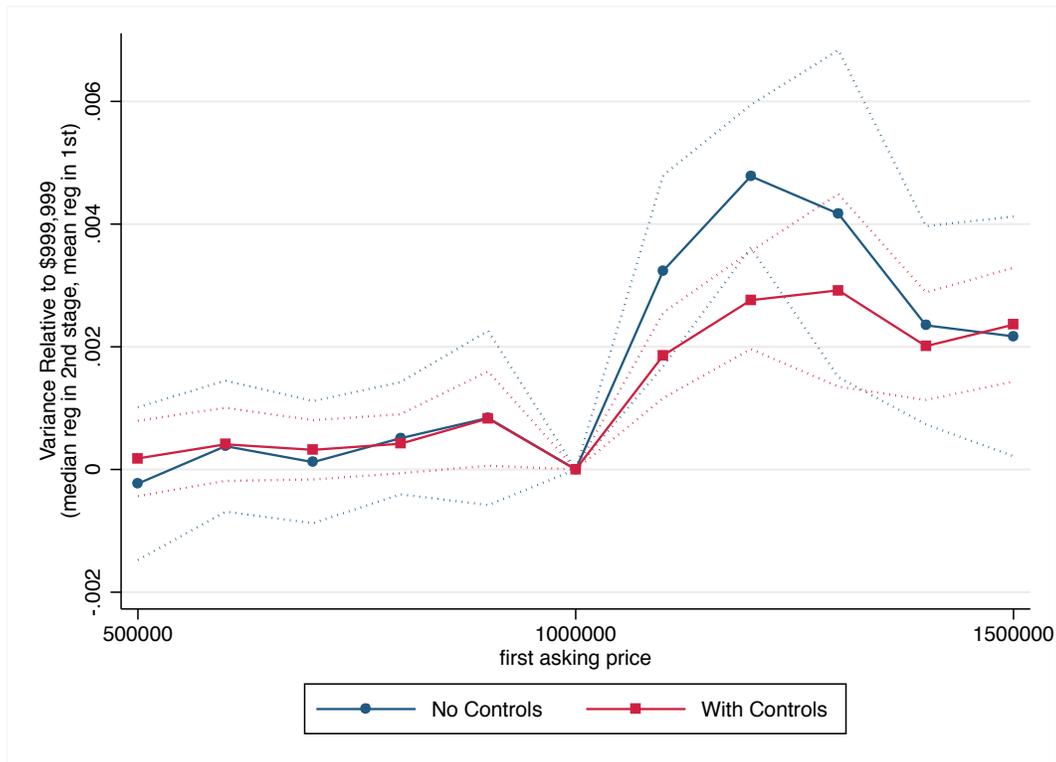
Note: Plot of the median and 75th percentile discount from initial asking to sale price ($= 1 - \text{final}/\text{initial}$) and initial asking to final asking price ($= 1 - \text{sale}/\text{initial}$) per \$25,000 initial-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range \$500,000–1,500,000 that match to NYC DOF data.

Figure 2.9: Predicted Price Discounts by Initial Asking Price (Relative to \$1,000,000)



Note: Plot of difference between predicted discounts by initial asking price and prediction at \$1,000,000 from regressions of discount (first to final price or first to sale price) on a linear spline in initial asking price with \$100K knots between \$500,000 and \$1,500,000. Dashed lines represent 95% confidence intervals.

Figure 2.10: Predicted Dispersion of Log of Sale Price by Initial Asking Price (Relative to \$1,000,000)



Notes: Plots of the difference between predicted values at given initial asking price and predicted value at \$1,000,000 from the following procedure. The log of sale price is regressed on a linear spline in the log of initial asking price with \$100,000 knots between \$500,000 and \$1,500,000. Squared residuals from this first stage are then regressed on a linear spline in log of initial asking price (using median regression; results are sensitive to outliers). Controls in the indicated results include year of sale, zip code, building type, whether the sale is of a new unit, and the log of years since construction. Dashed lines represent 95% confidence intervals from 999 wild bootstrap replications of the two-stage procedure, resampling residuals in the first stage by asking-price clusters.

2.9 Appendix

2.9.1 Bargaining model

To identify buyer-seller pairs that move to the notch, we compare maximized surplus above the notch to surplus when price is at the notch. Surplus at the notch is given by $\beta \ln(H - s) + (1 - \beta) \ln(b - H)$, while maximized surplus when the tax is due is given by $\beta \ln(\beta(b - s - T)) + (1 - \beta) \ln((1 - \beta)(b - s - T))$. Transactions in this category that sell at the threshold satisfy:

$$f(b, s; \beta) \equiv \beta \ln(H - s) + (1 - \beta) \ln(b - H) - \beta \ln(\beta) - (1 - \beta) \ln(1 - \beta) - \ln(b - s - T) \geq 0$$

$$(1 - \beta)s + \beta b \geq H + \beta T \text{ and } b - T - s \geq 0, b \geq H + T, s \leq H$$

at the notch. That is, the surplus at the threshold has to be higher than under the alternative of selling with the tax, and the non-trivial case corresponds to transactions with positive surplus that could otherwise sell at a price higher than the threshold. We show the following results:

Lemma 1. Fix $0 < \beta < 1$ and consider matches (b, s) that satisfy $p(b, s) \geq H + \beta T$, $b - T - s \geq 0$, $b \geq H + T$, and $0 \leq s \leq H$

(a) For any value of $0 \leq s \leq H$ there exists $b(s) > -\frac{(1-\beta)}{\beta}s + \frac{H+\beta T}{\beta}$ such that $f(b(s), s; \beta) = 0$

(b) Matches (b, s) that satisfy $b \in \left[-\frac{(1-\beta)}{\beta}s + \frac{H+\beta T}{\beta}, b(s)\right]$ locate at the notch and those with $b > b(s)$ sell with the tax at $p(b, s) - \beta T$

(c) Matches (b', s') that would otherwise sell at the same price as $(b(s), s)$ (i.e., $p(b', s') = p(b(s), s)$) sell at the notch when $s' \leq s$ and sell with the tax otherwise

(d) $b(0)$ is finite, and matches (b, s) that absent the tax would sell at prices higher than the corresponding price $p(b(0), 0) = \beta b(0)$ will never bunch.

Proof. For part (a) and (b) note that for a transaction that would otherwise sell exactly at $H + \beta T$ the notch is preferred so that $f\left(-\frac{(1-\beta)}{\beta}s + \frac{H+\beta T}{\beta}, s; \beta\right) \geq 0$; that $\frac{\partial f}{\partial b} = \frac{1-\beta}{b-H} - \frac{1}{b-s-T} = -\frac{H+\beta T - (\beta b + (1-\beta)s) - T}{(b-H)(b-s-T)} < 0$ because $(1 - \beta)s + \beta b \geq H + \beta T$; and finally that $\lim_{b \rightarrow \infty} f(b, s; \beta) = \beta \ln(H - s) - \beta \ln(\beta) - (1 - \beta) \ln(1 - \beta)$

$\beta) + (1 - \beta) \ln\left(\frac{b-H}{b-s-T}\right) - \beta \ln(b-s-T) = -\infty$ because all but last term converge to finite values as b increases. Hence, for each $0 \leq s \leq H$, there is $b(s)$ that solves $f(b(s), s; \beta) = 0$, and $b(s)$ separates positive from negative values of $f(b, s; \beta)$.

Part (c): to evaluate the effect of a change in s holding $p(b, s)$ constant, substitute $b' = -\frac{1-\beta}{\beta} s' + \frac{p(b(s), s)}{\beta}$ into $f(b, s; \beta)$ and totally differentiate with respect to s to obtain $\frac{df}{ds} = -\frac{\beta}{H-s} - \frac{(1-\beta)^2}{\beta(b-H)} + \frac{1}{\beta} \frac{1}{b-s-T} = \frac{1}{\beta} \left(\frac{1}{b-s-T} - \frac{\beta}{(H-s)/\beta} - \frac{1-\beta}{(b-H)/(1-\beta)} \right)$. Note that convexity of $\frac{1}{x}$ implies that $\frac{\beta}{(H-s)/\beta} + \frac{1-\beta}{(b-H)/(1-\beta)} \leq \frac{1}{\beta(H-s)/\beta + (1-\beta)(b-H)/(1-\beta)} = \frac{1}{b-s}$ and since $\frac{1}{b-s} < \frac{1}{b-s-T}$ we have $\frac{df}{ds} > 0$.

Part (d): finiteness of $b(0)$ follows from part (a). When $p(b, s) > \beta b(0)$, then there is $b' < b$ such that $p(b', s) = \beta b(0) = p(b(0), 0)$. Part (c) implies that (b', s) sells with the tax because $s \geq 0$. Because $b > b'$, $\frac{\partial f}{\partial b} < 0$ then implies that (b, s) also does not locate at the notch. \square

The lemma establishes the existence and shape of the schedule $b(s)$, which is marked by a solid red line in Figure 2.6. Given the seller reservation value, matches with buyers to the left of this schedule bunch at the notch, while those above it sell with the tax. Parts (a) and (b) of the lemma establish that $b(s)$ exists and is unique for any s lower than the threshold. Part (c) shows that the slope of $b(s)$ is flatter than that of the constant-price schedules, and Part (d) shows that the schedule $b(s)$ intersects the horizontal axis at some finite value so that for sufficiently high original prices transactions will never bunch. We re-state this last observation in the following corollary:

Corollary 2. *Transactions at the notch satisfy $H \leq p(b, s) \leq \beta b^*(\beta) < \infty$ where $b^*(\beta)$ is defined as $f(b^*(\beta), 0; \beta) = 0$ or, explicitly, it solves $\beta \ln(H) + (1-\beta) \ln(b-H) - \beta \ln(\beta) - (1-\beta) \ln(1-\beta) - \ln(b-T) = 0$.*

The corollary follows from part (d) of the lemma. As the original price increases, the attractiveness of the notch declines so that only transactions with sufficiently high overall surplus (sufficiently low reservation price of the seller) continue to bunch. For some price, even the seller with zero reservation value will no longer agree to bunch at the notch and hence no matches corresponding to higher $p(b, s)$ will bunch either. The bound in the Corollary depends on β . Interestingly, one can show that there is a uniform and finite bound for all β , so that transactions above some finite price are never induced to bunch, regardless of the value of β .

Theorem 3. For any $\beta > 0$, transactions at the notch satisfy $H \leq p(b, s) \leq \beta b^*(\beta) < x(0) < \infty$, where $f(b^*(\beta), 0; \beta) = 0$ and $x(\beta) \equiv \beta(b^*(\beta) - T)$ for any $\beta \in (0, 1)$, $x(0) = \lim_{\beta \rightarrow 0} x(\beta)$ and the value of $x(0)$ is the solution to $\ln(H) - \ln(x) - \frac{H-T}{x} + 1 = 0$.

Proof. In what follows we change variables as $x = \beta(b-T)$ (because it turns out that $\lim_{\beta \rightarrow 0} b^*(\beta) = \infty$). Note also that part (a) of the Lemma applied to $(b^*(\beta), 0)$ implies that $x(\beta) = \beta(b^*(\beta) - T) > H$ so that we don't need to consider $x \leq H$. Define the net benefit of locating at the notch for a given x as $g(x, \beta) \equiv f(\frac{x}{\beta} + T, 0; \beta)$ or more explicitly

$$\begin{aligned} g(x, \beta) &\equiv \beta \ln(H) + (1 - \beta) \ln\left(\frac{x - \beta(H - T)}{\beta}\right) - \ln\left(\frac{x}{\beta}\right) - \beta \ln(\beta) - (1 - \beta) \ln(1 - \beta) \\ &= \beta \ln(H) - \ln(x) + (1 - \beta) [\ln(x - \beta(H - T)) - \ln(1 - \beta)] \end{aligned}$$

We are interested in properties of $x(\beta)$ that solves $g(x(\beta), \beta) = 0$.

Denote by x^* the solution of $\ln(H) - \ln(x) - \frac{H-T}{x} + 1 = 0$. x^* is independent of β and finite. Note that $g(x, \beta)$ is continuous in x on $[H, x^*]$ and that $g(H, \beta) = (1 - \beta) \ln\left[\frac{(1-\beta)H + \beta T}{(1-\beta)H}\right] > 0$ and $g(x^*, \beta) = \beta \ln(H) - \beta \ln(x^*) + (1 - \beta) \ln\left[\frac{x^* - \beta(H - T)}{(1 - \beta)x^*}\right] < \beta \ln(H) - \beta \ln(x^*) + (1 - \beta) \left[\frac{x^* - \beta(H - T)}{(1 - \beta)x^*} - 1\right] = \beta [\ln(H) - \ln(x^*) - \frac{H-T}{x^*} + 1] = 0$ (using $\ln(x) < x - 1$ and the definition of x^*). Hence, for every β , $g(x, \beta) = 0$ has a solution on (H, x^*) and, in particular, $\lim_{\beta \rightarrow 0} x(\beta)$ has to be finite.

For all $\beta \in (0, 1)$, $x(\beta)$ solves $g(x(\beta), \beta) = 0$ so that $g_x x'(\beta) + g_\beta = 0$. Note that $g_x = \frac{1-\beta}{x-\beta(H-T)} - \frac{1}{x} = -\frac{\beta(x+H-T)}{x(x-\beta(H-T))}$. Clearly, for any $x > H - T$ we have $\lim_{\beta \rightarrow 0} g_x = 0$. Because $\lim_{\beta \rightarrow 0} x(\beta) > H - T$ and is finite, we also have $\lim_{\beta \rightarrow 0} g_x(x(\beta), \beta) = 0$ and $\lim_{\beta \rightarrow 0} |x'(\beta)| < \infty$. Consequently, $0 = \lim_{\beta \rightarrow 0} g_x x'(\beta) + g_\beta = \lim_{\beta \rightarrow 0} g_\beta(x(\beta), \beta) = \lim_{\beta \rightarrow 0} \left\{ \ln(H) - \ln(x(\beta) - \beta(H - T)) + \ln(1 - \beta) - \frac{(1-\beta)(H-T)}{x(\beta) - \beta(H-T)} + 1 \right\} = \ln(H) - \ln(x(0)) - \frac{H-T}{x(0)} + 1$ as in the statement of the proposition. \square

Example. When $H = 1,000,000$, $T = 10,000$, and $x(0) \approx \$1,144,717$, transactions that absent the tax would sell above this value will not bunch regardless of the value of β .

The gray dashed line on Figure 2.6 illustrates the bound, which corresponds to the price $x(0)$. As β changes, the slope of the corresponding line will change but it will always correspond to the price of $x(0)$. The sharp bound for a given β (the solid green line) always lies to the left of this uniform bound and converges to it as β tends to zero. While this bound is irrelevant given β , it is of natural interest when β is unknown.

2.9.2 Proportional tax

While considering a lump-sum tax simplifies the analysis, transaction taxes, including the mansion tax, are typically proportional. However, the results are only slightly affected when the tax is proportional. Intuitively, incentives for bunching at the notch are always determined by the level of the loss due to taxation (both due to the tax itself and any distortions it might cause), rather than the rate of the tax (this is standard intensive/extensive margin distinction). The proportional tax induces re-ranking, but retains qualitative features of the solution described above. In the presence of the proportional tax the Nash bargaining outcome is given by $p^s = q(b, s; t) \equiv \beta \frac{b}{1+t} + (1-\beta)s$, where $q(b, s, t)$ denotes the seller's given types and the marginal tax rate, and $p^b = (1+t)p^s = \beta b + (1-\beta)s(1+t)$ so that the overall surplus from the transaction is equal to

$$\beta \ln(\beta) + (1-\beta) \ln(1-\beta) + \ln(b - s(1+t)) - \beta \ln(1+t)$$

As in the case of the lump-sum tax, the price $q(b, s; t)$ is linear in types so that the locus of matches with constant price remains linear (although the slope is affected by the tax rate), as in Figure 2.6.

It is also straightforward to show an analogous result to Lemma 1 for the proportional case. Holding s constant, the net benefit to locate at the notch declines with b and becomes negative for sufficiently high b (because $(1-\beta) \ln(b - H) - \ln(b - s(1+t))$ declines in b). Thus, a schedule analogous to $b(s)$ also exists in the proportional case. Similarly to part (c) of the Lemma, the net surplus from switching to the notch declines in s holding $q(b, s; t)$ constant, so that the price corresponding to $b(0)$ has to constitute the upper bound of the region affected by the presence of the tax.

Finally, there is a straightforward relationship between the bounds corresponding to the lump-sum and proportional tax. To see it, note that for a given match (b, s) the value of locating at the notch is the same regardless of whether the tax is proportional or lump-

sum because that allocation does not involve any tax. Consider $s = 0$, and the value of b , $b^*(\beta; T)$, that as before represents a match that is indifferent between locating at the notch given the value of T . Simple inspection of the surplus for the lump-sum and proportional tax cases shows that the indifference will hold for the proportional tax as well (because the surplus will be the same as under the lump-sum tax) when the marginal tax rate is such that $\ln(b^*(\beta; T)) - \beta \ln(1 + t) = \ln(b^*(\beta; T) - T)$ so that $(1 + t)^\beta = \frac{b^*(\beta; T)}{b^*(\beta; T) - T}$. Thus, given β , the bounds for T map into the bounds for t by this relationship.

Theorem 3 describes a uniform bound for prices corresponding to transactions that might be affected by the lump-sum tax of T . Because bounds for proportional and lump-sum taxes are related for any β , that Theorem can be adapted to identify the corresponding bound in the proportional tax case.

Theorem 4. *Given marginal tax rate t , define x^* as the solution to $\ln(H) - \ln(x) - \frac{H - x \ln(1+t)}{x} + 1 = 0$. For any $\beta > 0$, transactions at the notch need to satisfy $H \leq q(b, s; t) \leq x^*/(1 + t) < \infty$, and x^* is the lowest such bound.*

Proof. To obtain the analogue of Theorem 3, recall that given T the Theorem established the existence of the upper bound of undistorted prices below which transactions (might) relocate to the notch. Denote by $x(T)$ the uniform bound for βb identified in Theorem 3 for a given value of T . For any β and tax rate t that satisfy $\ln(x(T)/\beta) - \beta \ln(1 + t) \geq \ln(x(T)/\beta - T)$, $x(T)$ would equal or exceed the undistorted price bound for transactions relocating to the notch under proportional tax. We will find the value of t for which $x(T)$ is the smallest such a bound for any positive β . Rewrite this inequality as $\ln(x(T)) - \beta \ln(1 + t) \geq \ln(x(T) - \beta T)$. Note that it holds with equality when $\beta = 0$. Taking derivatives of both sides with respect to β , we obtain $-\ln(1 + t)$ and $-\frac{T}{x(T) - \beta T}$ respectively. In order for the inequality to hold in the neighborhood of $\beta = 0$, we need to have $\ln(1 + t) \leq \frac{T}{x(T) - \beta T}$ for small β and when that's the case the inequality will hold for any value of β because the left-hand side is constant while the right-hand side is increasing in β . The bound will be tight when we do in fact have equality at $\beta = 0$ so that $\ln(1 + t) = \frac{T}{x(T)}$. Accordingly, when this relationship holds, substituting $x \ln(1 + t)$ for T in the equation defining $x(0)$ in Theorem 3 (which leads to the formula in the statement of the proposition) and solving for x will yield exactly the same solution $x^*(t)$. Finally, this procedure shows that when $s = 0$, only matches with $b \leq x^*(t)/\beta$ may bunch. Correspondingly, only transactions that satisfy $q(b, s; t) \leq q(x^*(t)/\beta, 0; t) = x^*(t)/(1 + t)$ might bunch. \square

Clearly, the value x^* that solves this formula is also the solution to the equation in Theorem 3 when $T = x^* \ln(1 + t)$, and the proof makes it clear that this is the right “conversion” between the proportional and lump-sum tax cases. The theorem provides a bound in terms of prices distorted by the tax $q(b, s, t)$. However, because $q(b, s, t)(1 + t) \geq q(b, s, 0)$ it also provides a (weaker) bound in terms of prices that are not distorted ($t = 0$): transactions that bunch need to satisfy $H \leq q(b, s; 0) \leq x^* < \infty$.

Example. For the New York and New Jersey mansion tax, $H = 1,000,000$, $t = 0.01$, $x^* \approx \$1,155,422$. Hence, regardless of the value of β , transactions that absent the tax would occur at prices above \$1,155,422 will never bunch, while those that would occur below might bunch. Transactions that do bunch, would otherwise sell (in the presence of the tax) at no more than $x^*/(1 + t) = \$1,143,982$. This is the same bound as the one corresponding to the lump-sum tax of $1,155,422 \cdot \ln(1 + .01) \equiv 11496.83$

2.9.3 Data Appendix

New York City Department of Finance Annualized Rolling Sales. The New York City Department of Finance (NYCDOF) Annualized Rolling Sales files contain details on real-property transactions for the five boroughs from 2003 to the present (we use the data through 2011). The data are realized by the NYCDOF on a quarterly basis and are derived from the universe of transfer-tax filings (which are mandatory for all residential and commercial sales). Geographic detail for each sale includes the street address (and zip code), the tax lot (borough-block-lot number), and the neighborhood (Chelsea, Tribeca, Upper West Side, etc.). The Rolling Sales files contain limited details about the properties themselves, including square footage, number of units (residential and commercial), tax class (residential, owned by utility co., or all other property), and building class category (a more detailed property code—for example, one-family homes, two-family homes, residential vacant land, walk-up condo, etc.). Transaction details in the data include the sale price and date. A sale price of \$0 indicates a transfer of ownership without cash consideration (ex. from parents to children).

New York City properties are subject to the mansion tax if they are single-, double-, or triple-family homes, or individual condo or co-op units. We define taxable sales as those transactions of a single residential unit (and no commercial units) with a building classification of “one family homes,” “two family homes,” “three family homes,” “tax class 1 condos,” “coops - walkup apartments,” “coops - elevator apartments,” “special condo billing lots/condo-rental,” “condos - walkup apartments,” “condos - elevator apartments,” “condos - 2–10 unit residential,” “condos - 2–10 unit with commercial use,” or “condo coops/condops.” We define co-ops as a building code of “coops - walkup apartments,” or “coops - elevator apartments.” We define a commercial sale to be a transaction with at least one commercial unit (and no residential units) or a tax class of 3 or 4.

New York State Office of Real Property Service SalesWeb. The New York State Office of Real Property Service (NYSORPS) publishes sales records for all real-property

transactions (excluding New York City) recorded between 2002–2006 and 2008–2010 available through the “SalesWeb” database. Since deeds are recorded after the sale, this data includes a small number of sales from 2007. The database is compiled by ORPS from filings of the State of New York Property Transfer Report (form RP-5217).

The NYS deeds records indicate several details about each transaction and property. Transaction-specific details include the sale price and date, the date the deed was recorded (and recording details such as book and page number), the buyer’s, seller’s, and attorney’s name and address (often missing), the number of parcels included in the transaction, and details about the relationship between the buyer and the seller (whether the sale is between relatives, whether the buyer is also a seller, whether one party is a business or the government, etc.). Of particular interest to us is whether the sale is defined by the state as arms-length. The data dictionary defines an arms-length sale as “a sale of real property in the open market, between an informed and willing buyer and seller where neither is under any compulsion to participate in the transaction, unaffected by any unusual conditions indicating a reasonable possibility that the full sales price is not equal to the fair market value of the property assuming fee ownership”, which excludes sales between current or former relatives, related companies or partners in business, sales where one of the buyers is also a seller, or sales with “other unusual factors affecting sale price.” Property details include the square footage, assessed value (for property-tax purposes), address (including street address, county, zip code, school district), and the property class (one-family home, condo, etc.). We consider as subject to the mansion tax all single-unit sales with property class equal to one-, two-, or three-family residence, residential condo, or a seasonal residence.

New Jersey Treasury SR1A File. We make use of sales records from the New Jersey Treasury’s SR1A file for 1996–2011, which contains records of all SR1A forms filed at the time of sale (the form is mandatory in the state for all residential sales). Each record includes the sale price and date the deed was drawn, buyer and seller name and address (often missing), deed recording details (date submitted, date recorded, document number),

and whether there are additional lots associated with the sale. Property details include land value, tax lot, square footage, and property class. We define taxable sales as those with a residential property class.

New York City County Register Deeds Records. These data are collected from the county registers for the five counties in New York City: Bronx County, Kings County, New York County, Queen’s County, and Richmond County. The records were collected by an anonymous private firm and made available to us by the Paul Milstein Center for Real Estate at the Columbia Graduate School of Business.

These data include additional detail as compared to the Rolling Sales files, although at the expense of precision. Prices in this data set are rounded to the nearest \$100, which leads to misallocation of sales to one side of a tax notch. Transaction details include the sale price and date, an indicator for whether the unit is newly constructed, the number of parcels being sold, whether the purchase was made in cash (i.e. whether a mortgage is associated with the sale), and indicators for private lenders and within-family sales. Property details are limited to address, zip code, and county.

Data Cleaning. We begin by dropping all transactions with a price below \$100 (1,658,639 in NY State, 954,241 in NJ, and 274,118). The bulk of these transactions have a zero price, representing transfers of property between parties not associated with a proper sale (e.g., a gift or inheritance). This restriction is relatively innocuous, as our analysis focuses on sales around each tax notch (although this choice does affect the descriptive statistics). More importantly, we attempt to identify and discard all duplicate records. In New York State, we identify duplicates as sales that occur within 90 days of one another at the same street number in the same grid number (a unique tax lot id). Of these 48,073 duplicates, we always keep the later sale (in case duplicates are representative of updates to the records). For New Jersey, since we do not observe tax lots, we identify all duplicate sales that occur at the same standardized address within 90 days of one another and drop all but the final duplicate (343,221). Finally, for NYC we identify duplicates as properties in the same borough at

the same standardized address that sell within 90 days of one another (20,420). While these duplicates represent a large number of sales, and there are several ways one could define duplicates, our estimates are insensitive to whether and how we clean duplicates (e.g., cleaning NY state based on address or NYC based on tax lot).

Real Estate Board of New York Listings Service. We have collected residential real-estate listings from the Real Estate Board of New York’s (REBNY) electronic listing service. REBNY is a trade association of about 300 realty firms operating in Manhattan and Brooklyn. REBNY accounts for about 50% of all residential real-estate listings in these boroughs. A condition of REBNY membership is that realtors are required to post all listings and updates to the listing service within 24 hours.

Using the REBNY listing service, we have collected all “closed” (i.e. sold) or “permanently off market” residential listings posted between 2003 (when the electronic listings are first available) and 2010. REBNY listings include the typical details available on a real-estate listing: asking price, address, date on the market and a description of the property. Additionally, we observe all updates to each listing (and the dates of each update), which lets us see how asking prices evolve and determine the length of time a property is on the market. Finally, we observe the final outcome of the listing: whether the property is sold or taken off the market.

We create several variables for each REBNY listing. We define the initial asking price as the first posted price on the listing, and the final asking price as the last posted price while the listing is “active.” We identify the length of time that a listing spends on the market as the number of days between the initial posting and the date that the listing is updated as “in contract.” We define the discount between two prices as the percent drop in price— $\frac{p_0 - p_1}{p_0}$, where p_0 and p_1 are prices and p_0 is posted before p_1 .

