1-1-2015

Essays on Mortgage Choice and Housing Economics

Anthony DeFusco
University of Pennsylvania, anthonydefusco@gmail.com

Essays on Mortgage Choice and Housing Economics

Abstract

This dissertation consists of three self-contained chapters that study various interactions between the housing market, mortgage choice, and public policy. The first chapter studies how changes to the collateral value of real estate assets affect homeowner borrowing. While previous research has documented a positive relationship between house prices and home-equity based borrowing, a key empirical challenge has been to disentangle the role of collateral constraints from that of wealth effects in generating this relationship. To isolate the role of collateral constraints, I exploit the fully anticipated expiration of resale price controls created through an inclusionary zoning regulation in Montgomery County, Maryland. I estimate that the marginal propensity to borrow out of increases in housing collateral is between $0.04 and $0.13. The magnitude of this effect is correlated with a homeowner's initial leverage and additional analysis of residential investment and ex-post loan performance further suggests that borrowers used some portion of the extracted funds to finance current consumption and investment expenditures. These results highlight the importance of collateral constraints for homeowner borrowing and suggest a potentially important role for house price growth in driving aggregate consumption.

The second chapter, co-authored with Andrew Paciorek, provides novel estimates of the interest rate elasticity of mortgage demand by measuring the extent to which households "bunch" at a discrete jump in interest rates generated by the conforming loan limit. Our estimates imply that a 1 percentage point increase in the rate on a 30-year fixed-rate mortgage reduces first mortgage demand by between 2 and 3 percent. One-third of this response is driven by borrowers who take out second mortgages, which implies that total mortgage debt only declines by 1.5 to 2 percent.

The third chapter, co-authored with Joseph Gyourko, Fernando Ferreira, and Wenjie Ding, uses extensive micro data to investigate whether contagion was an important factor in the last housing cycle. Our estimates provide evidence of contagion during the housing boom, but not during the bust. We also find that contagion effects are greater when transmitted from a larger to a smaller market, and are more important for the most elastically-supplied markets. Local fundamentals and expectations of future fundamentals have limited ability to account for our estimated effect.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Applied Economics

First Advisor
Joseph Gyourko

Second Advisor
Fernando Ferreira

Subject Categories
Economics

ESSAYS ON MORTGAGE CHOICE AND HOUSING ECONOMICS

Anthony A. DeFusco

A DISSERTATION

in

Applied Economics

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2015

Supervisor of Dissertation

Co-Supervisor of Dissertation

Joseph Gyourko
Professor of Real Estate, Finance, and
Business Economics & Public Policy

Fernando Ferreira
Associate Professor of Real Estate and
Business Economics & Public Policy

Graduate Group Chairperson

Eric Bradlow, Professor of Marketing, Statistics, and Education

Dissertation Committee

Joseph Gyourko, Professor of Real Estate, Finance, and Business Economics & Public Policy

Fernando Ferreira, Associate Professor of Real Estate and Business Economics & Public Policy

Gilles Duranton, Professor of Real Estate

Todd Sinai, Professor of Real Estate and Business Economics & Public Policy

Nikolai Roussanov, Associate Professor of Finance

ESSAYS ON MORTGAGE CHOICE AND HOUSING ECONOMICS

© COPYRIGHT

2015

Anthony Alden DeFusco

ACKNOWLEDGEMENT

ABSTRACT

ESSAYS ON MORTGAGE CHOICE AND HOUSING ECONOMICS

Anthony A. DeFusco

Joseph Gyourko

Fernando Ferreira

This dissertation consists of three self-contained chapters that study various interactions between the housing market, mortgage choice, and public policy. The first chapter studies how changes to the collateral value of real estate assets affect homeowner borrowing. While previous research has documented a positive relationship between house prices and home-equity based borrowing, a key empirical challenge has been to disentangle the role of collateral constraints from that of wealth effects in generating this relationship. To isolate the role of collateral constraints, I exploit the fully anticipated expiration of resale price controls created through an inclusionary zoning regulation in Montgomery County, Maryland. I estimate that the marginal propensity to borrow out of increases in housing collateral is between $0.04 and $0.13. The magnitude of this effect is correlated with a homeowner's initial leverage and additional analysis of residential investment and ex-post loan performance further suggests that borrowers used some portion of the extracted funds to finance current consumption and investment expenditures. These results highlight the importance of collateral constraints for homeowner borrowing and suggest a potentially important role for house price growth in driving aggregate consumption.

The second chapter, co-authored with Andrew Paciorek, provides novel estimates of the interest rate elasticity of mortgage demand by measuring the extent to which households "bunch" at a discrete jump in interest rates generated by the conforming loan limit. Our estimates imply that a 1 percentage point increase in the rate on a 30-year fixed-rate mortgage reduces first mortgage demand by between 2 and 3 percent. One-third of this response is driven by borrowers who take out second mortgages, which implies that total mortgage debt only declines by 1.5 to 2 percent.

The third chapter, co-authored with Joseph Gyourko, Fernando Ferreira, and Wenjie Ding, uses extensive micro data to investigate whether contagion was an important factor in the last housing cycle. Our estimates provide evidence of contagion during the housing boom, but not during the bust. We also find that contagion effects are greater when transmitted from a larger to a smaller market, and are more important for the most elastically-supplied markets. Local fundamentals and expectations of future fundamentals have limited ability to account for our estimated effect.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

CHAPTER 1 : Homeowner Borrowing and Housing Collateral: New Evidence from
Expiring Price Controls

1.1. Introduction

By some accounts, over half of the mortgage debt accumulated in the U.S. during the run-up to
the Great Recession can be directly attributed to the effect of rapidly rising house prices on the
demand for home equity debt among existing homeowners (Mian and Sufi, 2011). This can be seen
clearly in Figure 1.1, which plots aggregate trends in house prices and home equity debt relative
to income over the period 1990–2014.[1] The pattern is stark. At the same time that house prices
were rising, existing homeowners were taking out an increasingly large amount of debt against
their homes—debt that they quickly began to off-load as prices collapsed.

Why do homeowners respond to rising house prices in this way? In standard models, an increase
in house prices may lead homeowners to take on additional debt due to both a direct effect on
household wealth—rising prices make homeowners feel richer—and an indirect effect on collat-
eralized borrowing capacity—rising prices relax previously binding borrowing constraints tied to
the value of the home. The goal of this paper is to isolate the empirical relevance of the latter
channel by studying how the borrowing behavior of individual homeowners responds to changes
in the collateral value of their homes.

Isolating the independent effect of collateral constraints from that of wealth effects is important
from the perspective of macroeconomic policy because the two mechanisms have markedly dif-
ferent implications for the way in which house price changes spill over into aggregate economic
activity. By propagating the effects of small shocks throughout the economy, increases in home-
owner borrowing driven by the relaxation of binding collateral constraints have the potential to
generate large swings in aggregate consumption (Bernanke and Gertler, 1989; Kiyotaki and Moore,
1997; Iacoviello, 2005). In contrast, increases in homeowner borrowing driven by wealth effects

---

[1]Greenspan and Kennedy (2005, 2008) present similar time-series evidence using a broader measure of home equity
withdrawal that includes the proceeds from cash-out refinances and home sales. The measure used in this figure includes
only equity extraction occurring through home mortgages secured by junior liens and home equity lines of credit.

are likely to have a limited impact since they will typically be offset by decreases among renters, for whom higher house prices represent a *negative* wealth shock (Sinai and Souleles, 2005; Campbell and Cocco, 2007; Buiter, 2008). Therefore, knowing whether and to what extent collateral constraints drive individual homeowner borrowing behavior is central to our understanding of how house price changes affect the real economy and to the debate over how monetary policy should respond to such changes in prices.

Existing empirical research, however, has struggled to provide direct estimates of the effect of collateral values on homeowner borrowing. Two key challenges have hindered progress. First, it is difficult to identify situations in which changes to the collateral value of a house occur independently from changes to the owner's housing wealth. As a result, most analyses have focused primarily on the *overall* effect of house prices on borrowing while attempting to infer the role of collateral constraints through the use of indirect proxies—proxies that in many cases conflate credit demand with credit supply (Gross and Souleles, 2002; Agarwal et al., 2007).[2] Second, as in most empirical analysis, omitted variables and simultaneity biases loom large. Aggregate shocks to joint determinants of house prices and homeowner borrowing, such as interest rates and expected future income, make it exceedingly difficult to draw causal inferences from naturally occurring changes in house prices, even when one is only interested in the overall relationship between prices and borrowing.[3]

In this paper, I make use of an alternative approach to contribute new empirical estimates of the

---

[2]Recent papers using this approach to study various determinants of equity extraction include Hurst and Stafford (2004); Yamashita (2007); Disney and Gathergood (2011); Mian and Sufi (2011, 2014); Cooper (2013) and Bhutta and Keys (2014). There are also a host of studies using this approach to study consumption responses to house price changes, many of which are reviewed in Bostic et al. (2009). Three important exceptions are Leth-Petersen (2010), Abdallah and Lastrapes (2012), and Agarwal and Qian (2014) who study explicit policy-induced changes in collateral constraints similar to the one studied in this paper. However, as discussed in more detail below, these studies use national and state-level policy variation, making it difficult to separately identify aggregate trends from household-specific changes to collateral constraints.

[3]A frequently proposed solution to this problem is to instrument for local house prices using Saiz's (2010) estimates of cross-city variation in physical constraints to building (Mian and Sufi, 2011, 2014; Aladangady, 2013; Mian et al., 2013). However, as cautioned by Saiz (2010) and further emphasized by Davidoff (2011, 2014), physical supply constraints are highly correlated with a host of other demand factors that might be expected to directly affect both house prices and homeowner borrowing. Moreover, as Davidoff (2014) demonstrates, physical building constraints were not correlated with changes in the size of the housing stock during the 2000s, suggesting that the correlation between house prices and building constraints was not necessarily operating through the constraints themselves during that period.

causal effect of housing collateral on home equity-based borrowing. To isolate the effect of collateral values from generalized wealth effects, I exploit a unique feature of local land use policy in Montgomery County, Maryland that drives a wedge between the value of a house as collateral and its value as a component of homeowner wealth. Since 1974, housing developers in Montgomery County have been subject to an inclusionary zoning regulation known as the Moderately Priced Dwelling Unit (MPDU) program. This policy requires developers to set aside at least 12.5 percent of all housing units in new developments to be made available at controlled prices to moderate-income households.[4] These housing units are subject to deed restrictions that cap their resale prices for a period of time ranging between 5 to 30 years. During this period, owners are not permitted to refinance or take on home equity debt for an amount that exceeds the controlled resale price. Once the price controls expire, however, owners are able to pledge the full market value of the home as collateral. Since the duration and stringency of the price controls are set by formula and known in advance at the time of purchase, their expiration has no effect on the owner's total expected lifetime wealth. However, expiring price controls directly affect the collateral value of the home through the relaxation of the borrowing restrictions. Leveraging this fact, I show within the context of a stylized model of home equity-based borrowing that differential changes in the propensity for MPDU homeowners to extract equity from their homes at the time the restriction is lifted contain explicit information regarding the effect of collateral values on homeowner borrowing. I then use that information to provide new estimates of both the extensive margin effect of relaxing collateral constraints on home equity extraction and the marginal propensity to borrow against a $1 increase in collateral value.

To conduct my analysis, I assemble a unique dataset containing the precise geographic location and detailed structural characteristics of every housing unit in Montgomery County as well as the full history of transactions and loans secured against each property during the period 1997–2012. I combine this information with administrative records from the Montgomery County Department of Housing and Community Affairs, which identify the restricted housing units and

---

[4]For a four-person household, the maximum income limit is set at 70 percent of the median family income for the Washington, D.C. metropolitan area as published annually by the U.S. Department of Housing and Urban Development. As of 2014, that limit is $75,000, which is roughly 17 percent higher than the national median family income.

the dates for which the applicable price controls were in effect. This dataset allows me to identify the effect of expiring price controls by comparing how the borrowing behavior and prices paid by owners of controlled housing units changes following the expiration of the price control relative to that of owners of nearby and observationally identical never-controlled units. It also allows me to track the borrowing behavior of a given homeowner over time, permitting a within-ownership spell comparison of equity extraction before and after the expiration of the price control. The added degrees of freedom afforded by the fact that controlled units are dispersed relatively evenly throughout the county and expire at different points during the sample period further allow me to control flexibly for aggregate trends affecting borrowing behavior and for unobservable but fixed differences across localities within the county.