One caveat to the REBNY listings is that the price is often not updated at the time of sale. To overcome this, we match REBNY listings to the NYCDOF data by address and date. Of the 48,220 closed REBNY listings for Manhattan, we achieve a match rate of 92%.

Non-matches fall in a number of categories. Sales in some condop buildings are missing from the DOF data due to a clerical error at the NYC DOF. Some transactions contain only street address or a non-standard way of specifying the apartment number (in particular, commercial units and unusual properties such as storage units fall in this category). Occasionally, the same building may have two different street addresses and a unit may be listed differently in the two databases. At the same time, of the 23,655 Manhattan listings that are not reported as closed in the REBNY listings database, we find 7,425 corresponding sales in the NYCDOF data. We treat such matches as an indication that the property was sold without the REBNY realtor (either sold by the owner or using another realtor).

2.9.4 Robustness of incidence estimates

Table 2.8 demonstrates that our estimates are quite robust to variety of estimation approaches. Incidence estimates are very consistent, and gap estimates vary somewhat but remain positive and large in most specification checks that we consider. Intuitively, there are good reasons for why results may vary as one adjusts the order of polynomials and the omitted region. Both incidence and gap estimated using cross-sectional data (the only exception to it our estimates for NJ that rely on pre/post comparison) involve prediction out of sample (into the omitted region). As the size of the omitted region increases, one has to predict far out of sample so that the “forecast” error is bound to increase. Furthermore, very flexible polynomials that can fit data in sample well are not restricted in their behavior in the omitted region and in some cases may generate non-monotonicity or explosive behavior within the omitted region — overfitting is not the right approach for predicting out of sample. On the other hand, the omitted region that is too small generates bias in the estimates of the counterfactual. Nevertheless, our results are robust to reasonable modifications of our baseline specification as discussed below.

While our preferred specification uses a third-order polynomial, our incidence estimates are not too sensitive to this choice. The second through fifth rows of Table 2.8 present

estimates that we obtain using different orders—the results are similar, although inspection of the fit of the data suggests that very low-order polynomials cannot capture properly the shape of the distribution, while very high-order polynomials (not reported) introduce very unrealistic behavior in the omitted region. As the result, there is a bit of sensitivity to the order of polynomials in the gap estimates, which are positive and significant for all specifications up to the fifth order polynomial, but shrink somewhat for higher orders.

The results are only somewhat sensitive to selecting a narrow omitted region. The estimates in the sixth through eighth rows of Table 2.8 illustrate that a smaller omitted region leads to smaller incidence estimate (\$3000 to \$5000 less than the baseline). We do not estimate \hat{Z} for the narrowest specifications since it does not make sense to restrict the gap to be so small, especially given the visual evidence of the width of the gap. Relatedly, the estimate of \hat{Z} using the omitted regions through \$1.1M are smaller than the baseline — understandable, given the visual evidence in Figures 2.1, 2.2 and 2.3 indicating that the gap extends further than that (consistently with the theoretical argument as well), and the fact that the counterfactual is bound to be biased downward when part of the “true” gap is relied on in estimation.

On the other hand, our results are robust to extending the omitted region beyond the baseline, as is seen in Rows 10 through 12. The estimates of \hat{Z} are consistently large, positive, and significant, although less precise as we use less and less data (and need to predict the counterfactual over a larger range). Reassuringly, the incidence estimates change little as we vary the upper bound of the omitted region. Similarly, none of the estimates are too sensitive to expanding the omitted region below the threshold. The results in rows twelve through fourteen of Table 2.8 show that both the incidence and gap estimates grow as the bunching region is expanded below the threshold, however differences in estimates are economically small and not statistically distinguishable from the baseline. Naturally, the standard errors also grow as the omitted region is expanded.

We also estimate our counterfactuals for bunching and gap separately using only data

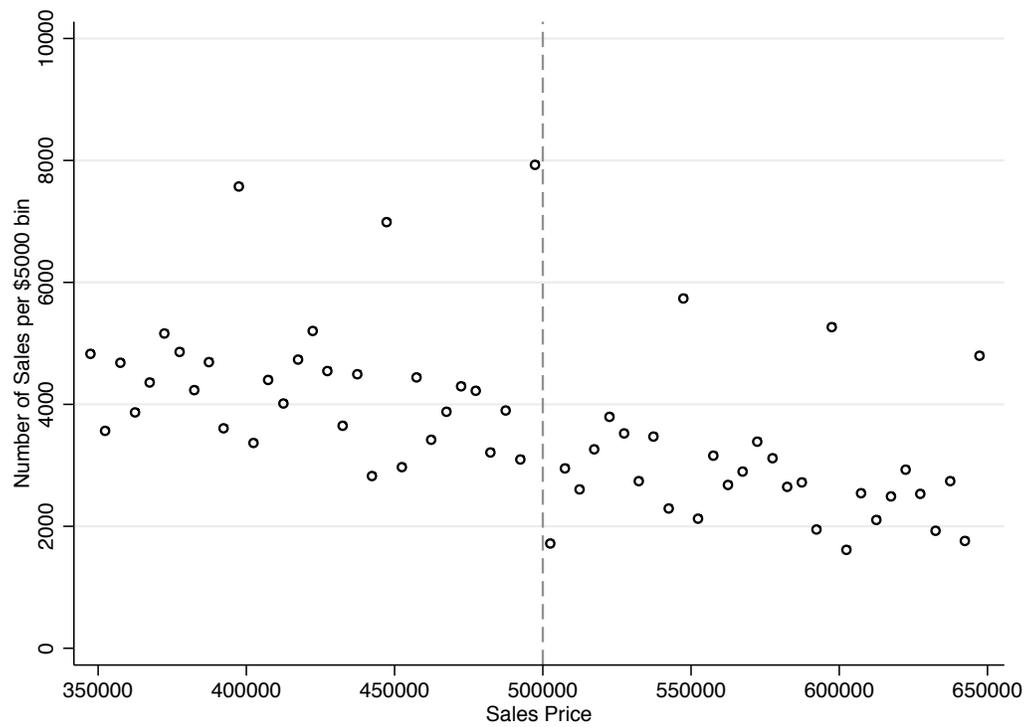
below and above the omitted region (respectively). We present in Row 16 our estimate using a 3rd order polynomial and data below the omitted region for the bunching/incidence counterfactual and a 1st order polynomial using data above the omitted region for the missing mass, and a 2nd order above the omitted region in Row 17. Again, incidence and \hat{Z} are comparable with our baseline. Furthermore, bootstrapped standard errors increase significantly as the order of polynomial increases, underscoring our earlier point that allowing for overfitting by estimating high order polynomials that are then used to project into the omitted region is a questionable approach. This observation (and visual inspection of the fit) justified our choice of the baseline specification that relies on the 3rd order polynomial and only a level shift at the threshold.

Our baseline estimate is also not sensitive to allowing for a discontinuity at the threshold. The baseline specification relies on the data both below and above the omitted region. Since the latter is distorted by the tax we rudimentarily control for it by allowing for a level shift in the distribution. The estimate in row 18 of Table 2.8 demonstrates that incidence and gap increase slightly when we do not allow for this discontinuity.

For completeness, we also estimate analogous specification by OLS—this is the standard approach in the recent public finance work on notches and kinks—but we note that any of these methods involves specifying the parametric density function and the maximum likelihood estimation is a natural choice that guarantees that the estimates satisfy the law of probability rather than the hard-to-interpret mean zero residual restriction. Additionally, by requiring the data to be binned, OLS will throw out information. We report the OLS results obtained by binning into \$5000 and \$10,000 bins in rows 17 and 18 and conclude that they are quantitatively similar to the baseline.

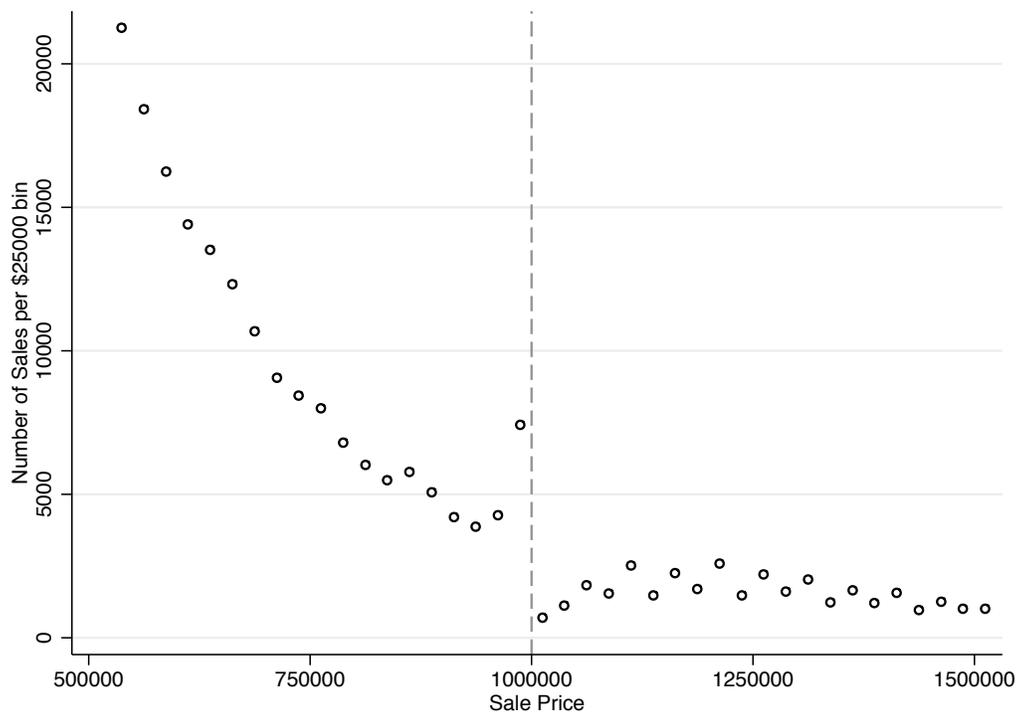
The placebo estimates in Table 2.9 show that our estimates for NYC are not spurious. Using the same procedure, we estimate the incidence and gap for all commercial sales (which are not subject to the mansion tax) and for residential sales at other multiples of \$100,000. In all cases, we find small negative incidence estimates and relatively small gap estimates.

Figure 2.11: Distribution of Sales in New York City around the \$500,000 RPTT tax notch



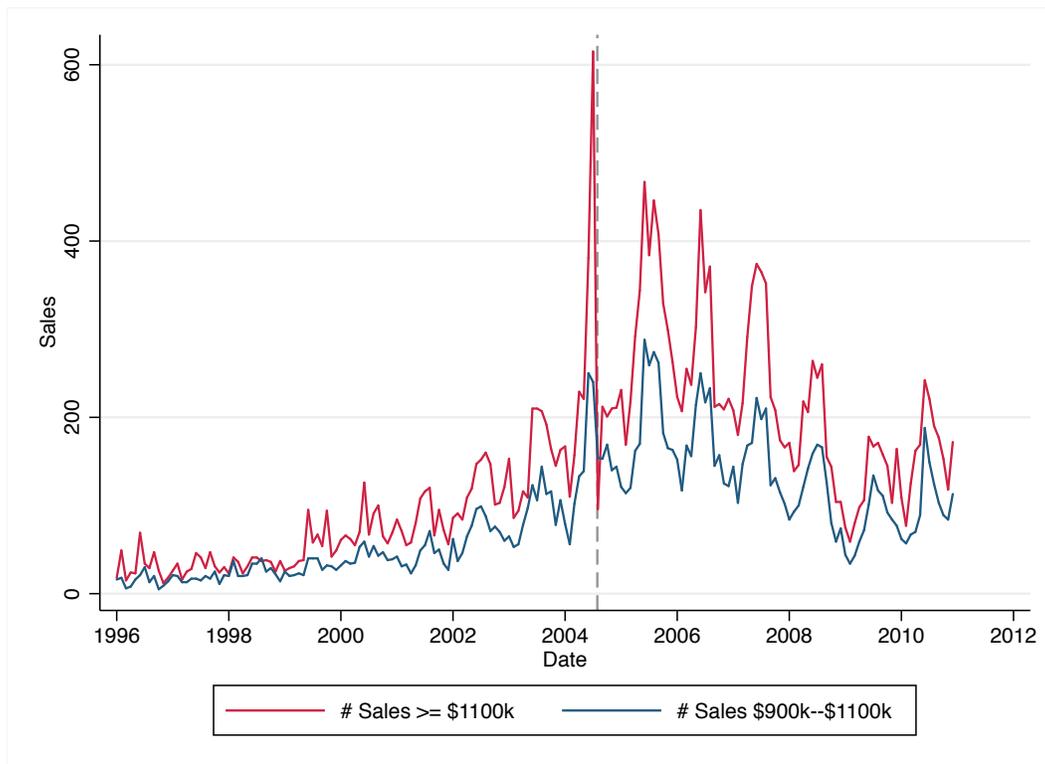
Notes: Plot of the number of sales in each \$5,000 price bin between \$350,000 and \$650,000. Data from the NYC Rolling Sales file for 2003–2011. Both commercial and non-commercial sales are subject to the NYC RPTT.

Figure 2.12: Distribution of Taxable Sales in New York State



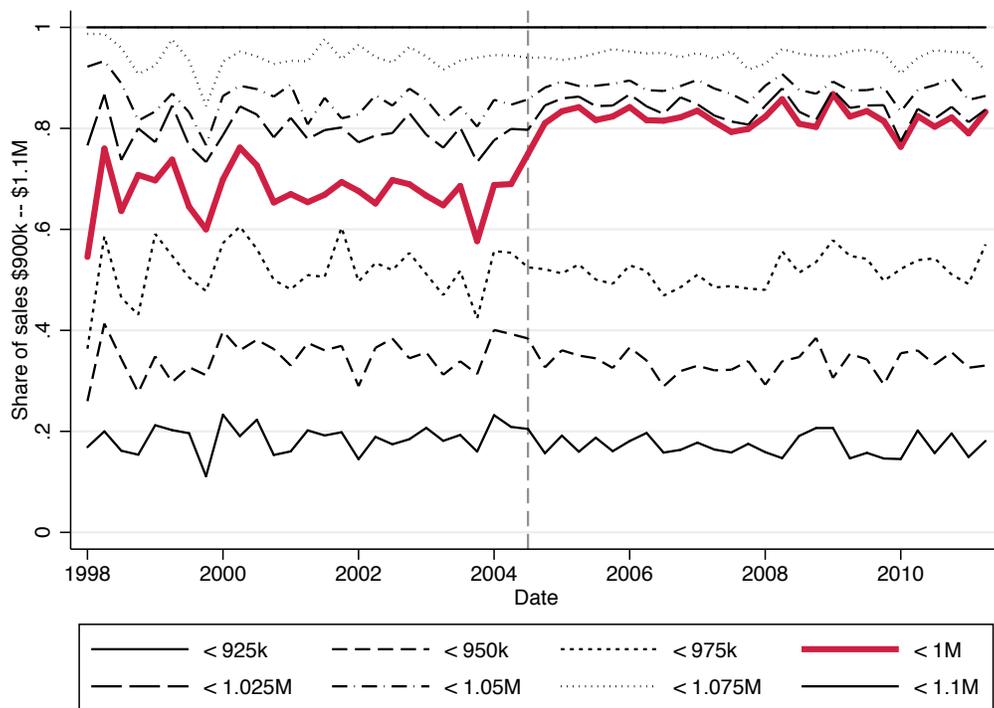
Notes: Plot of the number of mansion-tax eligible sales in each \$25,000 price bin between \$510,000 and \$1,500,000. Data from the NYC Rolling Sales file for 2003–2011 (taxable sales defined as single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) and from N.Y. State Office of Real Property Service deeds records for 2002–2006 and 2008–2010 (taxable defined as all single-parcel residential sales of one-, two-, or three-family homes).

Figure 2.13: New Jersey Monthly Sales Above \$990,000



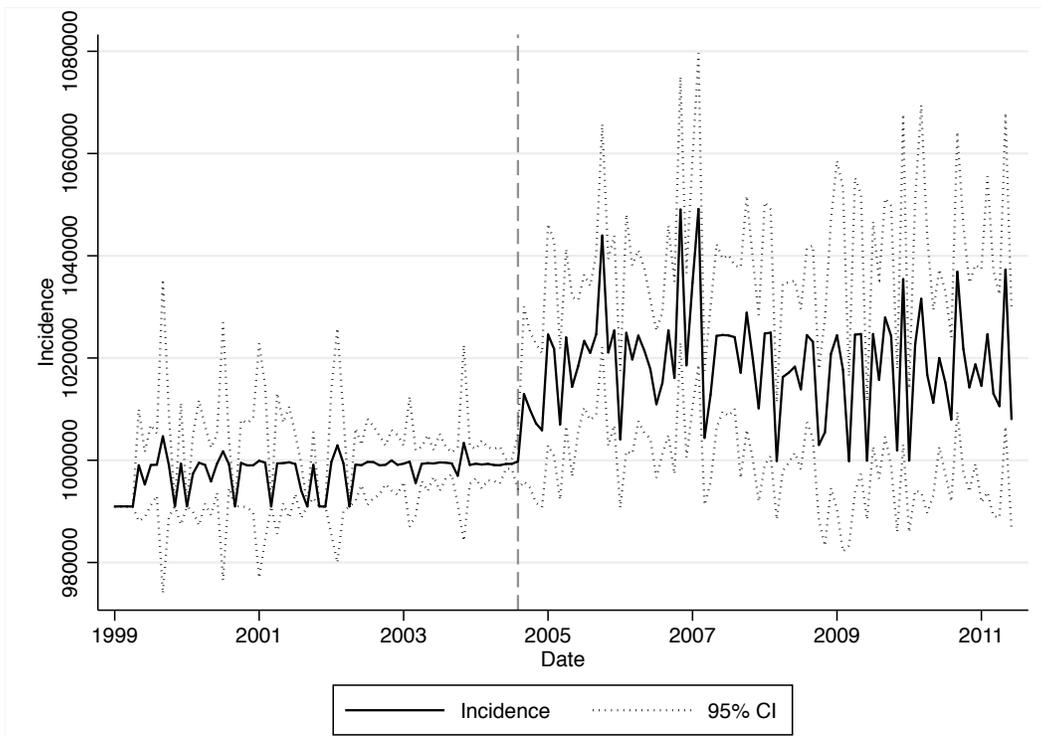
Notes: Total taxable NJ sales in given price range by month. Data from NJ Treasury SR1A file for 1996–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (indicated by dashed gray line).

Figure 2.14: Distribution of Monthly Sales in New Jersey (\$900k – \$1M)



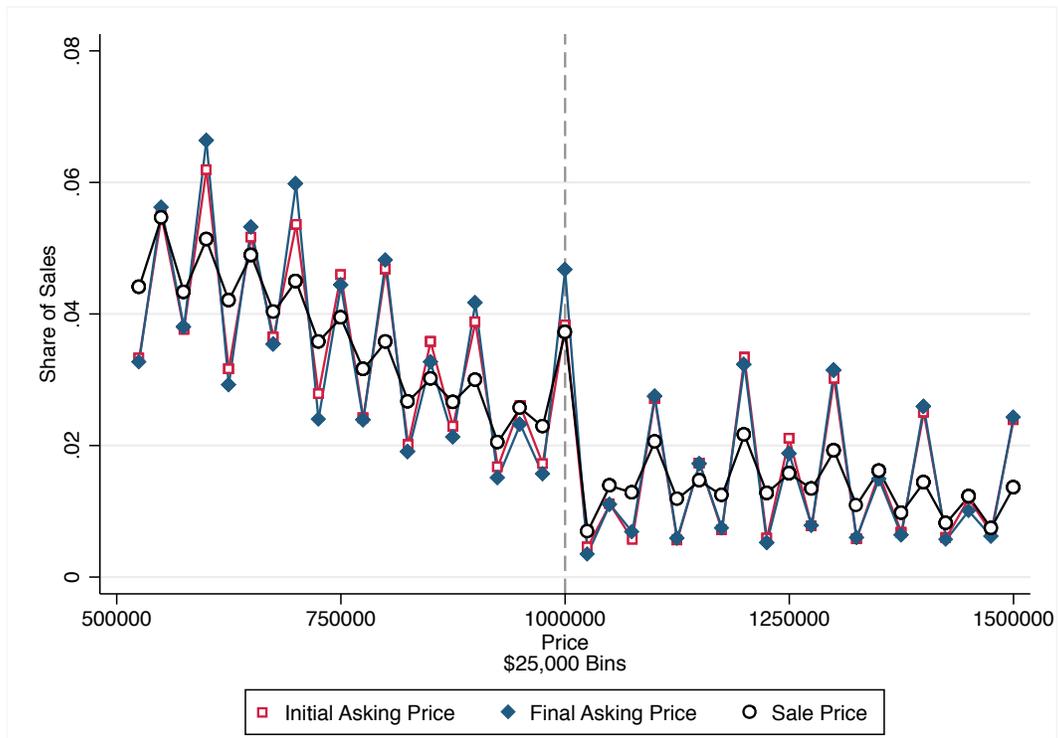
Notes: Number of taxable sales in given range as a share of total sales between \$900,000 and \$1,100,000 by month. Data from NJ Treasury SR1A file for 1998–2011 (taxable defined as any residential sale). Mansion tax introduced in August, 2004 (denoted by gray dashed line).

Figure 2.15: NJ Local Incidence Over Time



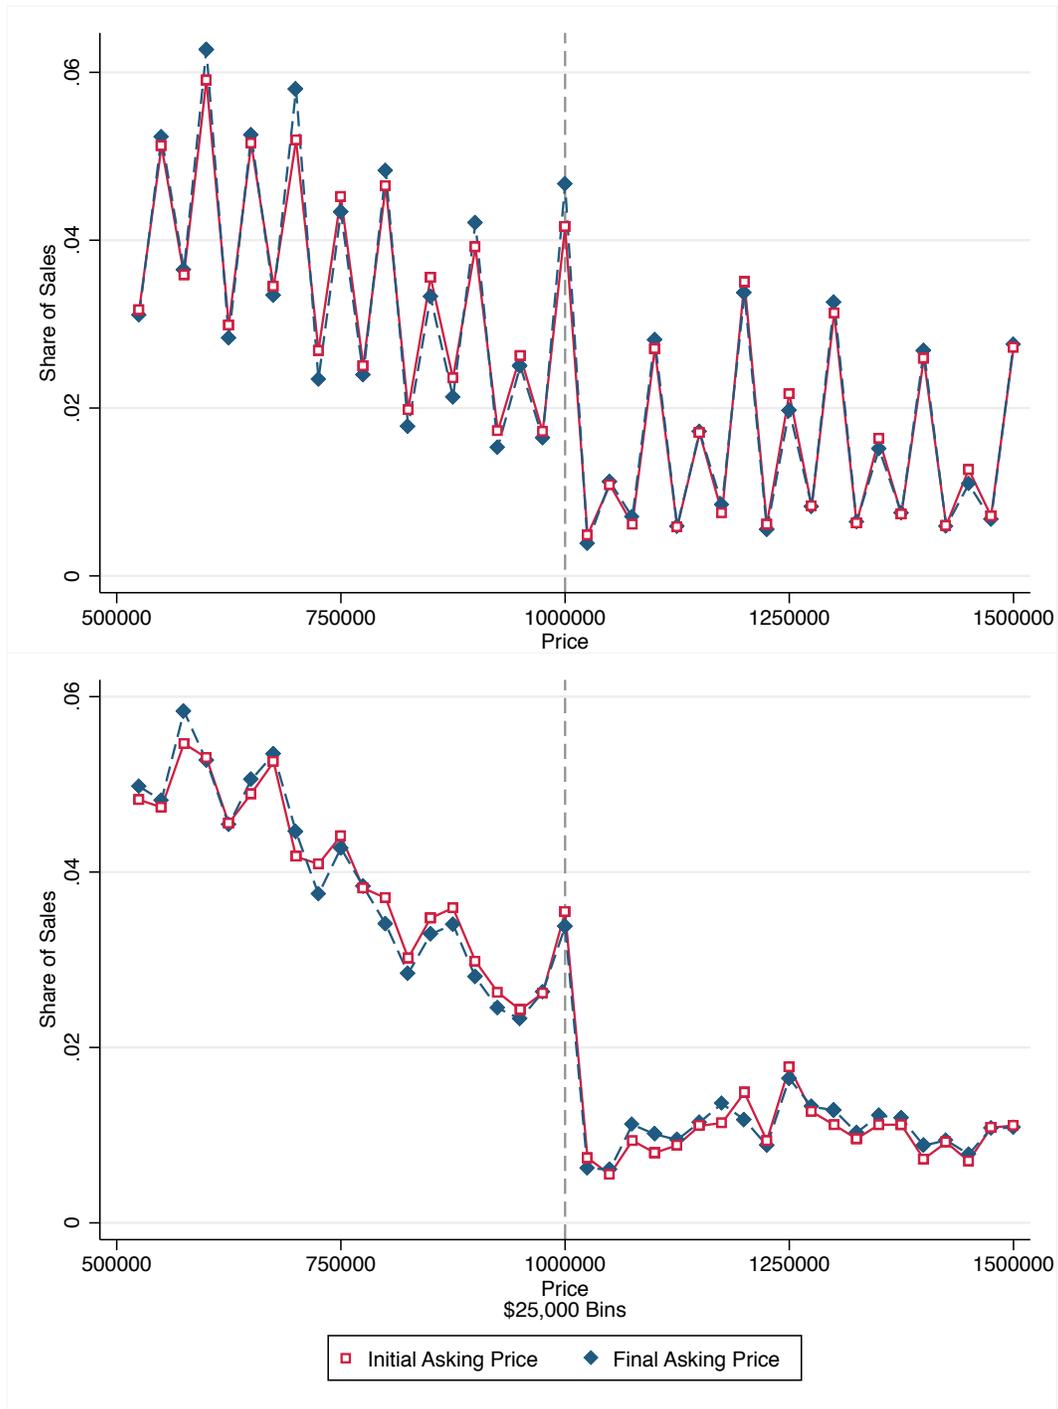
Notes: Monthly baseline local incidence estimates and 95% confidence intervals for NJ. Data from NJ Treasury SR1A file.

Figure 2.16: Distribution of Real-Estate Listing Prices in NYC (Sold Properties Only)



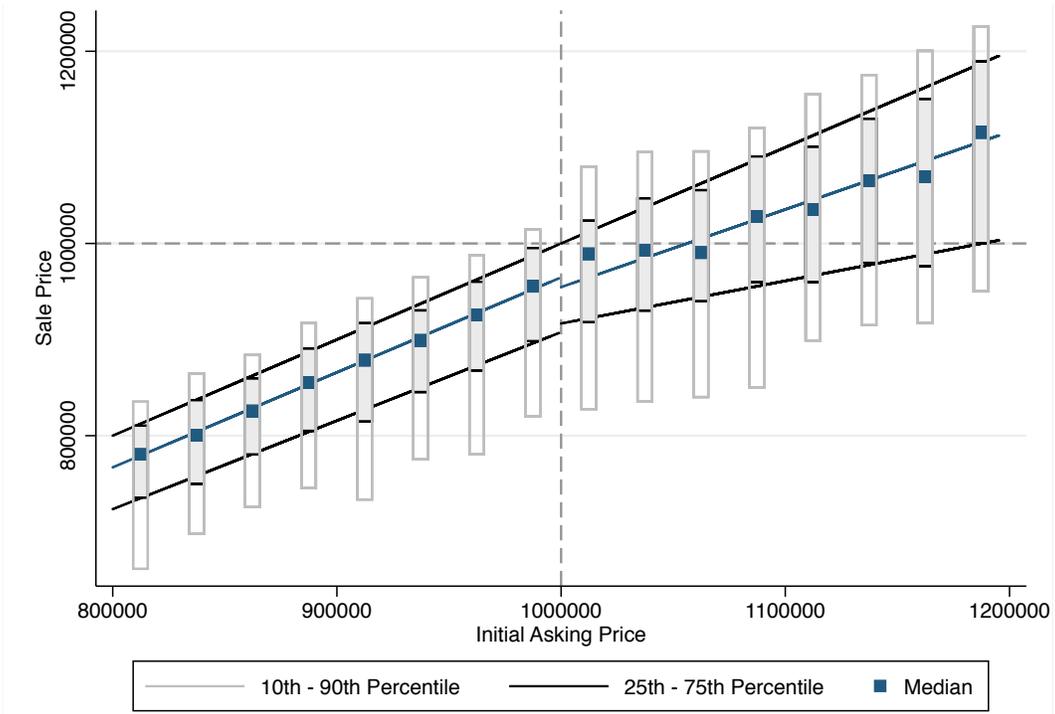
Notes: Data from REBNY listings matched to NYC Department of Finance sales records. Sample restricted to “sold” listings: last listing status is “closed” and property can be matched to NYC sales data. Plot of the number of listings per \$25,000 bin as a share of all sales between \$500,000 and \$1,500,000 (bins centered so that the threshold bin spans \$975,001–\$1,000,000).

Figure 2.17: Distribution of Real-Estate Listing Prices in NYC (All REBNY-Listed Properties)



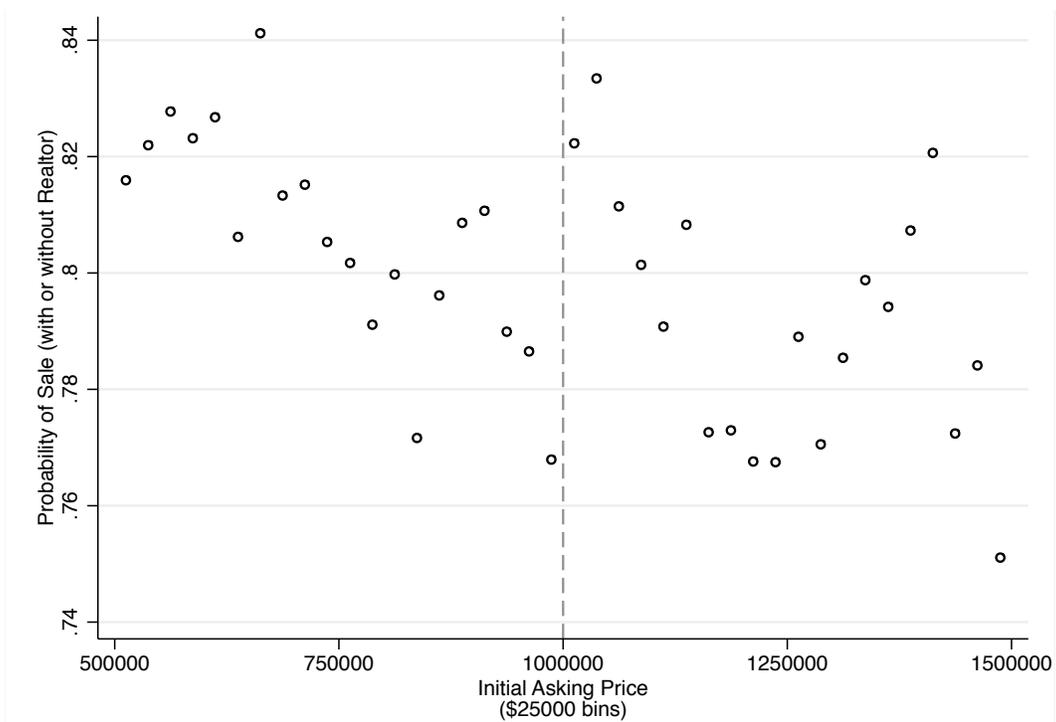
Notes: Data from REBNY listings. Sample includes all REBNY-listed sales in the given range. Panel (a) presents a plot of the number of listings per \$25,000 bin as a share of all listings between \$500,000 and \$1,500,000 (bins centered so that the threshold bin spans \$975,001–\$1,000,000). Panel (b) presents a smoothed plot of the distribution that accounts for round-number bunching: the log of the per-bin counts from panel (a) are regressed on a cubic in price and dummy variables for multiples of \$50,000 and \$100,000 interacted with the price. Predicted bunching for round-number bins are then subtracted from the corresponding counts.

Figure 2.18: Distribution of Sale Price by Initial Asking Price (with Quantile Regression)



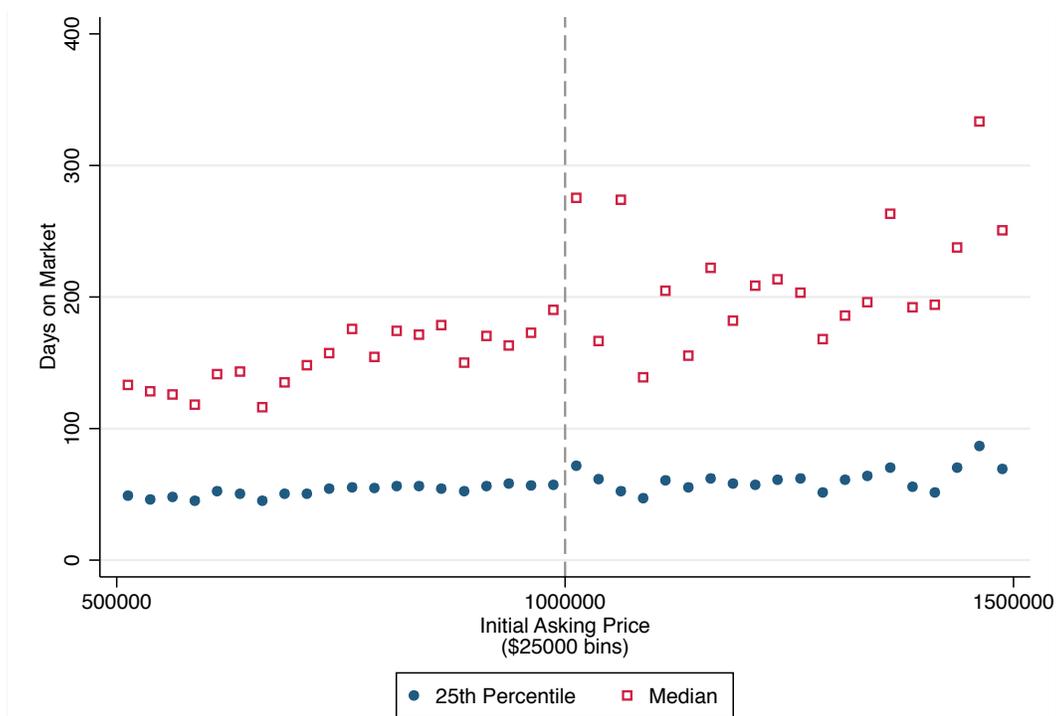
Notes: Plot of the median, 10th, 25th, 75th, and 90th percentiles of sale price per \$25,000 initial-asking-price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range \$800,000–1,200,000. Lines represent quantile regressions for the given range (\$800k–\$990k and \$1M – \$1.2M).

Figure 2.19: Probability that Listed Property Sells by Initial Asking Price



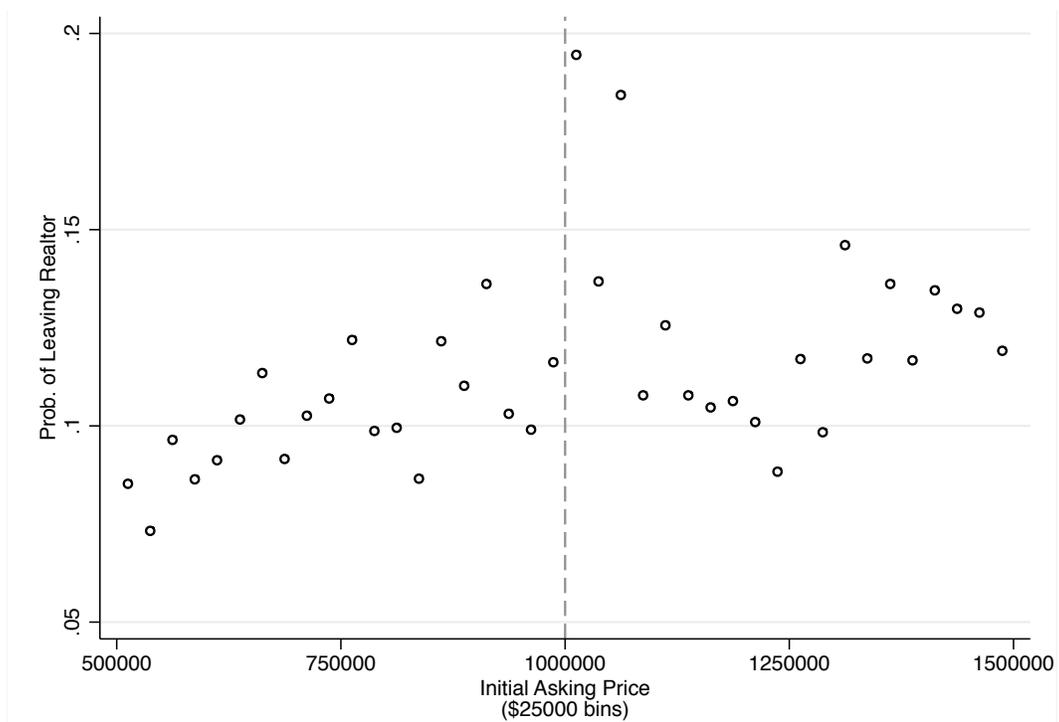
Notes: Plot of the share of REBNY-listed properties that close or are matched to a NYC DOF sale per \$25,000 bin. Data from REBNY listings—sample includes all listed properties in the range \$500,000–1,500,000. “Sold” defined as any property with a final listing status of “closed” or any listing that matches to NYC DOF sales.

Figure 2.20: Median Days to Sale by Initial Asking Price



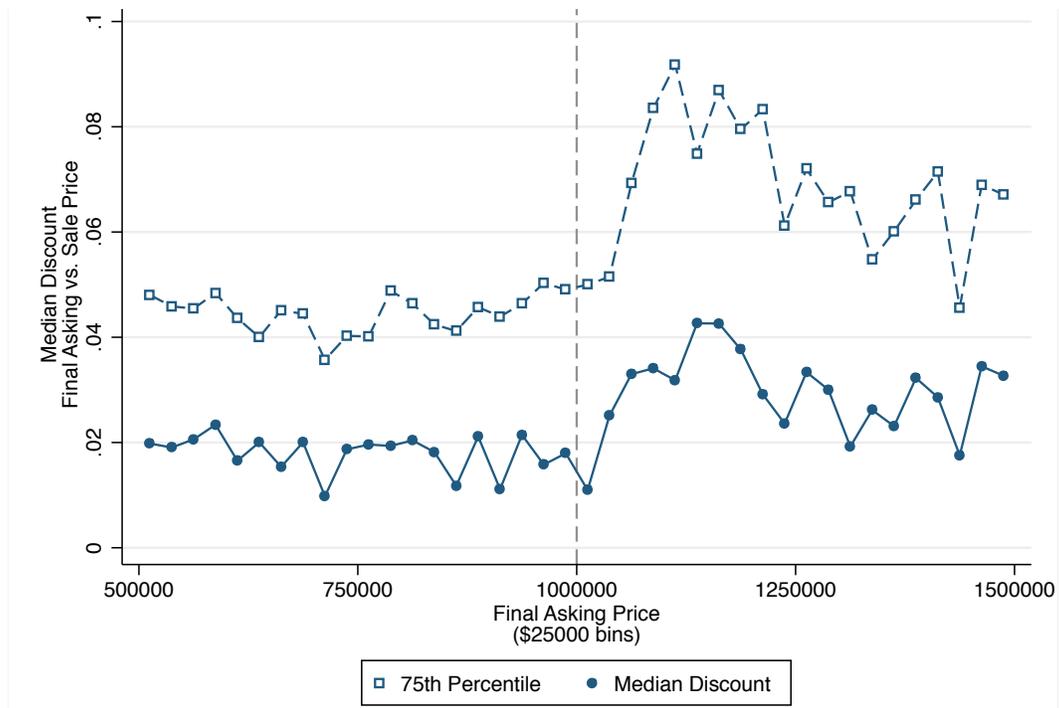
Notes: Plot of the median and 25th percentile of days to sale per \$25,000 initial-asking-price bin. Data from REBNY listings—sample includes all REBNY-listed properties in the range \$500,000–1,500,000. Days to sale defined as the number of days between initial listing of the property and buyer and seller entering into contract (defined as final status = “in contract”). Unsold properties are assigned a value of 999 days.

Figure 2.21: Probability of Selling Without REBNY by Initial Asking Price



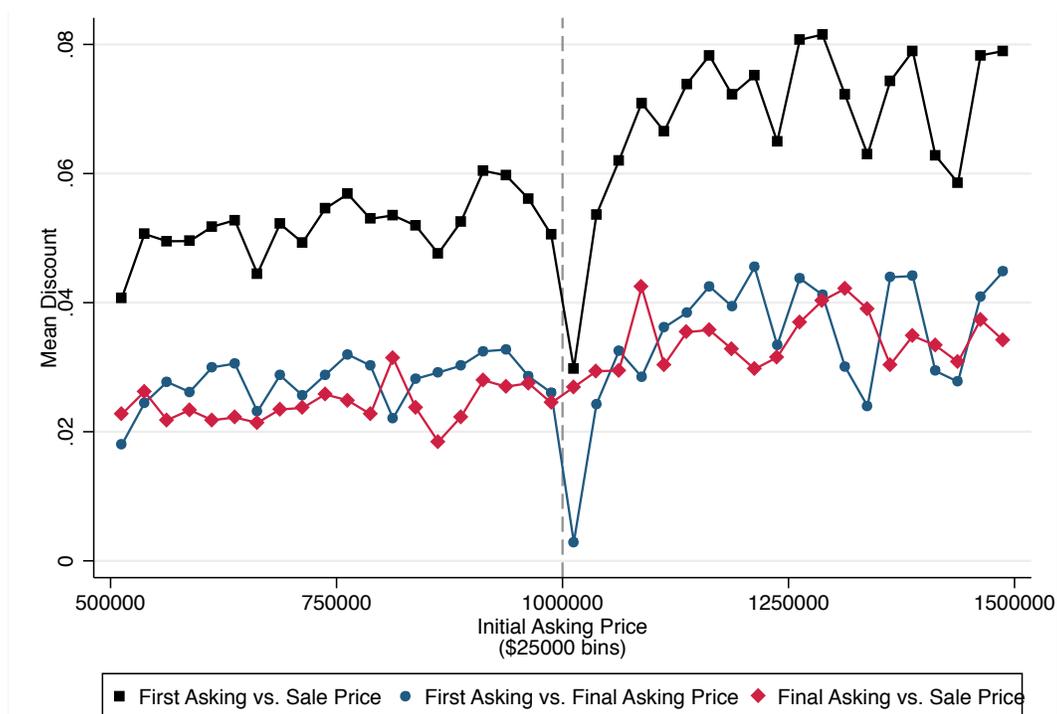
Notes: Plot of the share of REBNY listed properties that are sold in NYC DOF data, but are not listed as closed in the REBNY listing per \$25,000 initial-asking-price bin. Data from REBNY listings—sample includes all REBNY-listed properties in the range \$500,000–1,500,000.

Figure 2.22: Median Price Discount by Final Asking Price



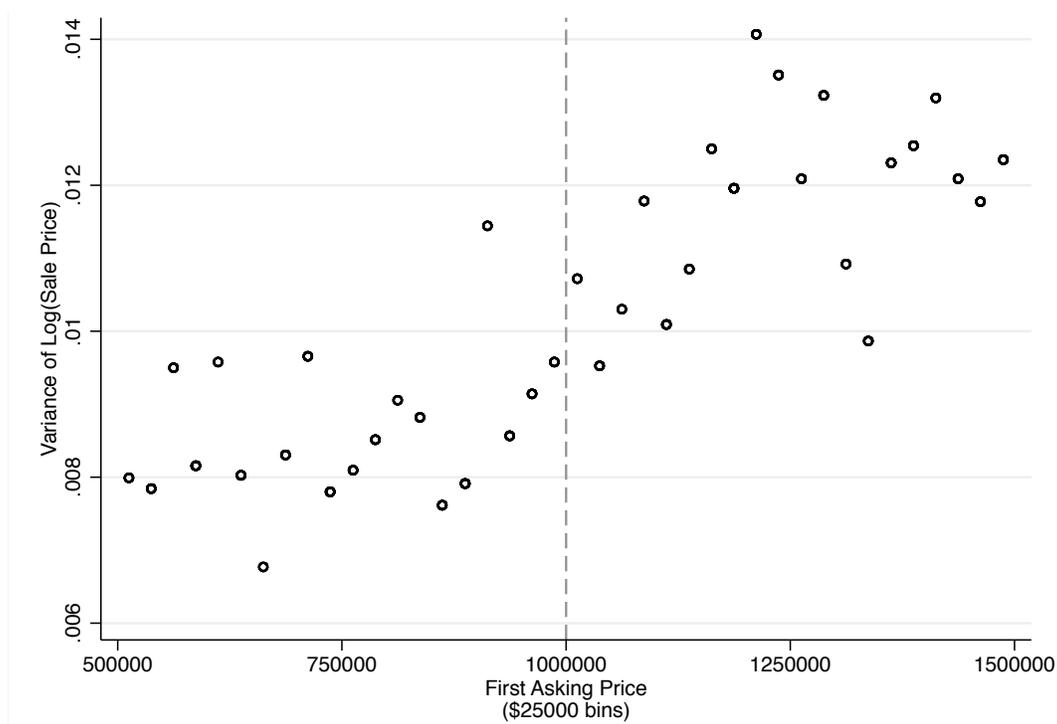
Notes: Plot of the median and 25th percentile discount from final asking price to sale price ($= 1 - \text{sale}/\text{final}$) per \$25,000 final-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range \$500,000–1,500,000 that match to NYC DOF data.

Figure 2.23: Decomposition of Mean Price Discounts by Initial Asking Price



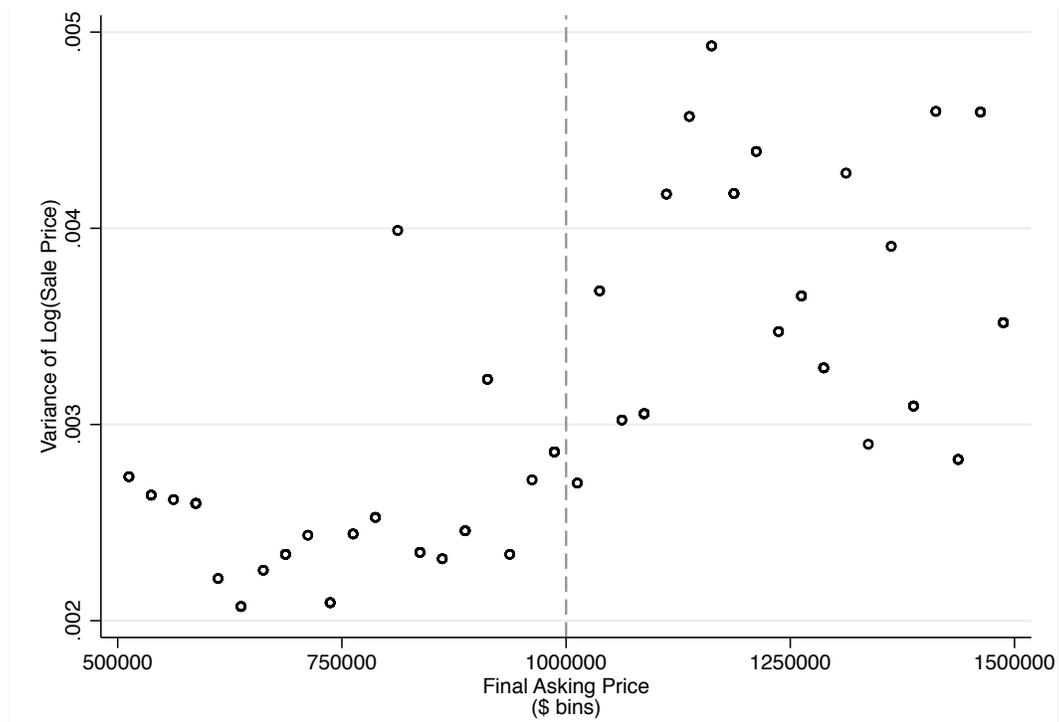
Notes: Plot of the average discount from initial asking to sale price ($= 1 - \text{sale}/\text{initial}$), initial asking to final asking price ($= 1 - \text{final}/\text{initial}$), and final asking price to sale price relative to initial asking price ($= (\text{final} - \text{sale})/\text{initial}$) per \$25,000 initial-asking-price bin. Data from REBNY listings—sample includes all closed REBNY-listed properties in the range \$500,000–1,500,000 that match to NYC DOF data.

Figure 2.24: Dispersion of Sale Price, Conditional on First Asking Price



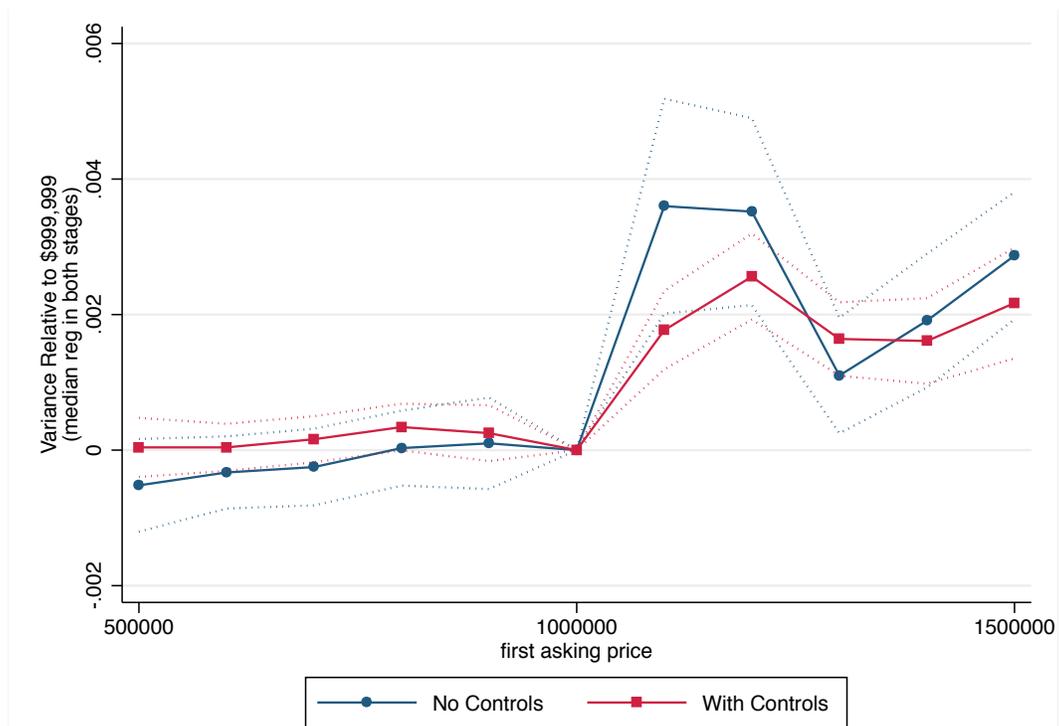
Notes: Plot of the variance of the log of sale price by \$25,000 initial asking price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range \$500,000–1,500,000. Observations with price discounts (first asking price to sale price) in the 1st or 99th percentile are omitted.

Figure 2.25: Dispersion of Sale Price, Conditional on Last Asking Price



Notes: Plot of the variance of the log of sale price by \$25,000 final price bin. Data from REBNY listings—sample includes all sold REBNY-listed properties (matched to NYC DOF) in the range \$500,000–1,500,000. Observations with price discounts (first asking price to sale price) in the 1st or 99th percentile are omitted.

Figure 2.26: Predicted Dispersion of Log of Sale Price (Median Regression in First Stage)



Notes: Plots of the difference between predicted values at given initial asking price and predicted value at \$1,000,000 from the following procedure. The log of sale price is regressed (median regression) on a linear spline in the log of initial asking price with \$100,000 knots between \$500,000 and \$1,500,000. Squared residuals from this first stage are then regressed on a linear spline in log of initial asking price (using median regression; results are sensitive to outliers). Controls in the indicated results include year of sale, zip code, building type, whether the sale is of a new unit, and the log of years since construction. Dashed lines represent 95% confidence intervals from 999 wild bootstrap replications of the two-stage procedure, resampling residuals in the first stage by asking-price clusters.

Table 2.8: Mansion Tax: NYC

Specification	Incidence	Std. Error	\hat{Z}	Std. Error	n
1. Baseline: 3rd Order, Omit \$990k – \$1.155M	21542.098	1124.005	43861.766	3990.977	102493
2. 1st Order	24095.152	475.318	106372.697	4807.543	102493
3. 2nd Order	22910.857	965.353	46988.518	4118.622	102493
4. 4th Order	21115.344	1160.988	37069.861	5734.794	102493
5. 5th Order	18444.127	1490.541	33652.867	5868.148	102493
6. Omit \$990k – \$1.01M	16308.129	974.380	.	.	108766
7. Omit \$990k – \$1.025M	16676.799	967.177	.	.	108378
8. Omit \$990k – \$1.050M	18168.244	1062.597	2311.782	1280.552	107635
9. Omit \$990k – \$1.1M	20163.154	1134.501	15145.428	2117.299	105549
10. Omit \$990k – \$1.2M	21978.363	1079.676	51725.536	5165.912	100846
11. Omit \$990k – \$1.255M	22122.063	1127.693	53596.361	8617.068	97673
12. Omit \$990k – \$1.3M	21917.696	1119.279	25178.003	10285.031	96094
13. Omit \$980k – \$1.155M	24980.343	957.747	43954.765	4051.863	100706
14. Omit \$970k – \$1.155M	26742.767	1634.857	44726.198	3810.847	99741
15. Omit \$960k – \$1.155M	37687.334	2374.628	49452.601	4938.454	98915
16. 3rd Order for LHS, 1st Order for RHS	20034.887	1372.811	32006.872	3891.654	102493
17. 3rd Order for LHS, 2nd Order for RHS	20034.887	1372.811	37243.921	9670.122	102493
18. No Discontinuity at Threshold	21601.728	1048.292	44349.086	2316.284	102493
19. OLS, \$5000 Bins	19539.143	2093.029	.	.	102493
20. OLS, \$10000 Bins	20956.556	1539.757	.	.	102493

Notes: Estimates by MLE as described in text (or OLS with binned data where indicated) using data from NYC Department of Finance Rolling Sales file for 2003–2011. Sample is restricted to all taxable sales (single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) with prices between \$510,000 and \$1,500,000. We do not estimate \hat{Z} for specifications 6 and 7 since our method mechanically restricts the missing mass/gap to the range between \$1M and the top of the omitted region. Since the omitted regions above the notch for these two specifications are small, the estimates of the missing mass in the gap are artificially small. Specification 15 estimates the counterfactual separately using data below the omitted region of \$990k–\$1,155,422 to estimate bunching and incidence, and data above the omitted region to estimate the gap and \hat{Z} .