I find compelling evidence that increases in collateral values lead homeowners to extract equity from their homes. In housing developments containing controlled units, transaction prices for the controlled units increase by roughly 40–65 percent relative to observationally identical non-controlled units in the same development following the expiration of the price controls. In response to these price gains, owners of controlled units are roughly four percentage points more likely to extract equity from their homes in a given year after the expiration of the price control relative to owners of non-controlled units. This effect is large, representing an almost 100 percent increase over the pre-expiration mean probability of equity extraction among owners of controlled units. Both the price effect and the increase in equity extraction among owners of controlled units are immediately present in the year the price control expires and almost perfectly offset the gap that exists between owners of controlled and non-controlled units during the imposition of the price control.

Using information on the size of individual loans, I convert these figures into an estimate of the marginal propensity to borrow against an increase in housing collateral. On average, I find that a $1 increase in collateral values leads homeowners to extract between $0.04–$0.13 in additional home equity debt. To put this into context, estimates from the literature of the *overall* effect of house prices on homeowner borrowing, which combine both collateral and wealth effects, range

between \$0.06–\$0.25.[5] Thus, my estimates imply that collateral constraints can explain a sizable fraction of the effect of house price increases on homeowner borrowing, even in the absence of any changes in perceived wealth.

To provide additional evidence that collateral constraints are the dominant force leading owners of price-controlled units to extract equity from their homes following the expiration of the price control, I also investigate heterogeneity in the response across individuals. In particular, I show that homeowners with high initial leverage (as measured by their loan-to-value (LTV) ratio at the time of purchase) are far more likely to respond to expiring price controls by extracting equity than homeowners with low initial leverage. I find no statistical or economically significant effects for homeowners in the bottom portion of the initial leverage distribution (LTV ≤ 0.7), whereas the effects for the most highly levered households (LTV > 0.95) are both statistically significant and roughly twice as large as the overall average effect. These results suggest that the increase in collateral values induced by the expiring price controls only affects borrowing behavior among the subset of homeowners for whom collateral constraints were likely to have bound prior to expiration.

My empirical strategy identifies the effect of collateral values on equity extraction under the assumption that borrowing behavior would have evolved similarly for owners of both controlled and uncontrolled units in the absence of the expiring price control. To probe the validity of this assumption, I conduct a range of different robustness checks. Most importantly, I provide direct graphical evidence showing that the trends in outcomes for controlled and uncontrolled units move together in the period prior to expiration and only begin to diverge once the price controls expire. To more formally assess the validity of the parallel trends assumption, I also conduct a series of placebo tests in which I randomly assign price control expiration dates to the controlled units and re-estimate the main specifications. The results of this exercise suggest that the effects I find are unlikely to have been generated by spurious correlation alone. The estimated effects are also robust to the inclusion of both ownership spell fixed-effects and subdivision-specific time

---

[5]See Haurin and Rosenthal (2006); Disney and Gathergood (2011); Mian and Sufi (2011, 2014), and Bhutta and Keys (2014).

trends, implying that any time-varying omitted factors driving the results must be present at both the level of the individual homeowner and the particular housing subdivision in which her home is located. Finally, to further address potential concerns regarding the comparability of controlled and never-controlled units, I also replicate the main analysis using a semi-parametric propensity score matching estimator.

While my results provide clear evidence that homeowners respond to increases in housing collateral by borrowing against their homes, the real effects of such borrowing depend on how the money is used. In particular, if homeowners simply reinvest the proceeds into more-liquid assets or use the funds to pay off other outstanding debt, then home equity-based borrowing induced by rising collateral values should not be expected to affect current consumption or investment expenditures. Although limitations of the data prevent me from being able to provide a full account of the uses of extracted funds, two pieces of evidence from the housing market suggest that at least some fraction of the borrowed money was used to fund current expenditures. First, using administrative data on building and home improvement permits issued by the Montgomery County Department of Permitting Services, I find that the annual likelihood of applying for a home improvement permit increases differentially by roughly 0.6–1 percentage points among owners of price-controlled units following the expiration of the price control. This effect represents an increase of approximately 60–100 percent over the pre-expiration mean and suggests that borrowers likely used some portion of the extracted equity to fund residential investment expenditures.[6] Second, the deeds data used to conduct the main analysis also contains information on home foreclosures. Using this information, I show that the three-year foreclosure rate associated with equity extractions secured against MPDU properties increases by roughly 1.5–2 percentage points rela-

---

[6]An alternative interpretation of this result is that the expiration of the price control increases the owner's incentives to invest in the home, as was documented by Autor et al. (2014) in the context of rent control, and may therefore explain both the increase in permitting activity and the increase in equity extraction even in the absence of any collateral effects. This is unlikely in my context for three reasons. First, the formula used to determine the controlled resale price is adjusted upward dollar-for-dollar to reflect documented home improvements and therefore generates little disincentive for investment during the control period. Second, for owners who plan to stay in the home beyond the end of the control period, the expected return from home improvements is determined based on the market price. For these owners, the exact timing of the expiration therefore has no effect on investment incentives. Finally, if the only factor driving the increase in equity extraction were changes to the demand for debt-financed home improvements, then one would not necessarily expect the equity extraction effect to be concentrated only among the set of borrowers with high initial leverage.

tive to equity extractions secured against non-MPDU properties following the expiration of the price control. This result is consistent with previous findings regarding the increased risks associated with house price-induced equity extraction and suggests that borrowers are unlikely to be reinvesting all of the proceeds into more-liquid assets, as their risk of foreclosure would presumably remain unchanged if that were the case (Mian and Sufi, 2011; Bhutta and Keys, 2014; Laufer, 2014).

Related Literature

My findings build on a large empirical literature studying the effect of house prices on household consumption, savings, and borrowing behavior. The most directly related strand of this literature has focused explicitly on the effect of house prices on home equity-based borrowing and has found consistent evidence that homeowners respond to increases in house prices by extracting equity from their homes (see, for example, Dynan and Kohn, 2007; Yamashita, 2007; Cooper, 2010, 2013; Disney and Gathergood, 2011; Mian and Sufi, 2011, 2014; Bhutta and Keys, 2014).[7] A different but closely related branch of this literature has also found consistently positive consumption responses to changes in house prices (see, for example, Skinner, 1989, 1996; Engelhardt, 1996; Lehnert, 2004; Benjamin et al., 2004; Case et al., 2005, 2013; Haurin and Rosenthal, 2006; Campbell and Cocco, 2007; Bostic et al., 2009; Gan, 2010; Carroll et al., 2011; Carroll and Zhou, 2012; Browning et al., 2013; Calormiris et al., 2013; Mian et al., 2013). Many studies have attempted to infer the role that collateral constraints play in generating these relationships through the use of indirect proxies for constraints. For example, a common approach has been to explore how households with different credit histories and incomes or at different points in the life cycle respond to similar changes in house prices.[8] While the results from these analyses are generally

---

[7]A related literature studies non-house price determinants of the demand for home equity debt, including income shocks (Hurst and Stafford, 2004) and various other sources of macroeconomic uncertainty (Chen et al., 2013). There is also an extensive literature studying the financial incentive to refinance an existing mortgage without taking on additional home equity debt in response to falling interest rates from an option theoretic point of view (see, for example, Agarwal et al., 2013; Keys et al., 2014, and references therein.)

[8]This approach is motivated by similar strategies that have been used to study the role of liquidity constraints in the vast empirical literature estimating consumption and borrowing responses to various forms of income receipt (see, for example, Zeldes, 1989; Jappelli, 1990; Shapiro and Slemrod, 1995; Jappelli et al., 1998; Parker, 1999; Souleles, 1999; Browning and Collado, 2001; Johnson et al., 2006; Agarwal et al., 2007; Stephens, 2008; Aaronson et al., 2012; Parker

suggestive of an important role for collateral constraints, the indirect nature of the proxy measures employed has left open the possibility that such estimates may be confounding differences in collateral constraints with differences in preferences.

In response to this limitation, an alternative approach has been to explore how borrowing and spending behavior responds to explicit policy-induced relaxation or tightening of collateral constraints.[9] Three recent papers taking this approach and which are closely related to this paper bear mentioning. Leth-Petersen (2010) and Agarwal and Qian (2014) provide evidence that homeowners' spending and borrowing behavior responds to national changes in the ability to use home equity debt for consumption purposes in Denmark and to the amount of time homeowners must live in their residences before being able to access home equity through cash-out refinances in Singapore, respectively. However, since both of these studies examine national policy changes that do not vary at the individual level, the authors again must rely on indirect proxy measures to determine which households would be most affected by changes in the ability to borrow against home equity. Moreover, the policy changes studied in these papers occur at the same time for all households, making it difficult to separate the effect of changing collateral constraints from other aggregate trends affecting borrowing behavior. In a similar analysis, Abdallah and Lastrapes (2012) study how retail spending at the county level in Texas responded to a statewide relaxation of a ban on the ability to use home mortgage debt for non-housing consumption. While this policy provides a clear control group, namely counties outside of Texas, the fact the policy change occurred at the same time for all households makes it difficult to control for aggregate trends, which may have driven part of the observed response. The current paper improves upon these studies in two dimensions. First, because the variation in collateral values used in this paper is specific to the individual housing unit and because the identity and location of price-controlled units are clearly demarcated in the data, I am able to avoid having to rely on proxy measures to identify the "treated" and "control" households. Instead, my approach directly compares borrowing outcomes for ob-

---

et al., 2013; Baker, 2014; Zhang, 2014 and many others as reviewed by Browning and Lusardi (1996), Browning and Crossley (2001), Jappelli and Pistaferri (2010), and Zinman (2015)).

[9]This approach is motivated by the work of Gross and Souleles (2002) and Alessie et al. (2005), both of which study the effects of supply side changes in access to credit in the consumer credit card market.

servationally identical housing units in the same housing development, some of which experience a relaxation of collateral constraints and some of which do not. Second, because the price control expiration dates that I study occur at various points throughout the business cycle, I am able to directly control for potentially confounding aggregate trends, something which is not possible when studying a single policy change that occurs at the same time for all households.

Finally, this paper also relates to a much broader literature studying the role of collateral in the macroeconomy. An important theoretical literature in macroeconomics emphasizes the role that collateral constraints can play in amplifying business cycle fluctuations through the effect of changes in asset prices on the borrowing capacity of both firms and households (see, for example, Bernanke and Gertler, 1989; Kiyotaki and Moore, 1997; Iacoviello, 2005). Given that real estate is such a large source of collateral for many households and businesses, a particular point of focus in the empirical literature studying the microeconomic foundations underlying this "financial accelerator" mechanism has been to examine how households and businesses respond to changes in the collateral value of their real estate assets.[10] This paper contributes new empirical evidence on the household side by documenting a strong positive relationship between housing collateral values and home equity-based borrowing.