Table 2.9: NYC Mansion Tax: Placebos

Cutoff	Incidence	Std. Error	\hat{Z}	Std. Error	n
Commercial	-634.223	69.3692	203.190	512.087	5616
600,000	-689.324	32.9582	430.942	386.221	74477
700,000	-708.755	30.5481	257.927	258.632	86026
800,000	-787.246	29.686	360.535	280.648	93148
900,000	-732.660	32.098	-985.343	197.414	98003
1,100,000	-937.686	141.5243	457.587	559.861	106344
1,200,000	-669.331	42.6621	346.163	591.304	103660

Notes: Data from NYC Department of Finance Rolling Sales file for 2003–2011. Commercial sales are defined as any transaction of at least one commercial unit and no residential units or a NYC tax class of 3 (utility properties) or 4 (commercial or industrial properties) and are not subject to the mansion tax. Placebo estimates are found using the baseline MLE procedure around the given cutoff.

Table 2.10: Local Incidence Over Time

Year	NYC				NYS				NJ						
	Incidence	Std. Error	\hat{Z}	n	Incidence	Std. Error	\hat{Z}	n	Incidence	Std. Error	\hat{Z}	Std. Error	n		
1996	-832.90	2010.30	591.59	4577.63	1929		
1997	-846.88	892.57	4672.45	3965.02	2308		
1998	-516.75	325.59	-7.22	2278.70	3176		
1999	-1651.36	1641.19	5799.47	5644.42	4193		
2000	-855.52	219.82	1610.77	751.77	5402		
2001	-890.48	362.37	4516.89	1772.83	6461		
2002	22203.32	3459.24	54264.66	14737.66	9596	739.28	105.02	2921.91	818.47	9606	
2003	14624.04	4208.77	40388.15	15073.28	7310	19003.62	3607.36	63687.59	12761.62	12462	-636.85	95.72	2995.21	668.84	12953
2004	23429.53	2510.09	19803.96	10841.69	10451	22087.23	2651.17	63869.40	11328.73	17389	-785.30	94.24	3425.06	736.19	10800
2004 (post)	5996.99	4145.92	38365.00	16506.71	8442		
2005	32614.47	4526.29	64561.58	12876.31	13554	24014.88	1833.18	47688.16	10156.43	21765	24229.46	1233.46	55087.72	10767.80	24921
2006	21603.44	2520.20	47836.37	11784.42	14728	24217.07	1562.11	28400.20	10617.07	17958	22080.59	2567.27	34109.31	9633.65	21963
2007	17487.07	2388.52	42113.91	8840.71	16477	22673.58	10770.93	16409.17	38609.60	948	24119.93	2025.23	27611.91	10546.27	18921
2008	17685.93	2587.64	34771.01	9461.67	12911	23867.94	2571.47	14847.10	11051.02	11143	19047.95	3479.87	31128.82	10877.80	12913
2009	20449.58	3434.69	1338.87	10748.73	8682	23930.51	3294.99	24523.63	14400.24	8370	20673.62	3852.98	24910.20	13266.07	10039
2010	17125.48	3781.44	74753.21	15731.87	8534	22620.92	3170.81	21216.03	12132.49	8831	19520.01	3566.17	37097.43	13552.64	10948
2011	21739.30	2853.97	70566.20	15725.26	9846	.	.	.	22460.74	5574.45	27266.46	21619.09	3789		

Notes: Baseline local incidence estimates by year for given geography. Data for NYC from Department of Finance Rolling Sales files, restricted to taxable sales (single-unit non-commercial sales of one-, two-, or three-family homes, coops, and condos) in the given year. Data for NYS from the Office of Real Property Services deeds records, restricted to taxable sales (all single-parcel residential sales of one-, two-, or three-family homes) in the given year. NYS ORPS data covers 2002-2007 and 2009-2010; observations in 2007 are from sales made in 2007, but recorded in 2008-2010. Data for NJ from NJ Treasury SR1A file.

Table 2.11: Predicted Price Discounts

	Discount: Initial Asking to Sale Price		Discount: Initial Asking to Final Asking	
	Slope	Value at Lower Knot relative to \$1,000,000	Slope	Value at Lower Knot relative to \$1,000,000
<i>Prediction at \$1,000,000</i>				
		<i>0.05117</i>	<i>0.02464</i>	
500000 to 600000	0.12785 (0.6455)	-0.00279 (0.00579)	0.71910 (0.44164)	-0.00355 (0.00410)
600000 to 700000	0.25113 (0.40788)	-0.00151 (0.00441)	-0.02462 (0.24894)	0.00364 (0.00270)
700000 to 800000	0.12379 (0.37161)	0.00100 (0.00416)	0.13650 (0.26585)	0.00340 (0.00266)
800000 to 900000	0.07008 (0.39359)	0.00224 (0.00419)	0.18161 (0.29979)	0.00476 (0.00282)*
900000 to 1000000	-0.29398 (0.46877)	0.00294 (0.00469)	-0.65794 (0.33941)	0.00658 (0.00339)*
1000000 to 1100000	1.95390 (0.57745)***	0	1.02142 (0.39059)***	0
1100000 to 1200000	0.38174 (0.68556)	0.01954 (0.00577)***	0.77419 (0.47691)	0.01021 (0.00391)***
1200000 to 1300000	0.29954 (0.65681)	0.02336 (0.00550)***	-0.49885 (0.45059)	0.01796 (0.00375)***
1300000 to 1400000	-0.58272 (0.65944)	0.02635 (0.00536)***	0.14506 (0.43487)	0.01297 (0.00350)***
1400000 to 1500000	0.72161 (0.74929)	0.02052 (0.00565)***	0.44734 (0.52392)	0.01442 (0.00367)***
1500000		0.02774 (0.00626)***		0.01889 (0.00431)***
<i>Prediction at \$1,500,000</i>				
		<i>0.07892</i>	<i>0.04354</i>	

Notes: Data from REBNY listings matched to NYC DOF sales records. Estimates from regression of given discount on linear spline in initial asking price (\$100k knots between \$500k and \$1.5M). Slope estimates are for the given interval; predicted values are the difference between the predicted discount at the lower knot of the given interval and the predicted value at \$1,000,000.

3 Response of Home Equity Debt to Mortgage Policy: Evidence from a Kink and a Notch

Abstract

I estimate the market-level response of the size of home equity loans to two mortgage policies—the home mortgage interest deduction and regulations in Title XI of the Financial Institutions Reform Recovery and Enforcement Act requiring independent appraisals of borrowers’ homes at loan origination. Using administrative data on home equity loan originations for California, Illinois, New Jersey, and New York, I employ recent empirical methods that study non-linearities in individuals’ budget sets (“kinks” and “notches”). I find robust evidence of bunching of loan principal at the thresholds beyond which mortgage interest is not deductible and licensed appraisals of homes are required. The extent of this bunching at the two policy cutoffs translates to a 20% market-level reduction in loan size in response to removal of the mortgage interest deduction, and a 23% reduction in loan size in response to appraisal requirements. Using data from the Survey of Consumer Finances, I find that the response to the mortgage interest deduction corresponds to an elasticity of debt to interest rates around -30.

3.1 Introduction

In the wake of the financial crisis of the late 2000s an emphasis has been placed on rethinking home loan policy in the United States. Economists have identified several possible causes of the crisis, including lax underwriting standards and a boom in credit demand and availability (Mian and Sufi (2009); Frame et al. (2008); Dell’Ariccia et al. (2012); Mayer et al. (2009)). In response, as part of the Dodd-Frank Wall Street Reform and Consumer Protection Act, lawmakers and regulators are requiring banks to conduct more thorough reviews of loan applicants and greater independence of home appraisers from banks.⁹⁰ Moreover, there is a strong movement toward reconsidering current policies that lower the cost of mortgage debt, especially the Home Mortgage Interest Deduction (MID).⁹¹ However, although these policy changes may increase the likelihood that home loans are repaid, they also change the relative price of home equity debt and may reduce the availability of such loans.

In this paper, I estimate how the size (principal) of home-equity loans taken by homeowners responds to two mortgage policies: the home mortgage interest deduction and home appraisal requirements of the Financial Institutions Reform Recovery and Enforcement Act (FIRREA). By allowing mortgage payments to be deducted from taxable income, the MID lowers the cost of debt to home owners, which may induce borrowers to take larger loans than they otherwise would.⁹² Similarly, independent appraisals of homes being used as collateral, as required by Title XI of the FIRREA, not only increase the cost of underwriting loans, perhaps reducing the size or availability of debt, but may also reveal information about bor-

⁹⁰For example, Dodd-Frank requires that, prior to origination, lenders must verify loan applicants’ income, employment, current debt obligations, and credit history in order to determine that borrowers can reasonably repay mortgages.

⁹¹The primary argument for removal of the deduction is that the policy, which is a substantial tax expenditure, benefits relatively wealthy households who are not on the margin of homeownership (Bourassa and Grigsby (2000); Glaeser and Shapiro (2003); Gyourko and Sinai (2004); Poterba and Sinai (2011)). Moreover, there is mixed evidence of the efficacy of the MID in encouraging home ownership. Rosen (1979), Green and Vandell (1999), Bourassa and Yin (2008), Hilber and Turner (2012), and Hanson (2012) find that the MID does not increase home-ownership, and may even reduce it through general-equilibrium housing price effects. To date, there are no studies of how the MID influences the quantity of home equity debt taken.

⁹²The MID reduces the effective interest rate: for every dollar of interest paid, a portion of that dollar is returned to the borrower in tax savings.

rowers and their property that might increase or decrease lenders' willingness to issue a loan. I use administrative data on home loan originations and take advantage of discontinuities in the application of both policies to estimate whether and by how much the size of home loans falls when home owners cannot deduct their mortgage interest and when independent licensed appraisals are required at origination. Understanding the market-level response of loan principal to these mortgage policies helps inform how such regulations distort the market for home loans.

The desirability of regulating loan markets depends on weighing the benefits from increases in loan quality against the distortionary costs of the regulations themselves. On one hand, mandating lenders to adopt more stringent underwriting standards and/or increasing the cost of debt by eliminating deductibility may improve the quality of loans (e.g., Keys et al. (2012)) and lower the probability of housing-debt-driven crises in the future.⁹³ For example, Mayer et al. (2009) and Frame et al. (2008) document an increase in low- and no-documentation mortgages and smaller down payments leading up to the financial crisis. Jiang et al. (2014) find that low-documentation loans perform much worse than full-documentation loans.⁹⁴

On the other hand, tighter regulation of appraisals may discourage lending, and higher after-tax interest rates (associated with the removal of the MID) may discourage borrowing, reducing the availability of an important source of household debt. Home equity debt is a commonly used tool for consumption smoothing in the U.S. (e.g., Abdallah and Lastrapes (2012); Lovenheim (2013); Johnson (2012)). However, there is relatively little empirical work estimating how regulation of lending standards and the removal of tax deductions for interest payments affect the market for housing debt. Pence (2006) examines state boundaries and concludes that foreclosure laws that favor borrowers result in mortgages that are three- to seven-percent smaller than in lender-friendly states. Maki (2001) and

⁹³Since I rely on loan origination data for my study, I cannot comment on how the MID and FIRREA appraisal requirements affect loan performance.

⁹⁴Mian and Sufi (2009) and Dell'Ariccia et al. (2012) use aggregate data and find similar evidence of low quality loans contributing to the housing crisis.

Dunsky and Follain (2000) study the removal of interest deductions for consumer credit in the Tax Reform Act of 1986, and find substantial portfolio shifting to home debt (which maintained deductibility). Similarly, Hanson (2012) finds that mortgage debt drops by 10–18% at boundaries between states that do and do not allow for deduction of mortgage interest payments. To date, however, there have been no studies of the response of home-equity debt to the home mortgage interest deduction.

A key empirical challenge to studying home loan regulation is an absence of identifying variation—much of the policy is set at the national level, with relatively few changes over time. Existing research has relied primarily on cross-sectional or time-series variation in policy. For example, Keys et al. (2009) exploit cross-sectional variation in underwriting regulations faced by different lending institutions, while Pence (2006) uses cross-sectional variation at the state level. Similarly, Hilber and Turner (2012) and Hanson (2012) exploit cross-sectional variation in state tax policy to study the MID, while others have studied the removal of mortgage interest deductions internationally often using low-income households or renters as a control group: Jappelli and Pistaferri (2007) find no effect of removing the deduction in Italy, while Hendershott and Pryce (2006), Fjærli (2004), and Saarimaa (2010) all find reductions in mortgage debt in the U.K., Norway, and Finland, respectively.

To identify the effects of the MID and FIRREA appraisal requirements on home equity debt, I exploit discontinuities in the applications of these policies by loan size. For home equity debt, the MID allows only interest payments on the first \$100,000 worth of debt to be deducted (interest on the 100,001st dollar of debt is not deductible). This limit creates a discontinuity in the marginal after-tax interest rate at \$100,000, and this discontinuous change in the price of debt creates a kink in borrowers’ budget sets at the cutoff. Similarly, the FIRREA appraisal requirements are limited: independent licensed appraisals are only required on home equity loans over \$250,000, while below the thresholds banks may perform an in-house “evaluation,” creating a discontinuity (or policy “notch”) in both the information that banks are likely to collect and the cost to banks of underwriting a loan (appraisal fees

are typically passed on to consumers).

I adapt non-linear budget set methods to exploit the threshold designs of these two policies. A clear benefit of these empirical methods is that they facilitate the study of policies that do not vary across time or geography, as is the case with many mortgage regulations. These techniques are commonly used to estimate behavioral responses to taxation by examining “bunching” of individuals at discontinuities in marginal tax rates (“kinks” in budget constraints) and average tax rates (discontinuities or “notches” in budget constraints)—examples include Saez (2010), Chetty et al. (2011), Ramnath (2012), Kleven and Waseem (2013), Bastani and Hakan (2014), and Kopczuk and Munroe (2014). Unlike regression discontinuity, which assumes no manipulation of the running variable (debt, in this case) around the threshold, non-linear budget set methods explicitly study how economic agents sort to one side of the cutoff. Specifically, these methods ask the question of how responsive individuals must be to the given price or policy change at the threshold to induce the amount of bunching seen at the threshold. Recent work studying tax notches (Kleven and Waseem (2013) and Kopczuk and Munroe (2014)) focuses not only on the excess mass bunching at the threshold, but also distortions to the distribution just above the threshold. I argue herein that the response to a notched policy (such as the FIRREA appraisal requirement) can be estimated by comparing the size of the mass bunching at the threshold to a missing mass above the threshold.

I find evidence of a substantial reduction of home equity debt in response to the removal of the mortgage interest deduction and the imposition of the FIRREA appraisal requirements. Bunching at the \$100,000 and \$250,000 thresholds imply that loans are reduced by 20% in response to the removal of the MID and by 22% in response to appraisal requirements. These estimates are robust to functional form assumptions, and I find no evidence of a similar response at various placebo thresholds (including estimates at other round numbers as well as estimates at the same thresholds in the distribution of refinances for which these discontinuities do not apply).

Studying the dynamics of the response to the FIRREA threshold over the late 1990s and early 2000s reveals a limited relationship between this policy and the lending boom. I do not find any evidence that FIRREA appraisal regulations have a dampening effect on lending during the credit expansion; for example, if the policy is effective in sorting out lower quality borrowers, then we might expect a larger response to the policy during the lending boom when more low-quality borrowers were applying for loans.

In addition to understanding the distortionary effects of the MID on the size of home loans, the response of home equity debt to the MID reveals information about the elasticity of debt to interest rates, since the policy reduces the after-tax interest rate. This elasticity is of considerable value for understanding the welfare costs of policies that tax or subsidize savings and debt (Bernheim (2002)). However, there are few credible estimates of this elasticity for the U.S. (see Bernheim (2002) for a summary of older estimates and the problems with these). Gross and Souleles (2002b) study the response of credit card debt to quasi-exogenous increases in credit limits and interest rates and find an elasticity between about -4.9 and -9.1.⁹⁵ Studying the market for auto loans, Attanasio et al. (2008) estimate an elasticity of debt to interest rates between 0 and -14, depending on credit constraints.⁹⁶ Relatedly, there is an ongoing literature studying the response of household savings to various policies in the U.S. (e.g., Duflo et al. (2006); Madrian and Shea (2001); Poterba et al. (1996); Weber (2012)). My research contributes to this literature by using a novel identification strategy to provide an estimate of the market-level response of debt to lending policies for the U.S. for a different, but important, source of debt.

Converting my estimate of response to removal of the MID to a market-level elasticity of

⁹⁵In their paper, Gross and Souleles (2002b) present elasticities relative to a change in r , while myself, Attanasio et al. (2008), and others present elasticities relative to $1 + r$. I rescale the elasticities from Gross and Souleles (2002b) using the mean interest rate in order to compare them to elasticities relative to $1 + r$. For example, the estimates (with respect to r) from Gross and Souleles (2002b) range from 0.7 to 1.3 and the average interest rate that they report is 16.6%. At this interest rate, a 1% increase in $1 + r$ requires a 7.02% increase in r —thus, I rescale the estimates by a factor of 7.

⁹⁶Using an interest-rate offer experiment in South Africa, Karlan and Zinman (2008) estimate an elasticity around -4, while Dehejia et al. (2012) find slightly smaller estimates studying variations in interest rates across branches of lending institutions in Bangladesh.

debt with respect to interest rates yields a relatively large elasticity. I first establish that the response to the removal of the MID increases with proxies for the implied change in after-tax interest rates at the kink: estimates increase with adjusted gross income (individuals who face higher marginal tax rates benefit more from the MID), are positively correlated with market interest rates over time (more interest payments means a larger benefit from the deduction), and state-year estimates of response at the kink are positively correlated with imputed state-year averages of the after-tax interest-rate change at the kink. Secondly, using marginal tax and interest rate estimates from the Survey of Consumer Finances, I convert the response to the MID kink to an elasticity of debt with respect to after-tax interest rates of around -30.⁹⁷ Interpretation of this elasticity estimate as a structural parameter representing borrowers' preferences relies on the assumption that there is no supply-side response to the policy (I cannot observe lender behavior in my data) and that there is no shifting of debt to other sources.⁹⁸

3.2 Policies

3.2.1 The Home Mortgage Interest Deduction

The Home Mortgage Interest Deduction (MID), which has existed in the U.S. since the introduction of federal income taxes, allows taxpayers to deduct interest payments on loans secured by primary or secondary homes from their taxable income.⁹⁹ The MID is the second-most commonly claimed tax deduction in the U.S., although the benefits of the deduction mainly accrue to wealthier taxpayers (Glaeser and Shapiro (2003)). Moreover, the deduction is the second-largest tax expenditure for the federal government, representing about 79.2 billion dollars of foregone revenue in 2010 (Office of Management and Budget 2012).

⁹⁷While this elasticity may seem large, remember that this is the elasticity with respect to $1 + r$. A one-percent change in $1 + r$ is about a one percentage-point change in r . For example, if r is 5%, a one-percent increase in $1 + r$ is equivalent to a 21% increase in r .

⁹⁸While I can only observe home-equity debt in my data, I argue herein that debt shifting is not likely to be very important since interest rates on home equity debt tend to be among the lowest.

⁹⁹See Ventry Jr (2010) for a detailed history of interest deductions in the U.S.

The MID is applicable to all loans secured by a home, which includes mortgages, refinances, and home-equity loans (HELs) and home-equity lines of credit (HELOCs), and reduces the effective interest rate faced by borrowers.¹⁰⁰ For every dollar of interest paid on an MID-eligible loan, a portion of that dollar is returned to the borrower in tax savings. Specifically, for an individual with a marginal tax rate of t borrowing at an interest rate of $1 + r$, the after-tax rate of return faced by this borrower is $(1 + (1 - t)r)$.

There are limits on the applicability of the MID, which create discontinuities in the marginal after-tax interest rate (budget-set kinks). Under the MID, home loans fall into two categories that determine the limit on deductibility: home-acquisition loans and personal-use loans. To be considered a home-acquisition loan, a loan must be used to purchase a primary or secondary home (as in a first mortgage or a refinance of a first mortgage) or used to build or improve a home (as in an HEL or HELOC specifically used for home renovation/construction). Personal-use loans are all other loans secured by a home, such as HELs used for personal consumption or consolidation of debt. Interest on home-acquisition loans is only deductible on the first million dollars borrowed. The interest on any money borrowed in excess of one million is not deductible. Similarly, for personal-use loans, only interest on the first \$100,000 is deductible.¹⁰¹ For example, if an individual borrows \$50,000 for personal consumption, then all the interest is deductible. On the other hand, only 50% of the interest paid on \$200,000 of outstanding personal-use debt is deductible.¹⁰² These

¹⁰⁰The IRS defines a “qualified home” as a main or second residence, including condominiums, coops, mobile homes and trailers, house boats or “any similar property that has sleeping, cooking, and toilet facilities.”(Internal Revenue Service 2011)

¹⁰¹Both limits are halved for individuals who are married and filing separately, allowing each half of the couple to claim half of the interest.

¹⁰²The limits apply to the average outstanding balance held by an individual over the course of the tax year. As such, the limits at the time of borrowing will be somewhat higher than \$100,000 and \$1,000,000, depending on whether or not enough of the balance will be paid down in the year. However, I find that the bulk of response occurs at \$100,000. This is consistent with individuals having only moderate information about the deduction at the time of taking their loans—the MID and the limits of \$100,000/\$1,000,000 are well advertised, whereas the details of how an average balance is calculated are much less visible. Moreover, individuals may have difficulty calculating exactly where their limit will be given the proposed interest rates and repayment schedules (if interest rates and repayment schedules are not certain until the loan is finalized and there is a cost to mental arithmetic, individuals may prefer to go with the concrete limits of \$100,000 and \$1,000,000).

limits create a jump in the effective interest rates that borrowers face from $(1 + (1 - t)r)$ to $(1 + r)$ at the thresholds, and, in turn, discontinuities in the price of borrowing create kinks in individuals' budget sets. I examine borrower behavior around the jump in the price of home-equity loans to back-out the response of debt to the MID.

3.2.2 FIRREA Appraisal Requirements

The Federal Institutions Reform, Recovery, and Enforcement Act of 1989 was enacted to tighten regulations of lenders (especially thrifts) in response to the Savings and Loan crisis. A primary focus of the FIRREA was to restructure the regulatory system governing thrifts.¹⁰³ The FIRREA also tightened regulations on the origination of home loans, including appraisal requirements. The act raised capital requirements for thrifts and established concrete regulations on the appraisal of homes being used as collateral for loans. Title XI of the act, which I focus on, requires that real estate appraisals of homes being used as collateral in “federally related transactions” be conducted in writing by an independent, licensed, appraiser. Title XI also provides minimum standards for appraisals.¹⁰⁴

Appraisals are conducted to establish an estimate of the value of a home being used as collateral for a loan, to ensure that borrowers have sufficient equity in their homes. Title XI of the FIRREA requires that banks employ an independent licensed appraiser: the individual performing the appraisal must be entirely independent from all other aspects of the underwriting process and must not have any direct or indirect interest in the property or loan. Licensing requirements typically include completion of several courses, an examination, and two years of experience (for example, assisting licensed appraisers). By law, appraisals are required to conform to the minimum standards outlined by the ASB.¹⁰⁵ While

¹⁰³Dissolving the Federal Savings and Loan Insurance Corporation and putting the Federal Deposit Insurance Corporation (FDIC) in charge of the Savings Association Insurance Fund, replacing the Federal Home Loan Bank Board with the Office of Thrift Supervision (OTS), and establishing the Resolution Trust Corporation to facilitate the closure and liquidation of failing thrifts.

¹⁰⁴Federally related transactions includes home loans made by federally regulated financial institutions (i.e. regulated by the FDIC, OCC, OTS, or NCUA).

¹⁰⁵Title XI established the Appraisal Subcommittee of the Federal Financial Institutions Examination Council to oversee licensing and appraisal standards, which are the respective responsibilities of state regu-

licensed appraisals reveal important information for loan underwriting—namely, the value of the collateral—they are costly to lending institutions (around \$400–\$500 for a typical residence), although this cost is typically borne by the borrower.

Like the MID, the requirement of a licensed appraisal depends on the size of the loan. For loans less than \$250,000, lenders have the option of completing an “evaluation” in lieu of a licensed appraisal. Evaluations may be performed by the lending institution and do not require the property to be visited or inspected—lenders have discretion in the methods they choose to estimate a property’s value. Thus, the cost of underwriting a home-equity loan (and the information uncovered in the process) jumps discretely at the \$250,000 threshold. To the extent that banks pass these costs on to consumers, for example through higher closing fees, interest rates, or by altogether denying loans above the threshold, borrowers also face a discontinuity at the FIRREA limit (creating a discontinuity or “notch” in the borrower’s budget constraint). Unlike home equity loans, mortgage refinancing does not face the FIRREA threshold. If borrowers refinance home loans with the same lender, no appraisal is required under the FIRREA. Thus, mortgage refinancing provides a useful placebo against which to check the effects I find for home equity loans.

3.3 Data

To study the market-level response of home equity debt to regulations, I use administrative data on the universe of recorded home loans in recent years in California, Illinois, New

lators and of the Appraisal Standards Board (ASB) of the Appraisal Foundation (a non-profit established by the appraisal industry and authorized by congress to regulate appraisal standards). Appraisers are required to identify the property and its intended use (e.g., residence vs. development) and establish an appropriate value for this use. In doing so, appraisers must account for local regulations that may affect the value (zoning, environmental or historical preservation, etc.) as well as anticipated neighboring public or private developments that may affect the value. Appraisers must explicitly identify all assumptions s/he must make to complete the appraisal. Appraisals generally require a visit to the property and the value may be based on one (or more) of three approaches that they deem most appropriate: using the sales price of comparable properties, based on the cost of the proposed use of the property, or based on the potential income stream from renting the property. All details of the appraisal are required to be written up and returned to the lending institution.

Jersey, and New York.¹⁰⁶ I study county-level home equity loan (including home equity lines of credit) and mortgage refinance records collected from county clerks from at least 2002 through 2008, with the majority of counties going back through 1995 (all home loans and deed transfers in the U.S. must be recorded with county officials).¹⁰⁷ These records contain basic information about each home loan, including the loan principal (rounded to the nearest \$100), type of loan (refinance vs. home equity), date, and zip code of the home being used as collateral. This data provides accurate information on the date of origination, location of property, and loan principal—the variable that determines the applicability of the MID and FIRREA appraisal regulations with a sufficiently large sample to study distortions in the distribution of loans around regulatory thresholds.