The remainder of this paper is organized as follows: Section 1.2 provides institutional background on the Montgomery County Moderately Priced Dwelling Unit Program. Section 1.3 discusses how MPDU price control expirations can be used to identify collateral effects in the context of a stylized model of home equity extraction. The data sources and method used to measure equity extraction are discussed in Section 3.3. Section 1.5 outlines the empirical research design. Sec-

---

[10]For example, several recent papers have provided empirical evidence on the firm side by documenting a sizable effect of real estate prices on corporate investment, capital structure, and credit terms (Benmelech et al., 2005; Gan, 2007; Chaney et al., 2012; Cvijanovic̀, 2014), although Deng et al. (2013) find no evidence of a real estate collateral channel on firm investment in China. A related set of empirical papers has studied the relationship between house prices and entrepreneurship to test whether access to collateralized debt through home mortgages is an important determinate of small business formation and employment (Hurst and Lusardi, 2004; Schmalz et al., 2013; Adelino et al., 2014; Jensen et al., 2014). Similarly, on the household side, Caplin et al. (1997) and Lustig and Van Nieuwerburgh (2010) provide empirical evidence for a link between falling house prices, collateral constraints, and the consumption responses to regional income shocks in the U.S., while Almeida et al. (2006) present cross-country evidence suggesting that both house prices and the demand for mortgage debt are more sensitive to income shocks in countries with more generous collateral constraints.

tion 1.6 presents the main estimates of the effect of expiring price controls on collateral values and borrowing behavior and discusses heterogeneity in the borrowing response across borrowers with differing initial leverage. Evidence on the uses of extracted funds is presented in Section 1.7, and Section 3.6 concludes.

1.2. The Moderately Priced Dwelling Unit Program

Established in 1974, the Moderately Priced Dwelling Unit (MPDU) Program in Montgomery County, Maryland is one of the oldest and most well-known inclusionary zoning policies in the United States. Inclusionary zoning policies are local land use regulations that either require or incentivize housing developers to set aside a fraction of their new developments to be sold or rented to low- and moderate-income households at below-market prices. Historically, these policies have been particularly popular in high-cost suburban areas; however, in response to rising house prices and concerns over increasing spatial segregation on the basis of income, inclusionary zoning policies have grown in popularity over the last 15 to 20 years and now exist in roughly 500 municipalities across 27 states, including several large urban centers such as New York, San Francisco, Washington, D.C. and Chicago (Hickey et al., 2014).

1.2.1. Developer Requirements

The MPDU program requires that any developer wishing to build a residential development within the county containing more than 20 housing units must set aside a minimum of 12.5 percent of those units to be sold or rented to income-eligible households at controlled prices.[11] Except in rare cases, the affordable units must be provided on-site and are subject to minimum quality standards and planning guidelines that encourage the developer to scatter MPDUs among market-rate units in the same development. Figure 1.2 provides an example of the spatial distribution of MPDUs in two representative subdivisions.[12] In general, MPDU units tend to be distributed throughout the subdivision though design standards can lead to some clustering in large subdivisions since MPDUs

---

[11]If the developer agrees to provide more than the required 12.5 percent of affordable units, they are also granted density bonuses that allow for the construction of more market-rate units than would otherwise be permitted under the pre-existing zoning code. In practice, developers rarely take this option.

[12]The data used to determine the location of the price-controlled units is described in detail in Section 3.3.

are typically smaller than many of the market-rate units and are therefore often placed alongside each other. While MPDUs are permitted to be smaller in terms of interior square-footage and the construction standards provide some allowances for lower quality interior finishes (e.g. countertops and bathroom fixtures), both the planning guidelines and the private incentives of developers encourage the exterior design of MPDUs to reflect that of the nearby market-rate units. This can be seen in Appendix Figure A.1, which provides pictures of several example MPDUs and nearby market-rate units. Since its inception, the program has resulted in the creation of roughly 14,000 housing units that were price-controlled at some point in their history.[13] As of the 2010 Census, these units represented roughly 3.7 percent of the total stock of housing units in the county.[14] Approximately 70 percent of the MPDUs were originally offered for sale as owner-occupied units while the remainder were marketed as rentals. In this paper, I restrict attention to the owner-occupied portion of the program.

1.2.2. Income Limits and Eligibility

Eligibility to purchase an MPDU is restricted to first-time homebuyers who qualify for a mortgage and whose annual gross household income falls within specified ranges published annually by the Montgomery County Department of Housing and Community Affairs (DHCA). In the frequent case in which more than one eligible buyer is interested in purchasing an MPDU, the right to purchase is allocated by lottery.[15] Minimum income limits are set at the same level for all households and are meant to reflect the minimum income required to qualify for a typical mortgage on an MPDU home. Maximum income limits are pegged to the median family income for the Washington, D.C. metropolitan area as published by the U.S. Department of Housing and

---

[13]This figure is based on aggregate counts published by the Montgomery County Department of Housing and Community Affairs (DHCA) at `http://www.montgomerycountymd.gov/DHCA/housing/singlefamily/mpdu/produced.html`. The aggregate counts disagree slightly with the number of units for which I was able to obtain microdata. This is likely due to changes in administrative record keeping that led some of the older units to drop out of the DHCA database from which my data are derived.

[14]This figure is less than the mandated 12.5 percent due primarily to the durability of the housing stock. According to the 2012 American Community Survey, the median housing unit in Montgomery County was built in 1977, implying that roughly half of the housing units in the county were built before the law went into effect. The remaining gap is likely made up by housing units in smaller subdivisions to which the law does not apply.

[15]The lottery gives preference to those living or working in the county and to those who have been on the waiting list for multiple drawings.

Urban Development (HUD). For a four-person household, the maximum income limit is set at 70 percent of the area median income for a household of the same size. That limit is then scaled by an adjustment factor to determine the income limits for households of other sizes. The income limits for 2014 are shown in Appendix Table A.1. In general, the income limits are quite high, reflecting both the relative affluence of the D.C. metropolitan area and the fact that the MPDU program is meant to specifically target *moderate*-income households. For example, the maximum income limit of $75,000 for a four-person household is roughly 17 percent higher than the 2014 national median income for a household of the same size. As will be discussed in the data section below, this leads purchasers of MPDU homes to look relatively similar to the typical homebuyer in the U.S., at least in terms of household income.

1.2.3. Price Controls and Borrowing Restrictions

The initial purchase price for an MPDU is set by the DHCA according to a schedule that is meant to reflect construction costs associated with housing units of various types and sizes. Adjustments are made on a square footage basis for unit sizes deviating from those specified in the schedule and various "soft cost" adjustments are made in order to take into account developer financing costs, overhead, and other miscellaneous fixed costs of construction.