Table 3.1 presents the breakdown of loans by type. After cleaning the data, I am left with 25,012,077 home loans.¹⁰⁸ Of these, there are about 16.6 million refinances and 8.4 million home-equity loans (including home-equity lines of credit). As would be expected, mean principal of mortgage refinances is larger than for home equity loans: \$386,240 for refinances (year 2000 dollars) versus \$157,287 for home equity debt. For both types of loans, median principal is smaller than the mean (e.g., \$61,606 vs. \$146,292 for home-equity loans), suggesting long right-tails to the distributions.

Looking at home equity loans over time illustrates the magnitude of the lending boom in the early 2000s. The upper panel of Figure 3.1 displays plots of the total amount of home equity debt at origination (in billions) and count of originations by year. There is a clear boom in lending in the early 2000s, as the total amount of debt originated per year increases from about \$50B per year in the late 1990s to more than \$200B in 2005. Similarly, the number of loans originated increases rapidly over this period from around 400,000 loans in

¹⁰⁶These data were collected from county records by an anonymous data provider and made available to me through the Paul Milstein Center for Real Estate at Columbia Business School.

¹⁰⁷I discuss county-level entry into the sample over time in the data appendix. My estimates are generally not sensitive to the inclusion of data prior to 2000 (when more than 75% of counties have data available).

¹⁰⁸In cleaning the data, I drop all duplicate loans (in address, date, and amount), loans with missing addresses, loans with loan amounts = 0 or missing, loans with non-rounded values, and all county-years where more than 20% of records have no associated address. See the data appendix for more details on the cleaning process.

2000 to 1.2M in 2005. At the same time, the beginning of the bust is evident—lending drops off after a peak in 2004/2005 falling somewhat through 2007.¹⁰⁹ The size of loans increases along with the volume of debt. The lower panel of Figure 3.1 shows that the median size of a home-equity loan grows rapidly in the early 2000s, doubling from \$50k in 2000 to \$100k in 2008.

Since the loan origination data does not include details about the borrower or about interest rates, I rely on the Survey of Consumer Finances to impute how the MID influences the cost of debt. The SCF is a nationally representative survey of households, conducted every three years. The survey asks detailed questions about households' income, wealth, and use of debt. Details about home equity loans collected in the survey include interest rates and loan principal. I restrict the sample to all households with a home equity loan, and pool this data for 1998, 2001, 2004, and 2007 (leaving a sample of 4144 households). Finally, I calculate marginal tax rates using NBER TAXSIM, accounting for the household's reported adjusted gross income, number of dependents, state of residence, and property taxes (Feenberg and Coutts (1993)).

3.4 Methodology

3.4.1 Estimating Response to Policy from a Kink

The discontinuity in marginal prices caused by a tax or policy kink (e.g., MID deductibility limit) may cause agents to bunch at this kink point. Existing work shows how the extent of bunching can be used to uncover the behavioral response to the given tax or policy (Saez (2010); Chetty et al. (2011); Weber (2012)). Intuitively, economic agents cluster just below the kink to avoid the change in policy—if agents are unresponsive to the marginal incentives that change at the threshold, then no one will bunch. The size of the mass bunching at the kink point is, thus, proportional to how responsive agents are to the policy itself. I briefly

¹⁰⁹Note that the number of originations in late 2008 is artificially low—many loans that originate in the latter half of 2008 will not have been recorded until 2009 and will not be present in the data set.

review the methodology herein.

Suppose that the MID applies everywhere (i.e., there are no limits to deductibility)—individuals will choose their optimal debt, which will generate some distribution of loan originations. Let \tilde{q} represent the optimal loan an individual takes when there is no limit to the deductibility of mortgage interest (and so the effective interest rate is $1 + (1 - t)r$ everywhere), and $f_0(q)$ represent the distribution of these loans. This distribution of debt may be generated by heterogeneity in income, need for debt (e.g., health shocks, children going to college), or preferences.

Following the argument outlined by Saez (2010), if the government imposes a deductibility limit at some level, q^* , individuals will either continue to choose debt equal to \tilde{q} or reduce their debt in response to the higher price, with some mass of borrowers reducing debt exactly to the kink point. At q^* , the marginal cost of debt increases discontinuously from $1 + (1 - t)r$ to $1 + r$. This change in price is represented for a simple two-period decision in Panel (a) of Figure 3.2, where q_t is debt taken for consumption in the first period, C_{t+1} is consumption in the second period, and Y is exogenously given second-period income. Let \check{q} denote the optimal choice of debt under this kinked policy. Borrowers with $\tilde{q} \leq q^*$ are unaffected by the kink and can achieve their optimal loan under the lower effective rate, $1 + (1 - t)r$, giving $\check{q} = \tilde{q}$.¹¹⁰ This is illustrated in Figure 3.2 as the individual borrowing $\tilde{q}_A = \check{q}_A$. When $\tilde{q} > q^*$ the optimal no-kink loan, \tilde{q} , is unattainable—individuals will reduce their demand for debt. In particular, for \tilde{q} sufficiently large, individuals will now choose $\check{q} = \phi\tilde{q}$, where $\phi > 0$ is a reduced-form parameter capturing the response of the market to the removal of the policy. This is illustrated in Figure 3.2 for the individual who moves from \tilde{q}_C to \check{q}_C . Note, however, that individuals who originally choose debt \tilde{q} “just above” the kink point will bunch at the threshold. Specifically, if $\tilde{q} \leq \frac{1}{\phi}q^*$, then borrowers will choose $\check{q} = q^*$; once $\phi\tilde{q}$ drops below the threshold, borrowers no longer face the higher marginal interest rate and choose the largest

¹¹⁰As discussed in Kopczuk and Munroe (2014), this relationship assumes that there are no general-equilibrium spillovers of the policy on the availability or cost of debt to those borrowing below the kink point.

loan (closest to \tilde{q}) that they can achieve without crossing the kink point. The “marginal buncher” appears in Figure 3.2 as the individual who moves from \tilde{q}_B to $\check{q}_B = q^*$. Notice that this creates a mass of individuals bunching at the kink point. Debt under the kinked policy can be summarized as:

$$\check{q} = \begin{cases} \tilde{q}, & \tilde{q} \leq q^* \\ q^*, & q^* < \tilde{q} \leq \frac{1}{\phi}q^* \\ \phi\tilde{q}, & \tilde{q} > \frac{1}{\phi}q^* \end{cases} \quad (17)$$

The corresponding distribution of debt under the kinked policy will feature bunching at the kink from which ϕ can be uncovered. Let $f(q)$ denote the empirical distribution of debt under the kinked policy. From the relationship between \check{q} and \tilde{q} given above:

$$f(q) = \begin{cases} f_0(q), & q \leq q^* \\ \int_{q^*}^{\frac{1}{\phi}q^*} f_0(x)dx, & q = q^* \\ \frac{1}{\phi}f_0\left(\frac{q}{\phi}\right), & q > q^* \end{cases} \quad (18)$$

Two important features of the empirical distribution are evident from this expression. Firstly, there is a mass bunching at the kink point: $B = \int_{q^*}^{\frac{1}{\phi}q^*} f_0(x)dx$. Secondly, there is a discontinuity in the distribution above q^* : $\lim_{q \rightarrow q^*_+} f(q) = f_0(q^*) \neq \frac{1}{\phi}f_0\left(\frac{q^*}{\phi}\right) = \lim_{q \rightarrow q^*_-} f(q)$. Thus, given an estimate of the excess mass bunching at the kink point, \hat{B} , and an estimate of the distribution $f_0(\cdot)$ above the kink point, an estimate of ϕ can be found by solving $\hat{B} = \int_{q^*}^{\frac{1}{\phi_k}q^*} \hat{f}_0(x)dx$ (I use the subscript k to denote the response estimated as for a kink (i.e., $\hat{\phi}_k$), and the subscript n to denote the response estimated as for a notch ($\hat{\phi}_n$), which I define below). In other words, the proportional response of debt to the removal of the MID can be estimated by answering the question: how much of the mass above the threshold under the counterfactual no-kink distribution is needed to explain the mass bunching at the kink?

As pointed out by Saez (2010), given an estimate of the function $f_0(\cdot)$, B can be estimated

by taking the difference between the observed volume of sales bunching at the threshold and the predicted volume of sales implied by $f_0(\cdot)$. In practice, I allow for individuals to bunch in a range of width $\underline{\delta}$ below q^* and estimate B accordingly: $\hat{B} = \int_{q^* - \underline{\delta}}^{q^*} (f(x) - \hat{f}_0(x)) dx$. The intuition is illustrated in Panel (b) of Figure 3.2. Bunching below the threshold is estimated as the distance between the observed distribution and the counterfactual. I then estimate $\frac{1}{\phi_k} q^*$ as the value of debt such that the integral under the counterfactual above the threshold equals the excess mass.

3.4.2 Estimating Response to Policy from a Notch

The observed distribution under a policy notch is similar to that of a policy kink, but has one important distinction—a gap above the threshold (see Kleven and Waseem (2013) for an in-depth discussion of policy notches). Because there is a discontinuity in the average cost of debt at the FIRREA notch, there may be a dominated region within which borrowers would prefer to bunch at the threshold, q^* , even if $\phi \tilde{q} > q^*$ (relying on the same notation from above, where $\tilde{q} \sim f_0(q)$ represents optimal debt in absence of the stricter FIRREA policy, and ϕ represents the response of debt above the threshold to the stricter underwriting standards).¹¹¹ In particular, there will be a borrower who chooses a value of debt in absence of the notch, \tilde{q}_m , such that under the notched policy the borrower is indifferent between borrowing q^* and avoiding the policy change or borrowing $\phi \tilde{q}_m$ under the stricter policy (notice that a kink is just a special case of the notch where $\tilde{q}_m = \frac{1}{\phi} q^*$). However, this borrower would always prefer to borrow q^* or $\phi \tilde{q}_m$ to $q \in (q^*, \phi \tilde{q}_m)$. This phenomenon is illustrated in Panel (a) of Figure 3.3, where notation is as in Figure 3.2 described above and the FIRREA policy increases the cost of debt by some fixed amount, k (although the discussion herein applies to the general case of a discrete jump in r). As with the kinked policy, those who initially

¹¹¹For expositional purposes, I discuss the response to policy as though lenders pass on the costs of the FIRREA appraisals to borrowers (anecdotally, they do) and there is no lender response. Nonetheless, the estimates outlined herein are general—in response to the stricter lending standards, debt (local to the threshold) will grow or shrink by some factor, $\phi > 0$, with the corresponding distortions to the distribution (although $\phi > 1$ will result in no bunching, but a depression around the threshold even for a kink). Thus, ϕ represents the market-level response to the changing policy.

borrow below the threshold do not change their behavior (q_A in Figure 3.3), while those borrowing well above the threshold reduce their demand for debt (q_C in Figure 3.3). Those who choose debt between q^* and \tilde{q}_m in absence of the notch will bunch at the threshold in the presence of the notch. Then, for the notched policy, the relationship between \check{q} and \tilde{q} is:

$$\check{q} = \begin{cases} \tilde{q}, & \tilde{q} \leq q^* \\ q^*, & q^* < \tilde{q} \leq \tilde{q}_m \\ \phi\tilde{q}, & \tilde{q} > \tilde{q}_m \end{cases} \quad (19)$$

The corresponding empirical distribution of debt is:

$$f(q) = \begin{cases} f_0(q), & q \leq q^* \\ \int_{q^*}^{\tilde{q}_m} f_0(x)dx, & q = q^* \\ 0 & q^* < q \leq \phi\tilde{q}_m \\ \frac{1}{\phi}f_0\left(\frac{q}{\phi}\right), & q > \phi\tilde{q}_m \end{cases} \quad (20)$$

Thus, there is a gap above the threshold with a width of $\phi\tilde{q}_m - q^*$. Like the kinked distribution, the distribution of debt under the notched policy displays bunching, equal to $B = \int_{q^*}^{\tilde{q}_m} f_0(x)dx$, and a discontinuity between the distribution below the notch and the distribution above the gap.

The response of debt to the FIRREA notch, ϕ , can be backed out by comparing the mass bunching at the threshold to the width of the gap. Notice that in the case of the notch, ϕ cannot be backed out from bunching alone. Rather, what is needed is knowledge of the right-hand edge of the gap, the point $\bar{q} \equiv \phi\tilde{q}_m$, which can be found by examining the gap, and of \tilde{q}_m itself, which can be inferred from bunching. In particular, given estimates of \hat{B} and $f_0(\cdot)$, an estimate of \tilde{q}_m can be found as solving $\hat{B} = \int_{q^*}^{\hat{q}_m} \hat{f}_0(x)dx$. Then an estimate of ϕ can be found by taking the ratio $\frac{\bar{q}}{\hat{q}_m}$.

In general, however, the gap above the threshold is not clean— \bar{q} cannot be observed from

inspection. Kleven and Waseem (2013) point out that in the presence of heterogeneity in the response to the policy (e.g., due to heterogeneity in preferences or in adjustment costs) the gap is not simply an empty space above the threshold. Rather, there will be a depression in the distribution above the notch that eventually disappears. In other words, there will be a region above q^* where the mass of loans is smaller than would be predicted by the distribution further to the right (i.e., over some range of width, Δ , $\int_{q^*}^{q^*+\Delta} f(x)dx < \int_{q^*}^{q^*+\Delta} \frac{1}{\phi} f_0(\frac{x}{\phi})dx$).

To estimate the value \bar{q} , I follow a similar procedure as that used to estimate \hat{q}_m . First, I estimate the distribution far above the threshold (to avoid bias from the gap). For notational ease, let $f_1(\cdot) \equiv \frac{1}{\phi} f_0(\frac{\cdot}{\phi})$, and $\hat{f}_1(\cdot)$ be an estimate of $f_1(\cdot)$. Second, I estimate the mass “missing” from the gap by taking the difference between the predicted mass were there no gap and the observed volume of sales for a fixed region of width $\bar{\delta}$ above the threshold: $\hat{G} = \int_{q^*}^{q^*+\bar{\delta}} (\hat{f}_1(x) - f(x)) dx$. Finally, I estimate \bar{q}_m by taking the integral under $\hat{f}_1(\cdot)$ up until the point that the integrated mass equals the size of the gap: $\hat{G} = \int_{q^*}^{\hat{q}} \hat{f}_1(x)dx$, and so $\hat{\phi}_n = \frac{\hat{q}}{\hat{q}_m}$. This is illustrated in Panel (b) of Figure 3.3. I estimate the mass of sales missing from the gap as the dark gray region below the dotted line ($f_1(q) = \frac{1}{\phi} f_0(\frac{q}{\phi})$) but above the solid line of the observed distribution. I then integrate under $f_1(q)$ to find the point \bar{q} where the integrated mass equals the missing mass.

3.4.3 Implementation

I approximate $f_0(\cdot)$ and $f_1(\cdot)$ by fitting a flexible global polynomial through the histogram of home equity loans.¹¹² There are several features of the empirical distribution that I accommodate when estimating these functions. First, since loan principal in the data is rounded to the nearest \$100 I treat this as a discrete distribution (i.e., count data). For every

¹¹²Fitting a global polynomial in this way is a commonly used approach to bunching estimation (e.g., Kopczuk and Munroe (2014); Chetty et al. (2011, 2013); Ramnath (2012); Kleven and Waseem (2013)). An alternative is to construct the counterfactual distribution locally by smoothing the histogram on either side of the bunching region (e.g., Weber (2012); Saez (2010)). While this non-parametric approach is appealing, it is impractical in cases where observations bunch at round numbers that coincide with the threshold—as they do in the case of home equity debt at multiples of \$25k—since it is not clear how to model this rounding non-parametrically.

\$100 interval of loan principal, q , I generate a count of the number of loans at that value, N_q . I use a poisson regression to relate these counts to a polynomial in loan principal.¹¹³ Second, similar to Kleven and Waseem (2013) and Kopczuk and Munroe (2014) I include a vector of dummy variables, $R_k = \mathbf{1}[q = z \cdot 100k : z \in \{1, 2, 3, \dots\}]$ for $k \in \{5, 10, 50, 250, 500\}$, to pick up bunching at multiples of \$500, \$1000, \$5000, \$25000, and \$50000. In the baseline specification, I interact these dummies with loan principal to further control for this rounding. Third, as noted above, there is a jump in the distribution at the policy thresholds (kink or notch). I allow for discontinuity at \$100,000 and \$250,000 by including an indicator for loan principal greater than each threshold, $D_t = \mathbf{1}[q \geq 1000t]$, $t \in \{100, 250\}$. I also allow for a constant discontinuity at all other multiples of \$50,000 to control for any possible jump in the distribution arising from the natural rounding that occurs at these numbers. Taken together, the specification is:

$$N_q = \exp \left(\beta_0 + \sum_{i=1}^j \beta_i q^i + \sum_{t \in \{50, 100, \dots\}} \delta_0 D_t + \delta_1 D_{100} + \delta_2 D_{250} + \sum_{k \in \{5, 10, 50, 250, 500\}} (\gamma_{0k} R_k + \gamma_{1k} R_k q) + \epsilon_q \right) \quad (21)$$

In my baseline specification I use a second-order polynomial ($j = 2$). However, I explore the sensitivity of my estimates to lower and higher order polynomials as well as varying the order of the interaction of q and R_k . For all specifications, I restrict the sample to loans between \$30k and \$475k.

I omit data around the MID and FIRREA thresholds so that the estimates are not influenced by distortions near these cutoffs. In particular, I omit all data between \$10,000 below and \$40,000 above each threshold when I estimate Equation 21 (i.e., I omit data between \$90,000 and \$140,000 and between \$240,000 and \$290,000). This restriction is important for two reasons. First, if bunching in response to a notch is imprecise, the distribution may be

¹¹³The estimates are generally insensitive to using the poisson regression vs. a log-linear regression, although the log-linear specification tends to be less robust. Moreover, the log-linear regression performs poorly when sub-setting the data—small subsets may have several \$100 bins with no data, which the log model cannot accommodate.

elevated below q^* and including this data would bias the estimate of $f_0(\cdot)$, which is based on the distribution to the left of each cutoff. Admittedly, this is unlikely to be a large problem in this case since there are few frictions in determining the value of a loan (as compared to bunching of taxable income—e.g., Saez (2010); Chetty et al. (2011, 2013); Kleven and Waseem (2013)). Second, since the response to a policy notch may open up a gap in the distribution above q^* , a range above the threshold should be omitted from estimation of the no-notch counterfactuals. I choose the omitted region by inspection of the graph of the log distribution of home equity loans and check the robustness of my estimates to the size of this region.¹¹⁴

I use the estimate of $f_0(\cdot)$ from Equation 21 to estimate the excess mass bunching below and gap above each threshold. Given an estimate of Equation 21, I define the counterfactuals for threshold q^* as: $\hat{f}_0(q) = \hat{N}_q$ for $q \leq q^*$, $\hat{f}_0(q) = \frac{\hat{N}_q}{\exp(\hat{D}_{q^*})}$ for $q > q^*$, and $\hat{f}_1(q) = \hat{N}_q$ for $q > q^*$.¹¹⁵ I find the excess bunching at each threshold by comparing the observed mass to the predicted mass in the omitted region below q^* (\$10,000 in the baseline). Specifically, I estimate:

$$\hat{B} = \sum_{q^*-10000}^{q^*} \left(N_q - \hat{f}_0(q) \right) \quad (22)$$

I find the gap above the threshold in the same way. Since the gap is the missing mass below the predicted distribution on the right of the threshold, I take the difference between the predicted density and observed data in the omitted region above the threshold (\$40,000 in the baseline):

$$\hat{G} = \sum_{q^*+100}^{q^*+40000} \left(\hat{f}_1(q) - N_q \right) \quad (23)$$

¹¹⁴Since it is not necessarily the case that the excess mass at the threshold will equal the missing mass in the gap, I cannot use the excess to determine where the end of the omitted gap region should be (unlike Kleven and Waseem (2013)). Estimates of the gap region tend to be small, rarely larger than \$10000 past the threshold.

¹¹⁵Recall that the discontinuity, \hat{D}_{q^*} , is included to allow for the shift in the distribution in response to the policy. To achieve a proper counterfactual for $f_0(\cdot)$ above the threshold, I must remove the discontinuity at q^* . Correspondingly, I define $\hat{f}_1(q) = \hat{N}_q \cdot \exp(\hat{D}_{q^*})$ for $q < q^*$, although this is only relevant in cases where the missing mass in the gap is negative.

I estimate \tilde{q}_m and ϕ_k by integrating under the relevant counterfactuals. To estimate \tilde{q}_m , I integrate under the counterfactual density, $f_0(\cdot)$, up until the point where the integrated mass equals the mass bunching at q^* . Specifically, I find the value, \hat{q}_m , that satisfies:

$$\hat{B} = \sum_{q=q^*+100}^{\hat{q}_m} \hat{f}_0(q) \quad (24)$$

Given this estimate of \tilde{q}_m , I find an expression for $\hat{\phi}_k$ as $\hat{\phi}_k = \frac{q^*}{\hat{q}_m}$. Notice that for a kinked policy change, $\tilde{q}_m = \frac{1}{\phi} q^*$, and so this estimate of $\hat{\phi}_k$ corresponds to the estimate given above.

As discussed above (and at length by Saez (2010); Weber (2012)), in the case of a kink, $1 - \hat{\phi}_k$, which is the quantity I estimate and present below, can be interpreted as the percent-reduction in debt in response to the higher marginal cost of borrowing above the kink for the marginal individual who moves to the threshold. In turn, this quantity can be used to estimate the elasticity of debt with respect to interest rates. As pointed out by Kleven and Waseem (2013) and Kopczuk and Munroe (2014), in the case of a notch $1 - \hat{\phi}_k$ has an intuitive interpretation, representing the average reduction in debt for marginal borrowers (i.e., those indifferent between $\phi\tilde{q}$ and q^*) who bunch at the threshold—how much is the marginal borrower willing to reduce debt to avoid the stricter policy? This stands in contrast to the value $1 - \hat{\phi}_n$, which represents the general response to the policy—how would the marginal borrower respond if the stricter policy applied everywhere?

Given the estimate of \hat{G} , I find \bar{q} by integrating under the predicted counterfactual for $f_1(\cdot)$ above the threshold and compare this to the estimate of \tilde{q}_m . I find \hat{q} as the value that satisfies:

$$\hat{G} = \sum_{q=q^*+100}^{\hat{q}} \hat{N}_q \cdot \exp(D_{q^*} \cdot \mathbf{1}[q < q^*]) \quad (25)$$

Using my estimates of \tilde{q}_m and \hat{q} , I arrive at my estimate of the global response of debt to the increase in borrowing costs above the threshold for the marginal borrower: $\hat{\phi}_n = \frac{\hat{q}}{\tilde{q}_m}$. Again, I present below the quantity $1 - \hat{\phi}_n$, which can be interpreted as the general percent reduction in debt in response to the stricter appraisal requirements. Of course, this parameter is specific

to those individuals borrowing in the neighborhood of the notch. I bootstrap standard errors for my estimates by resampling the original data and repeating the entire procedure 999 times.¹¹⁶

The identifying assumptions underlying this procedure are standard for bunching estimates. First, I assume that the observed data can be used to construct the counterfactuals $f_0(\cdot)$ and $f_1(\cdot)$. In particular, what this assumes is that the presence of the policy discontinuities does not create any general-equilibrium effects that distort the distributions away from the thresholds (in which case the observed data away from the threshold would not reflect the distributions $f_0(\cdot)$ below the threshold and $f_1(\cdot)$ above). Secondly, I assume that $f_0(\cdot)$ and $f_1(\cdot)$ can be reasonably approximated by fitting a global polynomial through the distribution, omitting data near the thresholds and allowing for discontinuity. I check the sensitivity of this assumption by experimenting with different functional forms. Third, I assume that the distribution, $f_0(\cdot)$ would be smooth in absence of the policy change at the threshold—if there is no kink or notch in the budget set then there is no bunching. In this particular case, this requires the assumption that the round number bunching at other multiples of \$50,000 can be used to construct an appropriate counterfactual for the excess mass at \$100,000 and \$250,000. I check the robustness of this assumption by estimating the response for placebo cutoffs at other multiples of \$50,000. I also estimate placebos using data on mortgage refinances where neither policy threshold exists.

3.5 Results

3.5.1 Graphical Evidence of Response

The empirical distribution of home equity loans displays distortions at \$100,000 and \$250,000 that are consistent with a response to the policy discontinuities. As outlined in Section 3.4,

¹¹⁶Owing to the very large sample size and a fairly stable distribution, my estimates of $1 - \hat{\phi}_k$ and $1 - \hat{\phi}_n$ are all very precise, and are significant at the 1% level. This is true whether I resample the raw data or whether I resample at a more aggregate level, for example resampling by zip code or year (i.e., randomly choosing a zip code and all the observations associated with it).

if the market for loans is sensitive to the removal of the MID at \$100,000 or the requirement of independent licensed appraisers above \$250,000, then there should be bunching at these threshold, with a possible depression (gap) in the density just above the thresholds, and a shift in the distribution further above the threshold. Figure 3.4 presents the histograms of home equity loan and mortgage refinance originations with \$5,000 bins. The distribution displays evident round-number bias; loans cluster at multiples of \$50,000. Despite this, there appears to be excess bunching of home equity loans at \$100,000 and \$250,000—the excess mass at these thresholds is visibly larger than neighboring multiples of \$50,000. Owing to the scale of this figure, the gap above the threshold is not apparent.

Taking the log of the histogram with \$1,000 bins, as in the upper panel of Figure 3.5, makes the distortions in the distribution at \$100,000 and \$250,000 more apparent. While taking logs compresses the excess mass at the round numbers, the bunching at \$100,000 and \$250,000 still appears out of line with that at other multiples of \$50,000. Moreover, a depression in the distribution is evident just above each policy threshold. At the same time, only minimal depressions can be seen above other multiples of \$50,000. As discussed above, I allow for discontinuity at these other round numbers to capture any distortions due to rounding at salient numbers. Finally, a salient feature of the log-histogram is that there is rounding in loan principal at many levels, not just multiples of \$50k. Hence, I include fixed effects to control for distortion at round numbers when estimating the counterfactual. I also estimate placebos at these other round numbers and find small or negative estimates.

One concern is that bunching at the policy thresholds is the result of behavioral rounding on the part of consumers at particularly salient numbers—perhaps \$100,000 and \$250,000 are special and the bunching is unrelated to the policies. If this is the case, however, we might expect similar behavior in the distribution of refinances. Refinances are loans taken to replace original mortgages, typically with better terms. It is common for borrowers to take a loan larger than the existing mortgage, using the difference both to pay for closing fees associated with the new loan, and often for consumption (similar to a home equity loan).

Thus, we would expect refinances to display at least some of the rounding behavior of home equity loans. The histogram for refinances over the same period in the same states is also plotted in Figure 3.4. While general bunching at multiples of \$50,000 is present, \$100,000 and \$250,000 do not appear to be outliers. The lower panel of Figure 3.5 displays the log of the histograms of refinances, which appears smooth throughout; there are no depressions above \$100k or \$250k. Despite clustering of loans at salient round numbers, there is no evidence of distortion at \$100,000 or \$250,000 in the distribution of refinances.