Owners of MPDUs are permitted to resell their homes. However, if the sale occurs before the end of the "control period" (a span of time ranging between 5 and 30 years depending on the initial purchase date), then the resale price is capped at the original price plus an allowance for inflation and a dollar-for-dollar adjustment that takes into account any documented major home improvements. The resale restrictions are enforced through deed covenants that are tied to the land and are released upon the first sale after the end of the control period (see Appendix Figure A.2 for an example deed covenant). Owners who sell before the end of the control period must sell their home either directly to the DHCA or to another income-eligible household on the waiting list. The first owner to sell the home after the end of the control period is permitted to sell to any buyer at the market price but is required to split any capital gains over the controlled price equally with the DHCA.

Appendix Table A.2 shows the history of rules governing the length of the control period. Prior to 2002, the control period was set as a fixed period of time from the date of the initial sale by the developer. Beginning in March 2002, the program was changed so that the control period now resets if the unit is sold at any time prior to expiration. The length of the control period was also extended from 10 years to 30 years in April 2005. These changes to the law are reflected in Figure 1.3, which plots the number of MPDU properties whose price controls expired or will expire in each year since the inception of the program. The shaded grey area marks the period of time during which I am able to observe transactions and loans.[16] The 1980s construction boom shows up as an increase in the number expiring price controls in the 1990s while the boom associated with the most recent cycle will not show up until approximately 2035.

Importantly, the owner's ability to borrow against the home is also restricted during the control period. In particular, MPDU owners are prohibited from refinancing their mortgages or taking on home equity debt for an amount that exceeds the controlled resale price. Thus, while the appraised market value of the home may be substantially higher than the controlled price, the owner is prohibited from pledging that equity as collateral until the expiration of the price control. This requirement is enforced by both the DHCA and by lenders themselves, who typically run title searches as part of the underwriting process, which would reveal any deed restrictions placed on the property. After the price controls have expired, MPDU owners are no longer restricted from borrowing more than the controlled price and, due to the shared profit agreement, are typically able to pledge up to half of the difference between the controlled price and the full market value as additional collateral.[17] As discussed in detail in Section 1.6, I estimate that the average discount for an MPDU home during the control period is between $66,000–$106,000, implying an increase in collateralized borrowing capacity of roughly $33,000–$53,000. Expiring price controls thus generally lead to a large increase in the collateral value of an MPDU owner's home, which, as the next section discusses, can be used to provide estimates of the effect of changes in housing collateral

---

[16]The transaction and loan data as well as the data used to determine the number of expiring price controls in each year is discussed in detail in Section 3.3

[17]This is because lenders are aware of the owner's obligation to the county and are thus often reluctant to extend credit for an amount beyond what the owner would receive in the event of sale.

on home equity-based borrowing.

1.3. Conceptual Framework

To illustrate how expiring price controls can be used to identify the effect of housing collateral on homeowner borrowing, this section presents a stylized model of a homeowner's equity extraction decision. I begin by considering a baseline model in which there are no price controls in order to highlight the difficulties associated with disentangling collateral effects from wealth effects using natural variation in house prices. I then show how the borrowing restrictions associated with MPDU price controls can be used to address these difficulties. The basic structure of this model draws heavily on Bhutta and Keys (2014), who use the same framework to study the effect of interest rates on equity extraction. To keep the model simple and focus the discussion on distinguishing collateral from wealth effects, I abstract from several issues that might be present in a more fully-specified life-cycle model but would otherwise not permit an analytical solution.[18] Most importantly, I assume that house prices and income are known with certainty, that households enter the world endowed with a house and a mortgage, and that they only live for two periods during which they may use their home as a source of collateral and as a source of wealth to fund consumption but from which they receive no direct utility.

1.3.1. Baseline Case: No Price Controls

Consider a household that lives for two periods, $t \in \{0, 1\}$, and is endowed with a house of value H and outstanding mortgage debt M_0 in the first period. The household has log preferences defined over non-housing consumption in each period, $u_t(c_t) = log(c_t)$, receives per-period income, y_t, and may extract equity by borrowing against the home in the first period, b_0, at the going mortgage interest rate, r, and up to an exogenous collateral constraint $\lambda H - M_0$, $\lambda \in [0, 1]$. The household chooses consumption in each period to maximize total lifetime utility

$$\max_{c_0, c_1} U(c_0, c_1) = log(c_0) + \beta log(c_1), \tag{1.1}$$

---

[18]For examples of fully specified life-cycle models that incorporate the home equity extraction decision, see Hurst and Stafford (2004) or Chen et al. (2013).

subject to constraints which are given in the baseline case by

$$c_0 = y_0 + b_0 \tag{1.2}$$

$$c_1 = y_1 - (1+r)(M_0 + b_0) + \omega H \tag{1.3}$$

$$0 \le b_0 \le \lambda H - M_0, \tag{1.4}$$

where $\beta \in [0,1]$ is the discount factor and $\omega \in [0,1]$ captures, in a reduced form way, the household's desire to consume out of housing wealth.

Differences in the parameter ω across households can arise from various sources. For example, in a life-cycle model with finitely lived households, ω will vary according to age. Younger households who plan to continue living in the same house for a longer period of time likely have lower values relative to older households who may choose to downsize in the near future (Campbell and Cocco, 2007). Similarly, ω may vary within age group due to differences in bequest motives, which lead households who wish to leave more to the next generation to consume less of their housing wealth before death. Or, as in Sinai and Souleles (2005), ω may vary by expected tenure length and by the correlation in house prices across markets to which a household is likely to move in the future. Modeling the housing wealth effect in this way, while somewhat *ad hoc*, greatly simplifies the discussion and is meant to capture these sources of heterogeneity without needing to specify a particular mechanism through which the wealth effect arises. The important point is that for some households who plan to consume part of their housing wealth before death, increases in house prices will lead to a desire to smooth consumption across periods.

Substituting the per-period budget constraints (1.2) and (1.3) into the objective function (2.5) and solving for the optimal level of equity extraction yields the solution

$$b_0^* = \begin{cases} b^* \equiv \dfrac{y_1 + \omega H - (1+r)(M_0 + \beta y_0)}{(1+r)(1+\beta)} & b^* < \lambda H - M_0 \tag{1.5a} \\[2ex] \lambda H - M_0 & b^* \ge \lambda H - M_0 \tag{1.5b} \end{cases}$$

where b^* denotes the optimal level of borrowing in the absence of the collateral constraint. This

expression highlights the empirical difficulties associated with using natural variation in house prices to disentangle collateral effects from wealth effects. To see this, consider the effect of an exogenous increase in house prices on equity extraction. For unconstrained borrowers, this effect is given by the partial derivative of (1.5a) with respect to H and is a pure wealth effect, whereas for constrained borrowers, it is equal to the partial derivative of (1.5b) and operates entirely through the collateral constraint. The empirically observable change in borrowing is therefore given by:

$$\frac{\partial b_0^*}{\partial H} = \begin{cases} \dfrac{\omega}{(1+r)(1+\beta)} & b^* < \lambda H - M_0 \\ \lambda & b^* \geq \lambda H - M_0, \end{cases}$$

(1.6a)

(1.6b)

where (1.6a) is the wealth effect for unconstrained borrowers and (1.6b) is the collateral effect for constrained borrowers. Without prior knowledge of b^*, it is impossible to know which of these two conditions applies for a given household and therefore impossible to know how much of the observed average change in borrowing in response to a change in house prices is due to wealth effects or collateral constraints.

The typical approach to solving this problem has been to infer the role of collateral constraints by examining how the magnitude of the borrowing response varies across different populations for whom one might expect either (1.6a) or (1.6b) to be the more relevant condition. For example, we might expect that not only are younger households more likely to have values of ω close to zero, but they are also more likely to have larger values of b^* as a result of steeply sloped life-cycle wage profiles that generate a gap between current and future income ($y_1 > y_0$). Thus, if younger homeowners are observed to increase borrowing more than older homeowners in response to similar increases in house prices, this could be taken as evidence for the importance of collateral constraints. However, as noted by Jappelli (1990), younger households might also be expected to be *less* constrained than older households if the rate of time preference is low relative to the real interest rate and consumption profiles are increasing in age. This ambiguity highlights the difficulty of using indirect proxy measures such as age to infer the role of collateral constraints and may explain why studies that do so have found mixed evidence (Campbell and Cocco, 2007; Mian

and Sufi, 2011; Bhutta and Keys, 2014).[19]

Another common approach is to examine heterogeneity in responsiveness across borrowers with different levels of prior debt utilization (M_0) or loan-to-value ratios ($\frac{M_0}{H}$), which shift the right-hand side of the inequality determining whether (1.6a) or (1.6b) applies. In this case, a finding that borrowers with higher loan-to-value ratios or prior debt utilization rates are more responsive to changes in prices is often taken as evidence in favor of collateral constraints (Disney and Gathergood, 2011; Mian and Sufi, 2011; Mian et al., 2013).[20] Similarly, several authors have investigated heterogeneity based on differences in income, liquid assets, or credit scores, with high-income, more-liquid, and high-credit score households expected to be less affected by collateral constraints (Yamashita, 2007; Mian and Sufi, 2011, 2014; Cooper, 2013; Bhutta and Keys, 2014). While these proxies, and loan-to-value ratios in particular, are more direct measures of collateral constraints than a homeowner's age, such proxies are nonetheless limited by their reliance on relatively strong *a priori* assumptions that are required in order to identify the set of potentially constrained households—assumptions that in many cases conflate credit demand with credit supply (Gross and Souleles, 2002; Agarwal et al., 2007). For instance, it is unclear whether homeowners with higher initial LTVs or fewer liquid assets borrow more in response to house price increases because they were unable to borrow prior to the change in prices (collateral constraints) or simply because they have stronger consumption smoothing motives that led them to carry more debt in the first place (wealth effects). More generally, as these examples illustrate, indirect proxy measures are inherently limited in their ability to distinguish differences in collateral constraints from other potential sources of heterogeneity, thus highlighting the need for estimates that are based on a more direct approach.

---

[19] Appealing to similar reasoning, Cooper (2013) finds larger consumption responses among households who experience higher realized future income growth and argues that this is evidence in favor of the role of collateral constraints.

[20] In related work, Hurst and Stafford (2004) also use loan-to-value ratios as a proxy for collateral constraints in studying how equity extraction responds to changes in interest rates.

1.3.2. Identifying Collateral Effects Using MPDU Expiration Dates

The borrowing restrictions associated with MPDU price controls drive a wedge between the value of a home as collateral and its value as a component of homeowner wealth that allows for a direct test of the role of collateral constraints. To see this, note that during the control period, an MPDU owner is prohibited from borrowing against the full market value of the property and therefore faces a more stringent collateral constraint so that equation (1.4) becomes

$$0 \leq b_0 \leq \lambda(H - \eta) - M_0, \tag{1.7}$$

where $\eta \geq 0$ denotes the MPDU price discount. For an MPDU owner who plans to stay in the home beyond the end of the control period, the eventual resale value of the home in the second period, H, remains unchanged and the optimal level of borrowing can be found by replacing equation (1.4) with equation (1.7) and resolving the utility maximization problem:[21]

$$b_0^* = \begin{cases} b^* \equiv \dfrac{y_1 + \omega H - (1+r)(M_0 + \beta y_0)}{(1+r)(1+\beta)} & b^* < \lambda(H-\eta) - M_0 \tag{1.8a} \\[2ex] \lambda(H-\eta) - M_0 & b^* \geq \lambda(H-\eta) - M_0. \tag{1.8b} \end{cases}$$

In this framework, an expiring price control is equivalent to lowering the value of η to zero in the first period while leaving the eventual resale price of the home in the second period, H, unchanged. To see how this affects borrowing, note that the effect of a decrease in η on b_0^* is given by:

$$-\frac{\partial b_0^*}{\partial \eta} = \begin{cases} 0 & b^* < \lambda(H-\eta) - M_0 \tag{1.9a} \\[2ex] \lambda & b^* \geq \lambda(H-\eta) - M_0. \tag{1.9b} \end{cases}$$

This expression makes immediately clear that borrowing should only respond to an expiring price control through the behavior of households who were collateral constrained prior to expiration. Comparing (1.9a) with (1.6a), we can see that there is no longer any role for wealth effects. Thus,

---

[21]Here, H should be thought of as the owner's expected proceeds from selling the home, net of the profit sharing agreement with the county. While the price control affects the amount of profit sharing, the key point is that for owners who plan to stay in the home beyond the end of the control period, that effect is fully anticipated so that the actual timing of the price control expiration has no effect on the expected proceeds from selling the home.