3.5.2 Baseline Estimates and Robustness

Using the approach described in Section 3.4.3, I find a significant market-level response of home equity loans to the removal of the MID and to FIRREA appraisal requirements. The first row of Table 3.2 presents the baseline estimates of response at both thresholds.¹¹⁷ The first two columns present the “kink” estimates $(1 - \widehat{\phi}_k)$ for each policy change—the percent reduction in the marginal loan that bunches at the threshold—while the third column presents the “notch” estimate $(1 - \widehat{\phi}_n)$ for the \$250,000 cutoff. Recall that for the kinked policy at \$100k, the kink estimate can be interpreted as the general response to the policy for the marginal borrower, while for the notched policy at \$250k, the notch estimate represents the general response (for loans in the neighborhood of \$250k) were the tighter appraisal requirements applied everywhere.

The baseline estimates suggest substantial reductions in loan size in response to the policies. The estimates imply that loan principal drops by 19.96% at \$100,000 in response to removal of the MID. This suggests that the deductibility of interest payments encourages home owners to borrow substantially more than they might otherwise. For the FIRREA, the kink estimates imply that marginal borrowers face a reduction of about 23.53% to the

¹¹⁷In the baseline I fit a global 2nd-order polynomial through the count of home loans (in \$100 bins) by poisson regression, omitting data between \$90,000 and \$140,000 and \$240,000 and \$290,000. I allow for discontinuity in the distribution at each of these thresholds as well as a constant discontinuity at each multiple of \$50k. I include fixed effects for multiples of \$500, \$1000, \$5000, \$25000, and \$50000 as well as interacting these fixed effects linearly with loan principal.

threshold in order to avoid the more stringent appraisal requirements above \$250,000. The notch estimates imply that when borrowers do not (or cannot) avoid the policy by moving below the threshold, the stricter appraisal requirements reduce the loan principal by 22.02%. Clearly, these threshold policies distort the decisions of agents in this market—a substantial amount of debt is foregone due to the policy changes. After discussing the robustness of these estimates in the remainder of this section, I return to the discussion of these estimates in the following two sections.

The estimates are insensitive to the specification of the model. I vary the order of the polynomial in loan principal (j in Equation 21) and display the estimates in the second panel of Table 3.2. At \$100,000, the kink estimate varies between about 16.6% (for 1st, 5th, and 6th order) and 20% (for 3rd and 4th order). Estimates at the \$250,000 threshold are a bit more variable—ranging from a local reduction of 16.7% to 27% and between a 15% and 25.5% general response to the policy. Despite the differences in magnitude, the estimates offer a qualitatively similar conclusion of a large reduction in debt in response to the policies. I also vary the order of the interaction between the round-number dummies and loan principal and display the estimates in the third panel of Table 3.2. Again, the estimates are somewhat sensitive, although consistently display a large response of similar magnitude to the baseline. The estimates are much less sensitive to varying the width of the omitted region around q^* —see the third and fourth panels of Table 3.2. I estimate the model allowing only for discontinuity at \$150,000 and \$200,000 to be used as a counterfactual for the two thresholds, since visual examination of the histogram suggests that the discontinuity is more prominent here than at other multiples of \$50,000. I report these estimates at the bottom of Table 3.2. This procedure results in a slightly larger estimate of the response at \$100,000 (23%) and a slightly smaller estimate at \$250,000 (18%).

I find no evidence that the response at \$100,000 and \$250,000 is driven by the salience of large round numbers or is an artifact of the empirical approach itself. A concern is that rounding at salient numbers—multiples of \$50,000, in this case—is difficult to model.

Since I am using the round-number bunching at other multiples of \$50,000 to construct a counterfactual for \$100,000 and \$250,000, the estimates I find could be spurious if this round-number counterfactual is incorrect. To test this assumption, I estimate placebos at other multiples of \$50,000 where there are no policy thresholds. Specifically, I estimate the baseline specification (Equation 21) treating the given placebo cutoff as the policy threshold and report these estimates in Table 3.3. The kink estimates in the first column (the first two rows replicate the baseline estimates for the sake of comparison) are consistently negative (with the exception of local response for \$200,000)—the empirical strategy tends to overstate the excess mass at the threshold rather than understate. These negative estimates likely occur because of, and provides additional evidence of, excessive bunching at \$100,000 and \$250,000 that overstates the round-number counterfactuals for these estimates. While still negative, notch estimates at the placebo cutoffs, presented in the third column, tend to be close to zero (with the exception of a negative estimate at \$150,000). Relative to these placebos, the baseline estimates at \$100,000 and \$250,000 are clear outliers.

A similar concern is that \$100,000 and \$250,000 are “special” round numbers—individuals (or lending institutions) round their loans to these numbers more than they would at other multiples of \$50,000, making the round number placebos uninformative. To address this, I estimate placebos using the distribution of mortgage refinances. Since the MID cutoff only applies to non-home-acquisition debt and since independent licensed appraisals are generally not required when mortgages are refinanced, there should be no unusual bunching or distortion around \$100,000 and \$250,000 in the refinance distribution. The second and fourth columns of Table 3.3 display these placebo estimates for refinances at the various cutoffs. Even though Figure 3.4 shows bunching at round-numbers in the distribution of refinances, the estimation procedure finds no evidence of excess bunching or distortion at \$100,000 or \$250,000 (or the other multiples of \$50k): the kink and notch estimates range between -2.9% and 4.7%. Thus, it does not appear to be the case that \$100,000 and \$250,000 are “special” round numbers in the context of home loans.

3.5.3 Response of Debt to the Mortgage Interest Deduction

The kink estimates for the MID threshold imply a large (about 20%) reduction in debt in response to removal of the deduction. This response suggests that, at least for those taking home equity loans in the neighborhood of \$100k, eliminating the MID will induce a substantial reduction in the quantity of housing debt that individuals take on (or, conversely, that the presence of the MID induces individuals to borrow substantially more than they otherwise would). Of course, understanding the economic magnitude of this response depends on how large the change in after-tax interest rates is when the MID does not apply. In this section, I show that the response at the MID threshold covaries with marginal tax rates and interest rates, and, thus, the magnitude of the kink. I then present implied estimates of the elasticity of home-equity debt with respect to the after-tax interest rate.

I investigate whether the response to the MID threshold varies with the size of the kink (i.e., the magnitude of the change in after-tax interest rates) by comparing the size of the response to local measures of income. At the threshold, the effective interest rate on the marginal dollar of debt increases from $1 + (1 - t)r$ to $1 + r$; the absolute size of this reduction ($-tr$) is increasing in both the marginal tax rate and the interest rate. Correspondingly, the kink estimate at \$100,000 should be larger when this change in effective rates is larger. To explore this, I estimate the percent drop in debt by quartile of adjusted gross income (AGI), using this as a proxy for marginal tax rates—individuals in higher-income neighborhoods will, in general, face higher marginal tax rates. I match home equity loans in each zip code to data from the IRS Statistics of Income (SOI), which reports mean zip-code-level AGI among all tax filers within the zip code (the IRS SOI is only available for 1998, 2001, 2002, and 2004–2008). Within each year, I find the quartile of AGI to which each zip code belongs. I then group all loans originated in each quartile and estimate the response to the MID kink as in the baseline specification and present these results in Table 3.4. I find that the estimated reduction of debt at \$100k is increasing with AGI, from a reduction of 15.22% in the lowest quartile to 19.94% in the highest. These effects suggest that the magnitude of

the response may correlate with tax rates. Reassuringly, when I repeat this exercise (also in Table 3.4) for the \$250,000 threshold—where the magnitude of the policy change does not vary with tax rates or income—the estimates do not display the same increasing pattern.

I also find that the response of debt to removal of the MID moves over time with national interest rates. As with the marginal tax rate, a higher contract interest rate means a larger change in after-tax interest rates at \$100,000. In Figure 3.6, I plot annual estimates of $1 - \phi_k$ against the national interest rate for a 30-year fixed rate mortgage.¹¹⁸ In general, response to removal of the MID moves with the interest rate. Conversely, we see that the same baseline estimates for the \$250k cutoff (where the size of the notch does not depend on interest rates) do not move with the rates and tend to grow over time. While there are certainly unobservable factors that are changing over time and influence the borrowing decision, this co-movement of the estimates with the interest rate and the positive correlation between the response and zip-code-level AGI is consistent with bunching at \$100,000 being caused by the removal of the MID.

To investigate these relationships more formally, I relate state-year estimates of local response to state-year averages of the imputed change in after-tax interest rates from the SCF. For each state (CA, IL, NJ, NY) and each year from 1995 through 2007, I estimate $1 - \phi_k$ as in the baseline specification.¹¹⁹ I then take the full sample of home-equity loan holders from the SCF (pooling the 1998, 2001, 2004, and 2007 survey years) and use NBER TAXSIM to impute the marginal tax rate for this sample were they residing in each state in each year. By using the nationally representative sample, I isolate variation in marginal tax rates due to differing state policies, rather than household selection into states/years (Currie and Gruber (1996)). I then match the state-year tax rates to the annual 30-year fixed mortgage rate, and calculate the average percent-change in after-tax interest rates at the kink. Variation in

¹¹⁸Mortgage rates are from the Freddie Mac mortgage rate survey (www.freddiemac.com/pmms/index.html). I use the 30-year fixed rate because Freddie Mac does not collect data on home equity loan rates. Nonetheless, home loan rates covary over time and so I do not expect the annual pattern to differ substantially for home equity rates.

¹¹⁹As discussed in the data appendix, not all counties have data available in 1995 through 2000. I find similar results for these state/year regressions using only 2000-2007, although the estimates are less precise.

the rate change is thus coming from two sources: changes in interest rates over time at the national level, and differences across states and over time in marginal tax rates.¹²⁰ Finally, I regress the kink estimates for each state/year on the average imputed after-tax interest rate change. I find standard errors by bootstrapping the entire procedure 999 times (resampling SCF data, estimating imputed change in after-tax interest rates, resampling home equity loans, estimating state/year response, and then regressing estimates on imputed change).

Regressing state/year estimates at \$100k on the mean change in interest rates shows a positive relationship between $1 - \widehat{\phi}_k$ and the change in $1 + (1 - t)r$. I present the results of these regressions in Table 3.5. The first column presents a bivariate regression of the state/year estimates on the state/year proxies for the change in after-tax interest rates and shows a positive relationship between the two—a larger effect of the removal of the MID on the marginal cost of borrowing means substantially more bunching at the kink. The second and third columns include state, and state and year fixed effects, respectively, and again show a large positive relationship.

I repeat this regression exercise for kink estimates at the \$250k cutoff to confirm that this correlation between bunching and the interest rate change is not spurious. Since the size of the notch at the \$250k cutoff does not depend directly on interest rates or tax rates, there is no reason to expect a positive relationship between the imputed MID interest rate change and the \$250k kink estimates. Indeed, the estimates, presented in the fourth through sixth columns of Table 3.5, show that there is not a positive relationship between the state/year estimates of $1 - \widehat{\phi}_k$ and the imputed change in rates at the \$250k cutoff.

Under the assumption that only the demand side of the market responds to the removal of the MID at \$100k, the response of debt at the kink can be reinterpreted as the elasticity of home equity debt with respect to interest rates (Saez (2010)). In particular, the baseline estimate suggests that the marginal “buncher” reduces his/her debt by 19.96% to arrive at

¹²⁰An alternative is to use the reported interest rates in the SCF for each household, although this eliminates any variation over time in interest rates (since households are only surveyed once and the survey is only conducted every three years). Estimates using this approach are similar (although larger) than using the 30-year fixed rate.

the MID threshold of \$100,000. Given an estimate of the change in the after-tax interest rate, this estimate of $1 - \phi_k$ can be converted to an elasticity. Since the loan records used in estimation provide no information on interest rates, I turn instead to the SCF. I calculate marginal tax rates for each household with a home equity loan using NBER TAXSIM, as discussed in Section 3.3. For each household, I find the implied change in the after-tax interest rate from a removal of the MID ($\frac{tr}{1+(1-t)r}$) using the reported interest rate on their home equity loan and the imputed marginal tax rate.

The elasticity of debt with respect to interest rates implied by the response to the removal of the MID is large. Table 3.6 displays the interest rate and tax rate statistics from the SCF as well as the implied elasticity. Taking the raw change in after-tax interest rates for the home-equity loan households in the SCF, the implied elasticity is -10.505. Restricting the sample to households with loans in the neighborhood of the threshold (between \$75,000 and \$125,000) raises the implied elasticity to -12.475—while these households typically get lower interest rates, they also face higher tax rates.

Although these estimates are on par with existing studies of the elasticity of debt with respect to interest rates, they likely understate the true elasticity. Since the MID threshold of \$100,000 only applies to non-home-acquisition debt, loans being used for home improvements will not be affected by the threshold. However, I cannot distinguish these loans in the origination data. Thus, my sample is contaminated by such borrowers for whom the kink does not apply and so are unresponsive to the threshold. To account for this, I use a question from the SCF that asks borrowers what the primary use of their home equity loan is. I calculate the share of home equity loans used primarily for home improvement as reported by the household (52.5% to 60.5%, depending on the size of the loan) and adjust the average implied change in after-tax interest rates, assuming no change for these home-improvement households (see Table 3.6).¹²¹ This adjustment of the imputed change in interest rates implies

¹²¹The reports in the SCF may overstate the number of home improvers if at loan origination individuals are either not aware of the exemption of limits for home acquisition loans or are not sure what they will use the money for. The true elasticity likely lies between the unadjusted and adjusted estimates.

substantially larger elasticities: -28.51 and -39.92 for the two samples.¹²²

These adjusted estimates suggest a larger elasticity of debt with respect to interest rates than has been found in existing studies. There are several reasons this may be the case. Firstly, it may be that information about the cost of debt is more visible when borrowers take home equity loans (e.g., required monthly payments are typically calculated for the borrower) than credit card loans (as in Gross and Souleles (2002a)). Secondly, it may be that individuals taking \$100,000 home equity loans are less credit constrained than those taking credit card loans or auto loans (as in Attanasio et al. (2008)). Both Gross and Souleles (2002a) and Attanasio et al. (2008) find larger elasticities for individuals facing lower constraints. Thirdly, it is possible that borrowers are shifting the marginal dollars of debt to non-housing debt, and so this is not the response of all debt with respect to interest rates, but of home equity debt to interest rates. However, contract interest rates for home equity debt are generally lower than other sources of debt. In the SCF, the average APR for housing debt is 7.9%, while the average for other lines of credit is 8.5% and for credit cards is 13.9%.¹²³ Finally, it may be that lenders are themselves responding to the presence of the policy kink, perhaps by restricting access to or increasing interest rates for loans greater than \$100,000. Unfortunately, I do not observe lender behavior. Thus, a conservative interpretation is that this is a market-level elasticity of debt to interest rates rather than a structural parameter of borrowers' preferences.

¹²²However, as noted in the introduction, elasticities expressed with respect to $1 + r$ tend to appear quite large, in part because a small change in $1 + r$ constitutes a large change in r , and consumers generally observe r , although $1 + r$ is the relevant price in theory (Bernheim (2002)). Expressing my elasticities as the change in debt relative to a percent-change in r gives -1.584 and -1.976 for the two samples (as compared to elasticities around -0.7 or -1.3 with respect to r found in Gross and Souleles (2002b)).

¹²³Brito and Hartley (1995) point out that even if rates are lower, if there are large origination fees for home equity loans and uncertainty about the future, borrowers may prefer other sources of debt. However, borrowers bunching at \$100,000 are already incurring the origination fees, and so this is unlikely to be an issue in this case.

3.5.4 Response of debt to FIRREA Appraisal Requirements

The baseline estimates at the \$250,000 threshold imply that stricter appraisal requirements reduce the size of loans. The estimate of local response to the \$250,000 FIRREA notch implies that loans are reduced by 23.53% in order to avoid a licensed appraisal. The general response is comparable—just above the threshold, there is a 22.02% drop in principal relative to how much debt would have been taken in absence of the policy. Taken at face value, requiring licensed appraisals of home loans appears to induce a substantial market-level response of home-equity debt.

This response is large relative to the monetary cost of complying with the tighter appraisal standards (around \$400 or \$500). However, it may be the case that there are non-pecuniary costs to an appraisal. For example, a licensed inspection and appraisal of a home may cause a costly delay in closing the loan (e.g., scheduling the appraisal, waiting for the report, etc.). Moreover, the formal appraisal introduces uncertainty into the lending process if there is a chance that the appraisal will return a home value different from the borrower and lender's prior (and, in turn, this different home value may affect the size of the loan available to the borrower). Finally, an appraisal might reveal information about lower quality borrowers that may cause lenders to increase interest rates or deny the loan altogether. Similarly, if mortgage originators are planning on selling loans on the secondary market, as was common during the lending boom of the early 2000s, they may avoid licensed appraisals in cases where the appraisal might reveal information that lowers the value of the mortgage.¹²⁴

However, the reduction in debt from appraisal requirements may be a desirable outcome— if tighter appraisal standards are costly to low-quality borrowers (or lenders seeking to sell low-quality mortgages on the secondary market), then this suggests that the policy is effective in identifying risky loans. While this will have the effect of reducing access to debt for these marginal borrowers, it may reduce the extent of subsequent loan defaults and foreclosures.

¹²⁴Recent research has focused on whether mortgage securitization led to lax underwriting standards; see Bubb and Kaufman (2009), Elul (2011), Jiang et al. (2010), and Keys et al. (2010).

If requiring licensed appraisals successfully weeds out risky borrowers and there are relatively more such borrowers taking loans during a lending boom, then we would expect the response of loans at the notch to be cyclical—costs are higher to these borrowers or lenders hoping to sell the loans on the secondary market. On the other hand, if the response to the FIRREA notch is purely due to the cost of appraisals, we would expect little movement in the response over time. However, I find the opposite—response at the threshold appears counter cyclical during the lending boom of the 2000s. I estimate the general response to the FIRREA policy at the \$250k threshold for each year from 1995 to 2008 and plot the estimates along with the total value of home equity loan originations (in billions of dollars) per year on Figure 3.7. The percent reduction in debt in response to the tighter appraisal requirements is negatively correlated with the number of home loans.

The FIRREA appraisal requirements may fail to dampen lending during a boom because the policy is ineffective in sorting out marginal borrowers. For example, it could be that appraisals exclusively reveal home value and nothing about borrowers' ability to repay (and during the boom home values were increasing sufficiently rapidly that this information was not a strong constraint on lending). Or it may be that during the boom, lenders concealed the information revealed in appraisals from the secondary market (e.g., fraudulent appraisals).

3.6 Conclusion

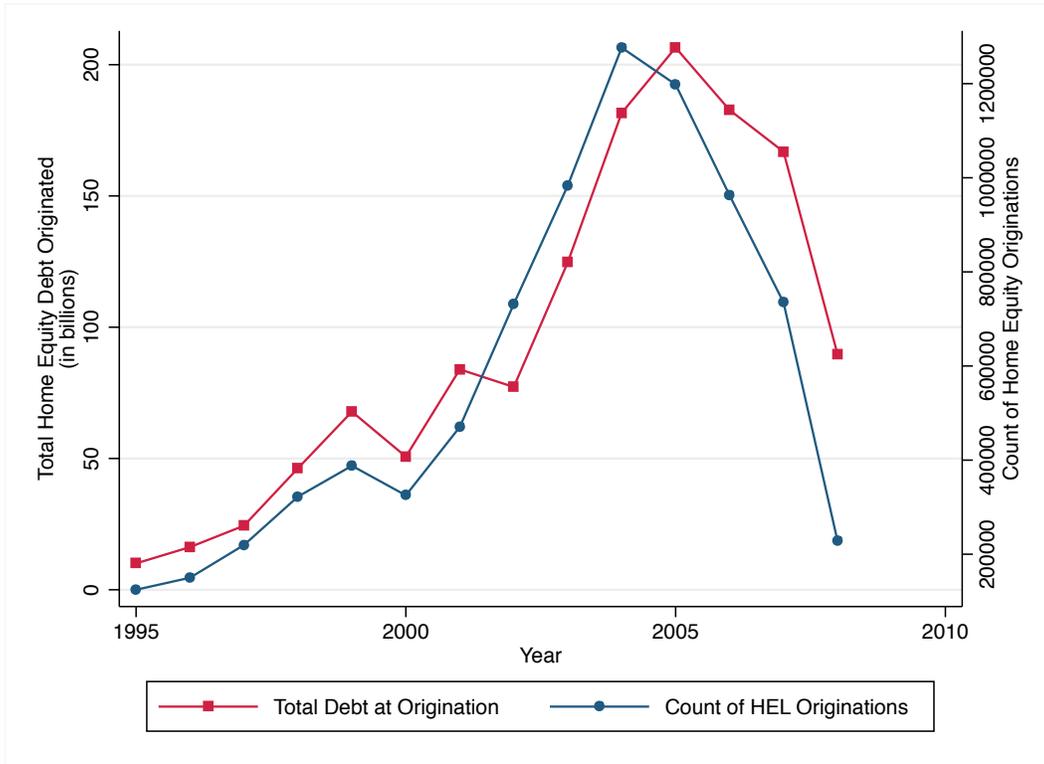
In this paper, I use a policy kink and a policy notch to estimate how home equity debt responds to two mortgage policies: the home mortgage interest deduction and the licensed appraisal requirements of the FIRREA. Using administrative data on home loan originations, I find evidence of substantial bunching at the limit of mortgage interest deductibility at \$100,000, and at the threshold of \$250,000 beyond which home equity loans require a licensed appraisal of the associated property. The corresponding estimates suggest a substantial reduction in debt in response to these policies: removing the mortgage interest deduction reduces loan size by about 20% at origination (for those who take loans in the

neighborhood of \$100,000), while requiring licensed appraisals reduces loan size by about 22% (for those in the neighborhood of \$250,000). Interestingly, I find that response to the tighter appraisal requirement decreases during the lending boom of the 2000s, suggesting that licensed appraisals were less effective during that period.

Relating my estimated response of debt to the removal of the mortgage interest deduction, gives a large elasticity of debt to interest rates. I use data from the Survey of Consumer Finances (and NBER TAXSIM) to impute the change in the after-tax interest rate that a typical home-equity borrower would experience. The implied elasticity ranges from -28 to -40, which is substantially larger than estimates the few existing quasi-experimental studies, and suggests that there may be substantial welfare costs to the taxation or subsidization of savings. However, this interpretation relies on the assumption that lenders are not responding to the change in policy, which I cannot explicitly verify. Nonetheless, my estimates suggest a large market-level response of debt to the mortgage interest deduction and appraisal requirements.

Figure 3.1: Data Over time

(a) Volume of Home Equity Loans over Time



(b) Median Value of Home-Equity Loans Over Time

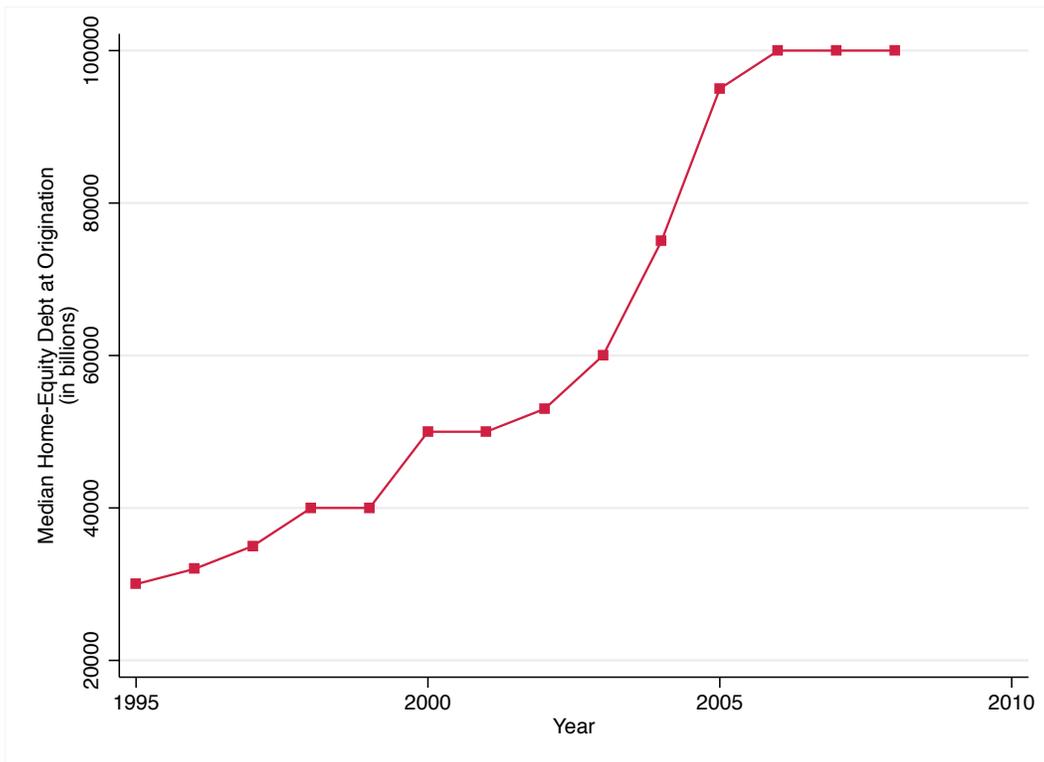
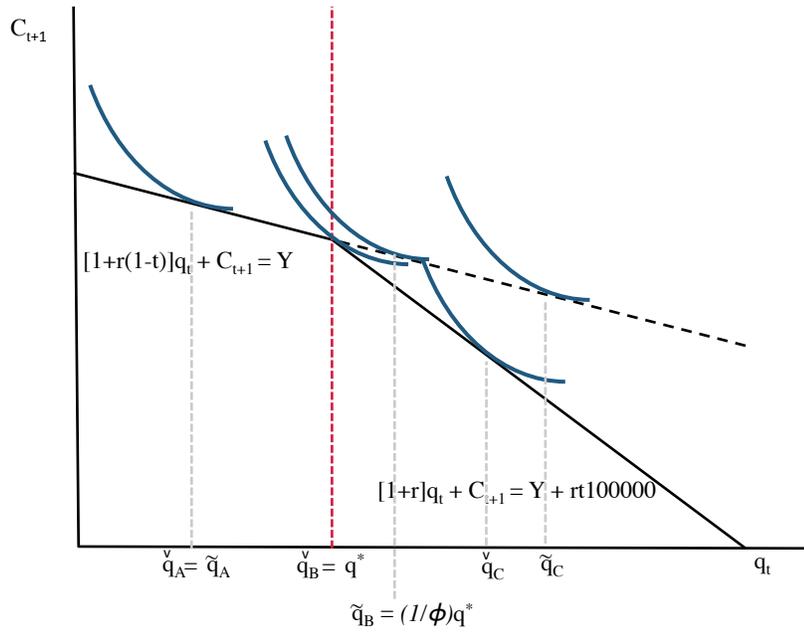


Figure 3.2: Conceptual Kink Figures

(a) Kinked Budget Set



(b) Kinked Distribution

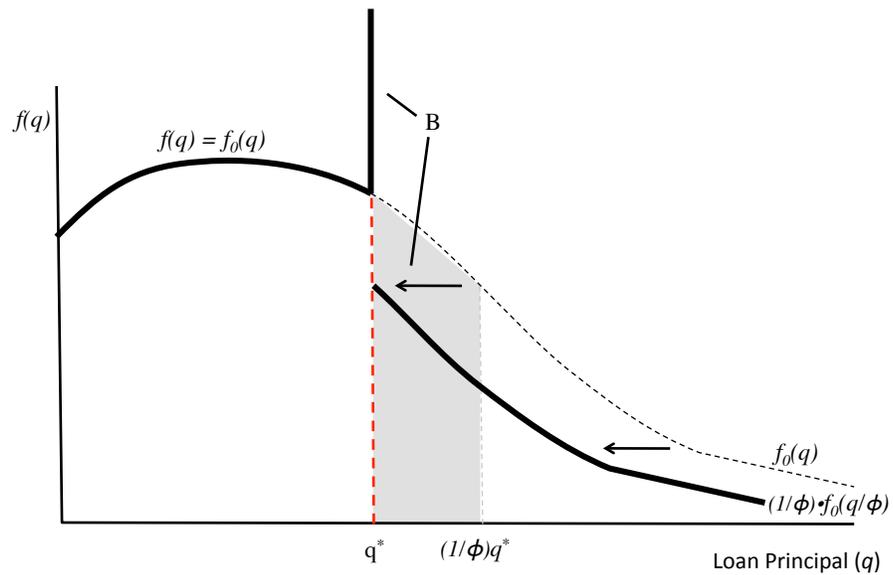
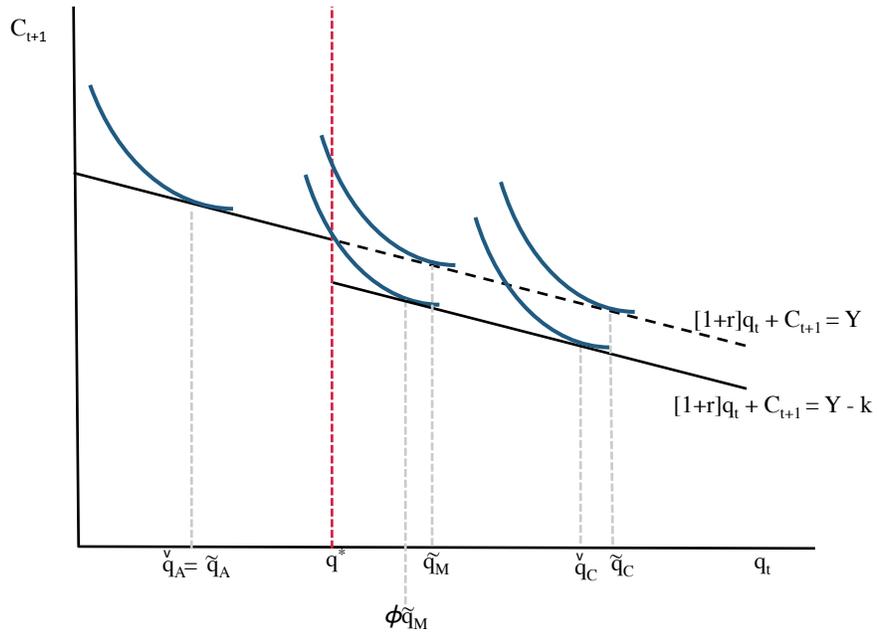


Figure 3.3: Conceptual Notch Figures

(a) Notched Budget Set



(b) Notched Distribution

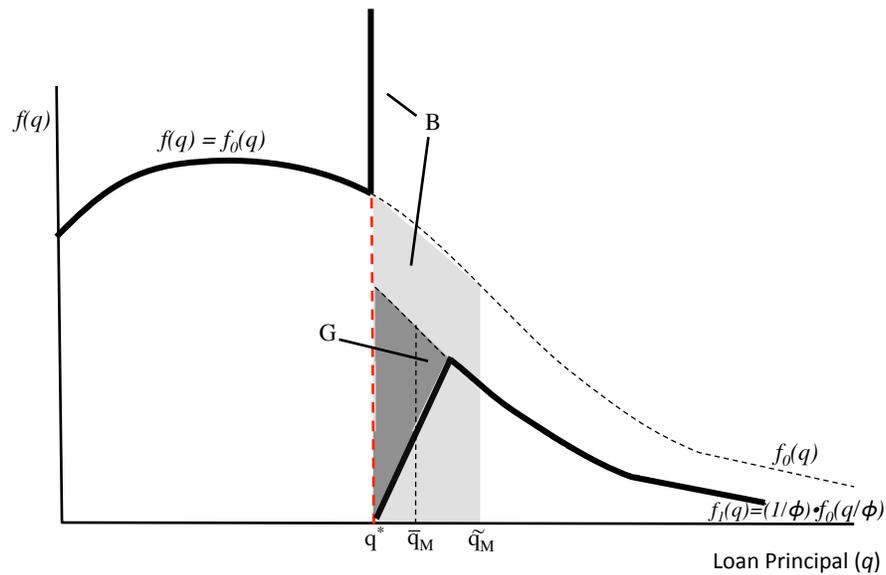
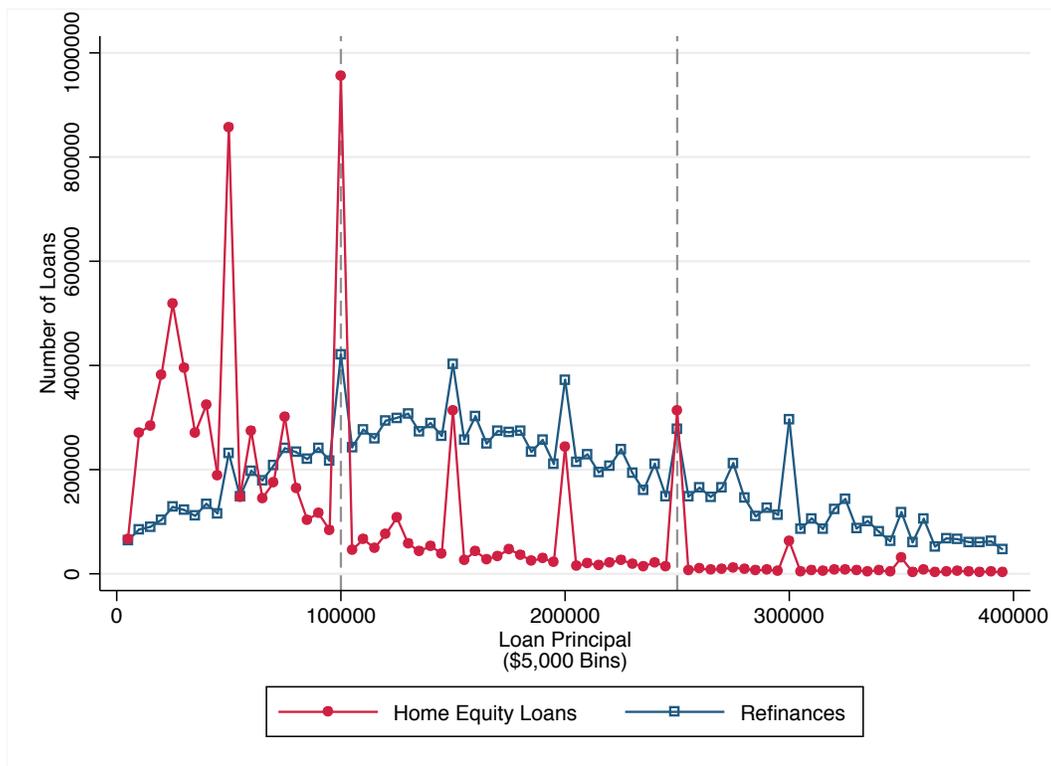


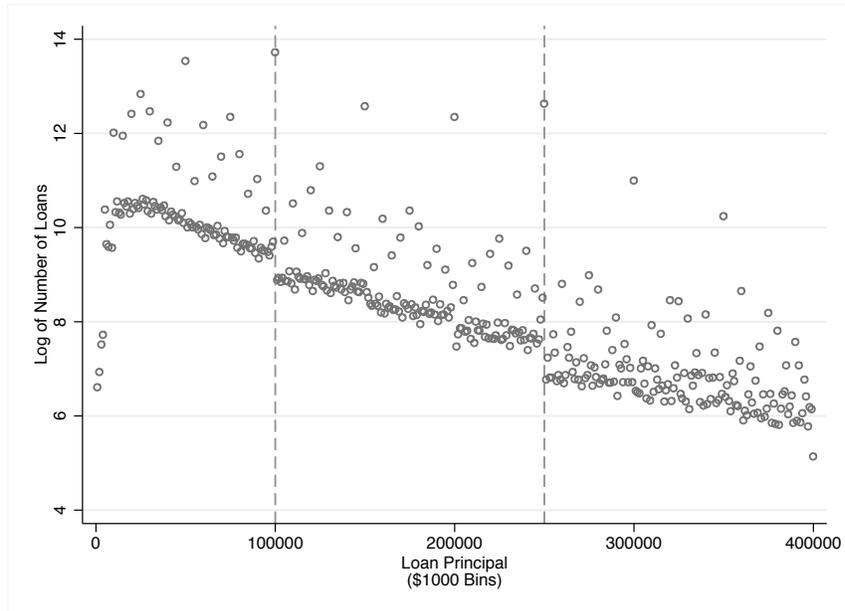
Figure 3.4: Histogram of Home Loans



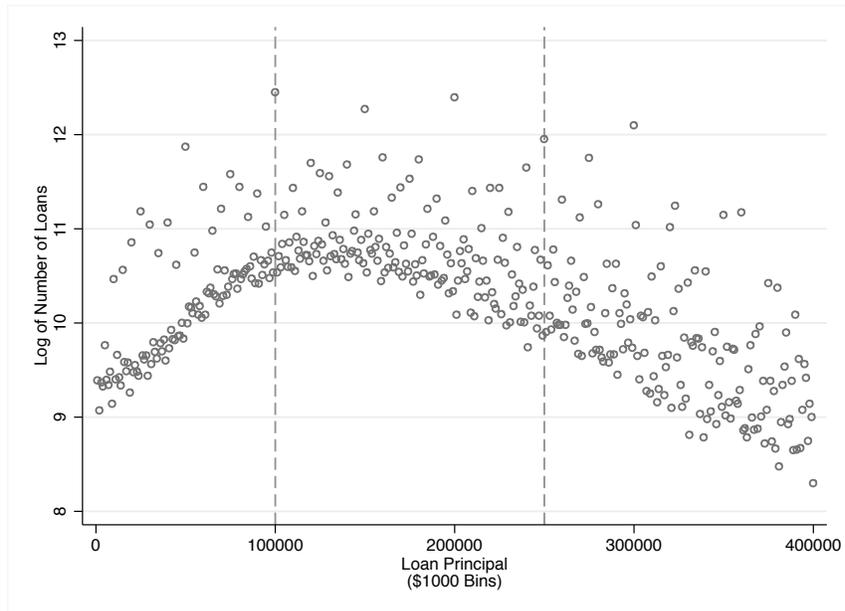
Notes: Histogram (\$5k bins) of home equity loan and mortgage refinance originations for CA, IL, NJ, and NY (1995–2008). Dashed lines represent the mortgage interest deduction (\$100k) and FIRREA appraisal (\$250k) thresholds.

Figure 3.5: Log of Histogram of Home Loans

(a) Log histogram of Home Equity Loans

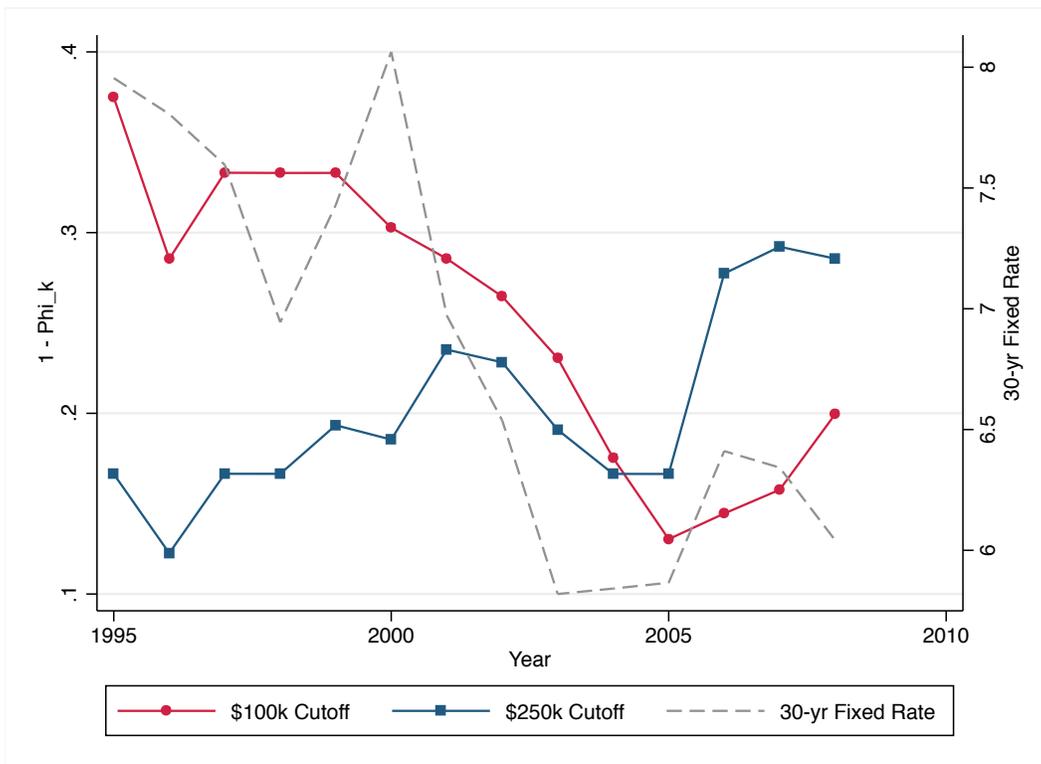


(b) Log Histogram of Refinances



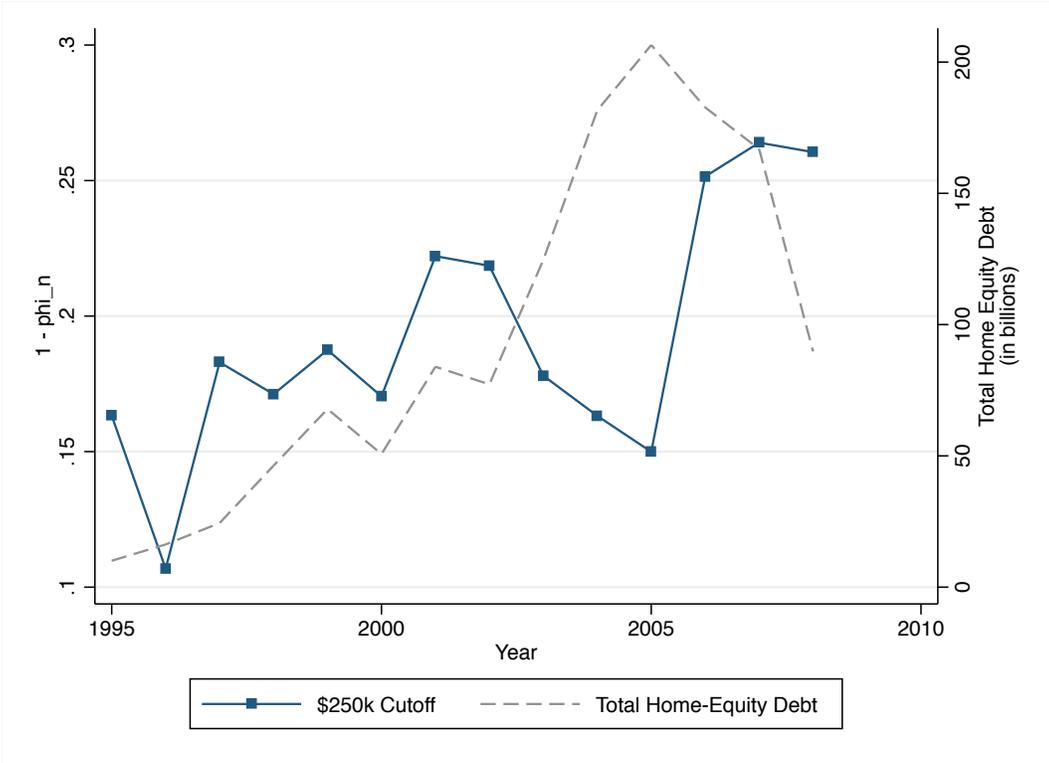
Notes: Histogram (\$1000 bins) of home equity loan (upper panel) and mortgage refinance (lower panel) originations for CA, IL, NJ, and NY (1995–2008). Dashed lines represent the mortgage interest deduction (\$100k) and FIRREA appraisal (\$250k) thresholds.

Figure 3.6: Kink Estimates Over Time



Notes: Baseline estimates of $1 - \phi_k$ over time for the \$100k and \$250k thresholds plotted against the 30-year fixed rate mortgage interest rate.

Figure 3.7: \$250k Notch Estimates Over Time



Notes: Baseline estimates of $1 - \phi_n$ over time for the \$250k thresholds plotted against the annual total of home equity debt originations (in billions).

Table 3.1: Home Loan Statistics (1995–2008)

	Home Equity Loans	Mortgage Refinances
Number of Loans	8,405,850	16,606,227
Loans ∈ 90k–110k	135,134	546,085
Loans ∈ 240k–260k	58,527	705,495
Mean Loan Principal (Nominal)	174,985	436,130
Mean Loan Principal (2000 Dollars)	157,287	386,240
Median Loan Principal (2000 Dollars)	62,288	173,207

Notes: “Loans ∈ 90,000–110,000” presents the number of the given loan type with (nominal-valued) principal at origination between \$90,000 and \$110,000.

Table 3.2: Baseline Estimates and Sensitivity

	$(1 - \widehat{\phi}_k)$		$(1 - \widehat{\phi}_n)$	
	\$100,000	\$250,000	\$250,000	n
Baseline	0.1996000 (0.0000167)	0.2353085 (0.0030037)	0.2201608 (0.0030477)	4201901
1st Order	0.1661398 (0.0000211)	0.2423946 (0.0020500)	0.2123417 (0.0021275)	4201901
Polynomial				
3rd Order	0.1998311 (0.0000173)	0.1666250 (0.0000046)	0.1546645 (0.0011514)	4201901
4th Order	0.1998019 (0.0000177)	0.1666340 (0.0000048)	0.1522683 (0.0012373)	4201901
5th Order	0.1663140 (0.0000455)	0.2752993 (0.0061452)	0.2552296 (0.0059724)	4485418
6th Order	0.1660396 (0.000864)	0.2259143 (0.0049487)	0.2106439 (0.0049402)	3797876
No Interaction with Round #s	0.1994315 (0.0000157)	0.2707455 (0.0030040)	0.2184322 (0.0031018)	4091119
2nd Order	0.1664815 (0.0000323)	0.1666452 (0.0000036)	0.1651159 (0.0016347)	3602382
Omit \$10k Below \$20k Above	0.1689187 (0.0020895)	0.2039459 (0.0023308)	0.1953821 (0.0023596)	4485418
\$35k Above	0.1994503 (0.0000152)	0.2282727 (0.0025010)	0.2160060 (0.0024762)	4260866
\$50k Above	0.1995467 (0.0000165)	0.2307022 (0.0009471)	0.2155439 (0.0010025)	3797876
Omit \$40k Above and \$5k Below	0.2288825 (0.0019185)	0.2519205 (0.0026387)	0.2371248 (0.0026761)	4331271
\$15k Below	0.1994722 (0.0000188)	0.2306717 (0.0008504)	0.2154378 (0.0008563)	4091119
\$25k Below	0.2258464 (0.0030612)	0.1934624 (0.0015077)	0.1775133 (0.0015316)	3602382
Discontinuities at \$100k, \$150k, \$200k, and \$250k only	0.2306726 (0.0005878)	0.1495348 (0.0020271)	0.1663582 (0.0021039)	4201901

Notes: Baseline specification omits data in the ranges \$90,000–\$140,000 and \$240,000–\$290,000, and estimates counterfactual by poisson regression of the count of loans per \$100 bin on a 2nd-order polynomial in loan principal, allowing for discontinuity at \$100k and \$250k as well as a fixed discontinuity at multiples of \$50k, and fixed effects for multiples of \$500, \$1000, \$5000, \$25000, and \$50000 plus linear interaction with loan principal. Bootstrap standard errors with 999 repetitions. As described in the text, kink estimates are found by taking excess bunching at kink (difference between mass of observed loans and loans predicted by counterfactual) and integrating under the no-kink counterfactual above the threshold to the point where integrated mass equals bunching ($1 - \widehat{\phi}_k$ is the percent reduction from this bunching point to the threshold). Notch estimates are found by finding the “bunching” point as with kink estimation, finding the missing mass from the gap by comparing the counterfactual above the notch to the observed mass of loans above the notch, integrating under the counterfactual above the notch to find the point where this integrated mass equals the missing mass, and comparing this gap point to the bunching point ($1 - \widehat{\phi}_n$ is the percent reduction from the bunching point to the gap point).

Table 3.3: Round Number Placebos

Cutoff	$1 - \phi_k$		$1 - \phi_n$		n
	Home Equity	Refinance	Home Equity	Refinance	
100000	0.1996000 (0.0000167)	0.0475814 (0.0000357)	0.1198919 (0.0000171)	-0.0087601 (0.000141)	4201901
250000	0.2353085 (0.0030037)	0.0194803 (0.0000108)	0.2201608 (0.0030477)	0.0350939 (0.0000999)	4201901
150000	-0.1109449 (0.0000379)	-0.0267686 (0.000068)	-0.1843628 (0.0000403)	-0.021499 (0.0005818)	5531743
200000	0.0697202 (0.0011408)	0.0025868 (0.0002833)	-0.0043468 (0.0012543)	-0.0070201 (0.0002845)	5743302
300000	-0.1995379 (0.0000011)	-0.0132077 (0.0000418)	-0.0398446 (0.0002592)	0.020402 (0.0000404)	6048969
350000	-0.1663163 (0.0000009)	-0.0203837 (0.0000621)	-0.0333204 (0.0011186)	-0.0289084 (0.0000898)	6091206

Notes: Estimates as in baseline specification from Table 3.2, except treating the given cutoff as the kink/notch (i.e., omitting data round that cutoff rather than \$100k or \$250k).

Table 3.4: Estimates by Quartile of AGI

Quartile of AGI	$\widehat{1 - \phi_k}$		$\widehat{1 - \phi_n}$	n
	\$100,000	\$250,000	\$250,000	
1	0.1521923 (0.006671)	0.1665482 (0.0027971)	0.1555081 (0.0027699)	192099
2	0.162068 (0.0027061)	0.1933219 (0.0038922)	0.1791202 (0.0045289)	368878
3	0.1663083 (0.0049238)	0.1709739 (0.0035651)	0.1546283 (0.0043205)	831104
4	0.1993969 (0.0000995)	0.1721729 (0.0035921)	0.1523333 (0.0046358)	1488503

Notes: Estimates as in baseline specification from Table 3.2, except restricting sample to zip codes in the given quartile of adjusted gross income based on the IRS SOI (1998, 2001, 2002, 2004–2008). Quartiles are calculated within the given year.

Table 3.5: State-Year Regressions of $1 - \phi_k$ on Percent-Change in After-Tax Interest Rate

	\$100k Threshold			\$250k Threshold		
Coefficient on Percent- Change in $(1 + (1 - t)r)$	0.844*** (0.060)	1.634*** (0.106)	7.705*** (1.003)	0.120* (0.064)	-0.706*** (0.101)	-0.463 (1.131)
State Fixed Effects		X	X	X	X	X
Year Fixed Effects			X			X

Notes: Regression of state/year estimates of $1 - \phi_k$ on average imputed state/year change in after-tax interest rates at the MID kink. Estimates of $1 - \phi_k$ using the baseline specification from Table 3.2, except restricting sample to each state in each year. Imputed state/year change in rates from the SCF—marginal tax rates for all households with home equity loans (pooling 1998, 2001, 2004, 2007 survey years) imputed from NBER Taxsim for each state and year (i.e., assume all households live in given state and given year); yearly interest rates for the 30-year fixed rate mortgages as reported in the Freddie Mac mortgage rate survey. Standard errors are found by bootstrapping over the entire procedure (resampling SCF data and home loan data) 999 times.

Table 3.6: Implied Elasticity of Debt to Interest Rates at \$100,000

	All Home-Equity Loan Households	Loans Between \$75,000–\$125,000
Mean Interest Rate	0.087	0.076
Mean Marginal Tax Rate	0.229	0.197
Percent Change in After-Tax Interest Rate	0.019	0.016
Implied Elasticity	-10.505	-12.475
Share Home Improvement	0.525	0.605
Adjusted Rate Change	0.007	0.005
Adjusted Elasticity	-28.514	-39.920
SCF Sample Size	4,144	397

Notes: Average interest rate and rate of use of home equity loan for construction from Survey of Consumer Finances (1998, 2001, 2004, 2007); marginal tax rate imputed for SCF sample using NBER TAXSIM. “All” denotes all SCF households with a home equity loan, \$75k–\$125k denotes all SCF households with a home equity loan in the given range. Marginal tax rate is calculated from survey data (household income, number of dependents, survey year) using NBER TAXSIM. “Home Improvement” indicates that the home equity loan was used for home repairs—change in rates for these households is set to zero to estimate the average adjusted rate change. Percent change in quantity is estimated as in the baseline case of Table 3.2.

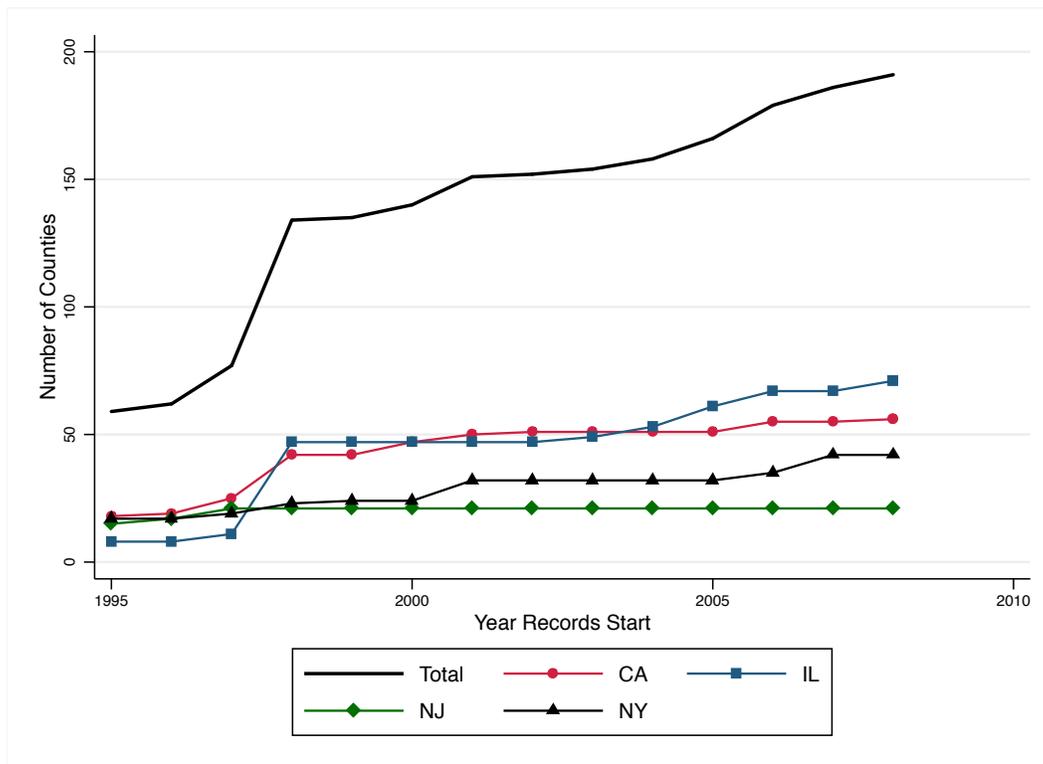
3.7 Data Appendix

My main data set is a collection of home equity loan and mortgage refinance origination records. These records have been collected by an anonymous private firm and made available to me by the Paul Milstein Center for Real Estate at the Columbia Graduate School of Business. The data indicate the size of the loan at origination, the date the loan is made, and the address of the associated property (including zip code). I drop all transactions with missing address or loan principal. I drop county-years in which fewer than 20% of records have valid addresses and counties that are not observed through 2008 (suggesting irregularities in data collection). I also drop duplicate records—multiple loans with identical principal that occur at the same property within 30 days of one another. I keep only loans on residential property.

One concern examining the data over time is that the county-year panels are not balanced—data for some counties is available as of 1995, while others are not available until the mid-2000s. Figure 3.8 presents the cumulative distribution of start dates for the counties in the sample (i.e., one observation per county). As of 1995, about 35% of counties are in the sample, and by 2000 more than three quarters of counties have data available. To explore county selection into the dataset I consider three distinct samples: the overall unbalanced data, home-equity loans from counties with data available as of 1995, and home-equity loans from counties with data available as of 2000. Panel a of Figure 3.1 presents the total number of loan originations by month. This figure demonstrates the dramatic increase in lending over this period—the number of loans increases over time with a slight dip through the 2000/2001 recession followed by a large run-up peaking in 2005. The two balanced samples track the unbalanced sample quite closely. Panel c presents the median loan principal (in year-2000 dollars) by year. Again, there is minimal difference between the total sample and the two balanced samples. Loan principal appears to increase steadily over time—the boom of the early 2000s is not as obvious here, suggesting that the credit expansion operated primarily through an increase in the number of home loans (as in panels a and b) and not the

magnitude of loans.

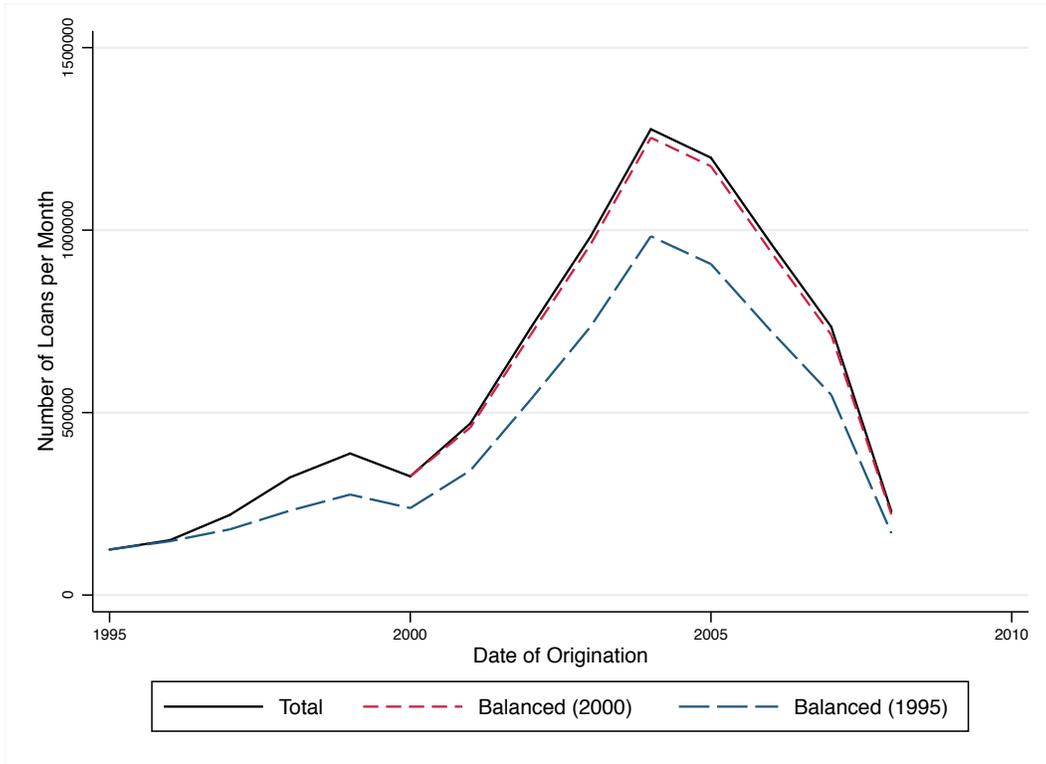
Figure 3.8: County Start Dates



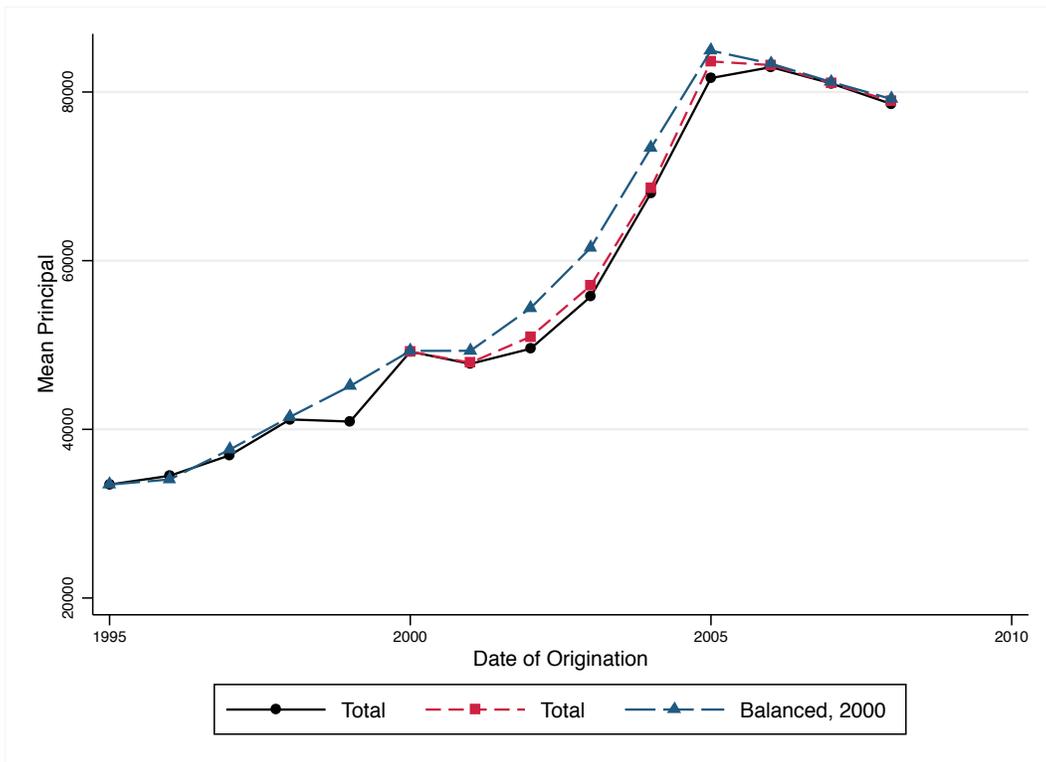
Notes: Historical availability of data varies by county. This figure presents a cumulative count of county start years by state. Each point represents the number of counties in the given state with data available as of the given year.

Figure 3.9: Loans over Time by Sample

(a) Number of Home Equity Loans Loans Over Time



(b) Mean Value of Home-Equity Loans Over Time



References

- Abdallah, C. S. and W. D. Lastrapes (2012). Home equity lending and retail spending: Evidence from a natural experiment in Texas. *American Economic Journal: Macroeconomics* 4(4), 94–125.
- Abrams, D., M. Bertrand, and S. Mullainathan (2008). Do Judges Vary in Their Treatment of Race? *American Law and Economics Association Annual Meetings* 93.
- Adelino, M., K. Gerardi, and P. Willen (2009). Why don't lenders renegotiate more home mortgages? redefaults, self-cures and securitization. *NBER Working Paper* (15159).
- Agarwal, S., G. Amromin, I. Ben-David, S. Chomsisengphet, and D. D. Evanoff (2011, December). The role of securitization in mortgage renegotiation. *Journal of Financial Economics* 102(3), 559–578.
- Anderson, J., J. Kling, and K. Stith (1999). Measuring interjudge sentencing disparity: Before and after the federal sentencing guidelines. *Journal of Law and Economics* 42(February), 271–307.
- Anenberg, E. and E. Kung (2014). Estimates of the size and source of price declines due to nearby foreclosures. *American Economic Review*.
- Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Attanasio, O. P., P. Koujianou Goldberg, and E. Kyriazidou (2008). Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review* 49(2), 401–436.
- Autor, D. and S. Houseman (2010). Do temporary-help jobs improve labor market outcomes for low-skilled workers? Evidence from “Work First”. *American Economic Journal: Applied Economics* 2(July), 96–128.
- Bastani, S. and S. Hakan (2014). Bunching and non-bunching at kink points of the Swedish tax schedule. *Journal of Public Economics* 109, 36–49.
- Benjamin, J. D., N. E. Coulson, and S. X. Yang (1993). Real estate transfer taxes and property values: The Philadelphia story. *The Journal of Real Estate Finance and Economics* 7(2), 151–157.
- Berdejo, C. and D. L. Chen (2010). Priming Ideology: Electoral Cycles Among Unelected Judges. *Working Paper* (January).
- Bernheim, B. D. (2002). Taxation and saving. In A. Auerbach and M. Feldstein (Eds.), *Handbook of Public Economics* (1 ed.), Volume 3, pp. 1173–1249. Elsevier.
- Bernheim, D. B. and J. Meer (2013). Do Real Estate Brokers Add Value When Listing Services Are Unbundled? *Economic Inquiry* 51(2).

- Besley, T., N. Meads, and P. Surico (2013). The incidence of a transactions tax: Evidence from a stamp duty holiday. London School of Economics, mimeo.
- Best, M. and H. Kleven (2013, April). Property transaction taxes and the housing market: Evidence from notches and housing stimulus in the uk. London School of Economics, mimeo.
- Bourassa, S. C. and W. G. Grigsby (2000). Income tax concessions for owner-occupied housing. *Housing Policy Debate* 11(3), 521–546.
- Bourassa, S. C. and M. Yin (2008). Tax deductions, tax credits and the homeownership rate of young urban adults in the united states. *Urban Studies* 45(5), 1141–1161.
- Brito, D. L. and P. R. Hartley (1995). Consumer rationality and credit cards. *Journal of Political Economy*, 400–433.
- Bubb, R. and A. Kaufman (2009). Securitization and moral hazard: Evidence from credit score cutoff rules.
- Cameron, C., J. B. Gelbach, and D. L. Miller (2011, April). Robust Inference With Multiway Clustering. *Journal of Business & Economic Statistics* 29(2), 238–249.
- Campbell, J. and J. Cocco (2011). A Model of Mortgage Default. *NBER Working Paper* (17516).
- Campbell, J., S. Giglio, and P. Pathak (2011). Forced sales and house prices. *American Economic Review* 101(5), 2108–31.
- Campbell, J. Y. (2013). Mortgage market design. *Review of Finance* 17(1), 1–33.
- Carrillo, P. E. (2012). An empirical stationary equilibrium search model of the housing market. *International Economic Review* 53(1), 203–234.
- Carrillo, P. E. and J. C. Pope (2012). Are Homes Hot or Cold Potatoes? The Distribution of Marketing Time in the Housing Market. *Regional Science and Urban Economics* 42(1), 189–197.
- Chang, T. and A. Schoar (2006). The effect of judicial bias in Chapter 11 reorganization. *Unpublished manuscript. Massachusetts Institute of Technology.*
- Chetty, R., J. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment costs, firm responses, and labor supply elasticities: Evidence from danish tax records. *Quarterly Journal of Economics* 126(2), 749–804.
- Chetty, R., J. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *American Economic Review* 103(7), 2683–2721.
- Currie, J. and J. Gruber (1996). Saving babies: The efficacy and cost of recent changes in the medicaid eligibility of pregnant women. *Journal of Political Economy*, 1263–1296.

- Dachis, B., G. Duranton, and M. A. Turner (2012). The Effects of Land Transfer Taxes on Real Estate Markets: Evidence from a Natural Experiment in Toronto. *Journal of Economic Geography* 12(2), 327–354.
- Dehejia, R., H. Montgomery, and J. Morduch (2012). Do interest rates matter? credit demand in the dhaka slums. *Journal of Development Economics* 97(2), 437–449.
- Dell’Ariccia, G., D. Igan, and L. Laeven (2012). Credit booms and lending standards: Evidence from the subprime mortgage market. *Journal of Money, Credit and Banking* 44(3), 367–384.
- Deng, Y., J. Quigley, and R. Order (2000). Mortgage terminations, heterogeneity and the exercise of mortgage options. *Econometrica* 68(2), 275–307.
- Ding, L., R. G. Quercia, and A. M. White (2009, January). State Anti-Predatory Lending laws: Impact and Federal Preemption Phase I Descriptive Analysis.
- Dobbie, W. and J. Song (2013). Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection.
- Doyle, J. (2007). Child protection and child outcomes: Measuring the effects of foster care. *The American Economic Review* 97(5), 1583–1610.
- Doyle Jr., J. J. and K. Samphantharak (2008). \$2.00 Gas! Studying the Effects of a Gas Tax Moratorium. *Journal of Public Economics* 92(3-4), 869–884.
- Duflo, E., W. Gale, J. Liebman, P. Orszag, and E. Saez (2006). Saving Incentives for Low- and Middle-Income Families: Evidence from a Field Experiment with H&R Block. *Quarterly Journal of Economics* 121(4), 1311–1346.
- Dunsky, R. M. and J. R. Follain (2000). Tax-induced portfolio reshuffling: The case of the mortgage interest deduction. *Real Estate Economics* 28(4), 683–718.
- Elul, R. (2011). Securitization and mortgage default.
- Feenberg, D. and E. Coutts (1993). An introduction to the taxsim model. *Journal of Policy Analysis and Management* 12(1), 189–194.
- Fjærli, E. (2004). Tax reform and the demand for debt. *International Tax and Public Finance* 11(4), 435–467.
- Foote, C., K. Gerardi, L. Goette, and P. Willen (2009). Reducing foreclosures: No easy answers. Technical report, National Bureau of Economic Research.
- Foote, C. L., K. Gerardi, and P. S. Willen (2008, September). Negative equity and foreclosure: Theory and evidence. *Journal of Urban Economics* 64(2), 234–245.
- Frame, S., A. Lenhert, and N. Prescott (2008). A snapshot of mortgage conditions with an emphasis on subprime mortgage performance.

- Frame, W. (2010). Estimating the effect of mortgage foreclosures on nearby property values: A critical review of the literature. *Economic Review, Federal Reserve Bank of Atlanta* (3).
- Gelpern, A. and A. Levitin (2009). Rewriting frankenstein contracts: The workout prohibition in residential mortgage-backed securities. *Southern California Law Review* 82, 1077–1152.
- Genesove, D. and L. Han (2012). Search and matching in the market for existing homes. *Journal of Urban Economics* 72, 31–45.
- Genesove, D. and C. Mayer (2001). Loss Aversion and Seller behavior: Evidence from the Housing Market. *Quarterly Journal of Economics* 116(4), 1233–1260.
- Gerardi, K., E. Rosenblatt, P. Willen, and V. W. Yao (2012). Foreclosure externalities: Some new evidence. *Federal Reserve Bank of Atlanta* (Working Paper 12).
- Ghent, A. C. and M. Kudlyak (2011). Recourse and residential mortgage default: Evidence from us states. *Review of Financial Studies* 24(9), 3139–3186.
- Glaeser, E. and B. Sacerdote (2000). The Social Consequences of Housing. *Journal of Housing Economics* 3, 1–23.
- Glaeser, E. and J. Shapiro (2003). The benefits of the home mortgage interest deduction. In *Tax Policy and the Economy*, Volume 17, pp. 37–82. National Bureau of Economic Research.
- Goodstein, R., P. E. Hanouna, C. D. Ramirez, and C. W. Stahel (2011). Are Foreclosures Contagious? *SSRN Electronic Journal*, 1–34.
- Goodstein, R. and Y. Lee (2010). Do Foreclosures Increase Crime? *Available at SSRN 1670842*.
- Green, R. K. and K. D. Vandell (1999). Giving households credit: How changes in the u.s. tax code could promote homeownership. *Regional Science and Urban Economics* 29(4), 419–444.
- Gross, D. and N. Souleles (2002a). An empirical analysis of personal bankruptcy and delinquency. *Review of Financial Studies* 15(1), 319–347.
- Gross, D. B. and N. B. Souleles (2002b). Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. *The Quarterly Journal of Economics* 117(1), 149–185.
- Guiso, L., P. Sapienza, and L. Zingales (2013). The Determinants of Attitudes toward Strategic Default on Mortgages. *The Journal of Finance* 68(4), 1473–1515.
- Guren, A. and T. McQuade (2013, April). How do foreclosures exacerbate housing downturns?

- Gyourko, J. and T. Sinai (2004). The (un)changing geographical distribution of housing tax benefits: 1980 to 2000. *Tax Policy and the Economy* 88(7-8), 175–208.
- Han, L. and W. Strange (2012). What is the role of the asking price for a house? University of Toronto, mimeo.
- Han, L. and W. Strange (2013). Bidding wars for houses. *Real Estate Economics*.
- Hanson, A. (2012). Size of home, homeownership, and the mortgage interest deduction. *Journal of Housing Economics*.
- Harding, J. P. J., E. Rosenblatt, and V. W. V. Yao (2009, November). The contagion effect of foreclosed properties. *Journal of Urban Economics* 66(3), 164–178.
- Hastings, J. and E. Washington (2010). The First of the Month Effect: Consumer Behavior and Store Responses. 2(2), 142–162.
- Haurin, D. R., J. L. Haurin, T. Nadauld, and A. Sanders (2010). List Prices, Sale Prices and Marketing Time: An Application to US Housing Markets. *Real Estate Economics* 38(4), 659–685.
- Hendershott, P. and G. Pryce (2006). The sensitivity of homeowner leverage to the deductibility of home mortgage interest. *Journal of Urban Economics* 60(1), 50–68.
- Hilber, C. and T. M. Turner (2012). The mortgage interest deduction and its impact on homeownership decisions. *mimeo*.
- Iaryczower, M. (2009). The Value of Information in the Court—Get it Right, Keep it Tight. *American Economic Review* 102(1), 202–237.
- Immergluck, D. and G. Smith (2006). The external costs of foreclosure: The impact of single-family mortgage foreclosures on property values. *Housing Policy Debate* 17(1), 57–80.
- Internal Revenue Service (2011, December). *Publication 936: Home Mortgage Interest Deduction*. Internal Revenue Service.
- Jappelli, T. and L. Pistaferri (2007). Do people respond to tax incentives? an analysis of the italian reform of the deductibility of home mortgage interests. *European Economic Review* 51(2), 247–271.
- Jia, P. and P. A. Pathak (2010). The Impact of Commissions on Home Sales in Greater Boston. *American Economic Review: Papers and Proceedings*.
- Jiang, W., A. Nelson, and E. Vytlačil (2010). Securitization and loan performance: Contrast of ex ante and ex post relations in the mortgage market.
- Jiang, W., A. Nelson, and E. Vytlačil (2014). Liar’s loan? - effects of origination channel and information falsification on mortgage loan delinquency. *Review of Economics and Statistics*.

- Johnson, R. C. (2012). The impact of parental wealth on college enrollment and degree attainment: Evidence from the housing boom and bust.
- Karlan, D. S. and J. Zinman (2008). Credit elasticities in less-developed economies: Implications for microfinance. *American Economic Review* 98(3), 1040–1068.
- Keys, B., T. Mukherjee, A. Seru, and V. Vig (2009). Financial regulation and securitization: Evidence from subprime mortgage loans. *Journal of Monetary Economics* 56(5).
- Keys, B., T. Mukherjee, A. Seru, and V. Vig (2010). Did securitization lead to lax screening? evidence from subprime loans. *Quarterly Journal of Economics* 125(1).
- Keys, B., A. Seru, and V. Vig (2012). Lender screening and the role of securitization: Evidence from prime and subprime mortgage markets. *Review of Financial Studies* 25(7), 2071–2108.
- Kiel, P. (2012, April). The american foreclosure story: The struggle for justice and a place to call home. *ProPublica*.
- Kleven, H. J. and M. Waseem (2013, May). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan. *Quarterly Journal of Economics* 128(2), 669–723.
- Kling, J. (2006). Incarceration length, employment, and earnings. *American Economic Review* 96(3), 863–876.
- Kopczuk, W. and D. Munroe (2014). Mansion tax: The effect of transfer taxes on residential real estate market.
- Leonard, T. and J. C. Murdoch (2009, May). The neighborhood effects of foreclosure. *Journal of Geographical Systems* 11(4), 317–332.
- Levitin, A. (2009). Helping homeowners: Modification of mortgages in bankruptcy. *Harvard Law & Policy Review Online* 3.
- Levitin, A. J. and T. Twomey (2011). Mortgage Servicing. *Yale J. on Reg.* (11).
- Levitt, S. D. and C. Syverson (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics* 90(4), 599–611.
- Lin, Z., E. Rosenblatt, and V. W. Yao (2007, November). Spillover Effects of Foreclosures on Neighborhood Property Values. *The Journal of Real Estate Finance and Economics* 38(4), 387–407.
- Lovenheim, M. (2013). The effect of housing wealth on college choice: Evidence from the housing boom. *Journal of Human Resources* 14(1), 3–37.
- Madrian, B. C. and D. F. Shea (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Quarterly Journal of Economics* 116(4), p1149–1187.

- Maestas, N., K. Mullen, and A. Strand (2013). Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt. *American Economic Review* 103(5), 1797–1829.
- Maki, D. M. (2001). Household debt and the tax reform act of 1986. *The American Economic Review* 91(1), 305–319.
- Marion, J. and E. Muehlegger (2011, October). Fuel Tax Incidence and Supply Conditions. *Journal of Public Economics* 95(9-10), 1202–12.
- Marx, B. M. (2013). Regulatory hurdles and growth of charitable organizations: Evidence from a dynamic bunching design. Columbia University, mimeo.
- Mayer, C., E. Morrison, T. Piskorski, and A. Gupta (2011). Mortgage Modification and Strategic Default: Evidence from a Legal Settlement with Countrywide. *NBER Working Paper*.
- Mayer, C., K. Pence, and S. M. Sherlund (2009). The rise in mortgage defaults. *Journal of Economic Perspectives* 23(1), 27–50.
- Merlo, A. and F. Ortalo-Magne (2004). Bargaining over residential real estate: Evidence from England. *Journal of Urban Economics* 56, 192–216.
- Mian, A. and A. Sufi (2009). The consequences of mortgage credit expansion: Evidence from the u.s. mortgage default crisis. *Quarterly Journal of Economics* 124(4), 1449–1496.
- Mian, A., A. Sufi, and F. Trebbi (2011). Foreclosures, house prices, and the real economy. *NBER Working Papers* (16685).
- Mian, A. R., A. Sufi, and F. Trebbi (2012). Foreclosures, House Prices, and the Real Economy. *Chicago Booth Research Paper* (13-41).
- Mishra, P., A. Subramanian, and P. Topalova (2008). Policies, Enforcement, and Customs Evasion: Evidence from India. *Journal of Public Economics* 92(10-11), 1907–1925.
- Mueller-Smith, M. (2013). Program evaluation with randomized screeners: Estimating heterogeneous response instrumental variable (hriv) models.
- Office of Management and Budget (2012, February). *Analytical Perspectives Budget of the United States Government Fiscal Year 2012*. Office of Management and Budget.
- Pence, K. (2006). Foreclosing on opportunity: State laws and mortgage credit. *The Review of Economics and Statistics* 88(1), 177–182.
- Pennington-Cross, A. (2006). The value of foreclosed property. *Journal of Real Estate Research* 28(2), 193–214.
- Persson, P. (2013, January). Social insurance and the marriage market. Columbia University, mimeo.

- Poterba, J. and T. Sinai (2011). Revenue costs and incentive effects of the mortgage interest deduction for owner-occupied housing. *National Tax Journal* 64(2), 531–564.
- Poterba, J. M., S. F. Venti, and D. A. Wise (1996). How Retirement Saving Programs Increase Saving. *Journal of Economic Perspectives* 10(4), 91–112.
- Ramnath, S. (2012). Taxpayers’ response to notches: Evidence from the saver’s credit. *Journal of Public Economics* 101, 77–93.
- Rosen, H. (1979). Housing decisions and the u.s. income tax. *Journal of Public Economics* 11, 1–23.
- Rothstein, J. (2010). Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence. *American Economic Journal: Economic Policy* 2(1), 177–208.
- Saarimaa, T. (2010). Tax incentives and demand for mortgage debt: evidence from the finnish 1993 tax reform. *International Journal of Housing Policy* 10(1), 19–40.
- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy* 2(3), 180–212.
- Saez, E., M. Matsaganis, and P. Tsakloglou (2011, February). Earnings Determination and Taxes: Evidence from a Cohort Based Payroll Tax Reform in Greece. *Quarterly Journal of Economics* 127(1), 493–533.
- Saez, E., J. B. Slemrod, and S. H. Giertz (2012, March). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economics Literature* 50(1), 3–50.
- Schuetz, J., V. Been, and I. G. I. Ellen (2008). Neighboring Effects of Concentrated Mortgage Foreclosures. *Journal of Housing Economics* (212), 0–36.
- Slemrod, J., C. Weber, and H. Shan (2012). The lock-in effect of housing transfer taxes: Evidence from a notched change in D.C. policy. University of Michigan, mimeo.
- Slemrod, J. B. (1990, Winter). Optimal Taxation and Optimal Tax Systems. *Journal of Economic Perspectives* 4(1), 157–78.
- Slemrod, J. B. (2010, September). Buenas Notches: Lines and Notches in Tax System Design. University of Michigan, mimeo.
- Towe, C. and C. Lawley (2013). The contagion effect of neighboring foreclosures. *American Economic Journal: Economic Policy* 5(2), 313–335.
- Van Ommeren, J. and M. Van Leuvensteijn (2005). New Evidence of the Effect of Transaction Costs on Residential Mobility. *Journal of Regional Science* 45(4), 681–702.
- Ventry Jr, D. (2010). The accidental deduction: A history and critique of the tax subsidy for mortgage interest. *Law & Contemp. Prob.* 73, 233–357.

Weber, C. (2012). Does the earned income tax credit reduce saving by low-income households?

Yang, C. S. (2012). Free At Last? Judicial Discretion and Racial Disparities in Post Booker Sentencing. *Working Paper*.