18

any observed changes in borrowing behavior at the time the price control is lifted can be entirely attributed to the effect of relaxing collateral constraints. This is the key insight underlying the empirical analysis. In the following sections, I provide empirical estimates of the magnitude of this response by studying how the borrowing behavior of MPDU owners changes around the time the price control expires relative to that of owners of observationally identical market-rate units in the same housing development for whom there is no corresponding change in collateral values.

1.4. Data and Measurement

To conduct the empirical analysis, I merge data at the property, transaction, and loan level using information from tax assessments, deeds records, and administrative data from the MPDU program. This section provides a brief overview of the data sources, variable construction, and sample selection procedures. Further details are available in Appendix A.2.

1.4.1. Data Sources

Property-Level Data

The basic structure of my dataset is organized around the 2011 Montgomery County property tax assessment file, which provides a single snapshot of all taxable properties in the county as of 2011. This file was purchased from DataQuick, a private vendor that collects and standardizes publicly available tax assessment and deeds records from municipalities across the U.S. The tax assessment file includes detailed information on the physical characteristics (e.g. square footage, number of bathrooms, number of stories, year built), use type (e.g. residential, commercial, single-family, condo), and street address for every property in the county. From this file, I drop all non-residential and multi-family properties as well as any properties with missing characteristics.[22] This leaves a "universe" of 286,484 single-family residential properties from which I select my anal-

---

[22]The assessment data in Montgomery County is of unusually high quality. Only 3,702 out of 290,186 single-family residential properties are dropped due to having missing characteristics. In 3,258 of these cases, it is the year built that is missing while other characteristics, such as square footage and number of bathrooms, are coded as zero, suggesting that many of these properties were vacant land at the time of assessment.

ysis sample. Each of these properties is geocoded and assigned a subdivision ID based on whether the geographic coordinates for the property fall within the boundaries of a particular subdivision, as delineated by the Maryland State Department of Assessments and Taxation (SDAT).[23]

To identify MPDU homes, I match the property assessment file with a list of MPDUs scraped from a publicly available online search portal hosted by the DHCA.[24] This data provides me with the street addresses for all MPDU properties in the DHCA administrative database as well as the price control expiration dates for those properties. MPDU properties were matched to the assessment file using a combination of exact physical location (geographic coordinates) and street address as described in Appendix A.2.1. Of the roughly 8,300 MPDUs in the DHCA database, I am able to match approximately 90 percent to a property in the assessment file.[25] Figure 1.4 maps the location of these properties as well as census tract-level population density for Montgomery County in 2010. In general, MPDU properties are evenly distributed across the non-rural regions of the county. One exception is the southern region of the county immediately bordering Washington, D.C., where MPDUs are underrepresented. This region contains the cities of Bethesda and Silver Spring and was developed much earlier than the rest of the county.[26] As a result, much of the housing stock in that area was not subject to the MPDU regulations at the time of development.[27]

---

[23]The subdivision boundary file was created using a parcel-level boundary file provided by the Montgomery County Planning Department. In addition to the geographic boundaries, this file also contains the SDAT subdivision ID for each parcel. The subdivision boundaries were constructed by dissolving the individual parcel boundaries into larger polygons based on whether they shared the same subdivision ID.

[24]The search portal can be accessed at http://www6.montgomerycountymd.gov/apps/DHCA/pdm_online/pdmfull.asp and was scraped using a script that exhaustively searched through and returned all possible MPDU addresses beginning with an alpha-numeric character.

[25]The match rate is lower than 100 percent largely due to poor quality record keeping in the DHCA database for some of the older MPDU properties. For example, when matching on street address, I require an exact match on the street number. Some of the older MPDU properties are missing street numbers and are therefore not included in the set of matches.

[26]This can be seen in Appendix Figure A.3, which replicates Figure 1.4 replacing population density with property age.

[27]Another reason for the underrepresentation of MPDUs in this region is that a larger fraction of the housing stock in the most densely populated areas (i.e. central cities) is composed of rental properties, which are not included in the MPDU data that I use.

Transaction and Loan-Level Data

To analyze how expiring price controls affect collateral values and homeowner borrowing, I merge the property-level file with two additional datasets from DataQuick. Both datasets are sourced from local deeds records and can be linked to properties in the assessment file using a unique property ID. The first dataset contains information on all housing transactions occurring in the county during the period 1997–2012. For each transaction, this dataset records the purchase price, buyer, seller, and lender names, as well as loan amounts on up to three loans used to finance the purchase. The second dataset contains information on all non-purchase loans secured against a property during the same period. This dataset records the initial loan amount and borrower and lender name for every refinance, junior lien, and home equity line of credit (HELOC) secured against a property. Together, these two datasets provide me with a highly granular and near complete picture of all mortgage borrowing and housing purchases occurring in the county during this period. Each dataset is cleaned as described in Appendix A.2.2 in order to ensure that the transactions represent true ownership-changing arm's length transactions and that the loan information is accurate and consistent.

1.4.2. Measuring Equity Extraction

Since the non-purchase loans dataset contains a combination of loan types but does not distinguish between them, several steps must be taken in order to construct an accurate measure of equity extraction. In particular, it is important to distinguish between three different types of non-purchase loans: (1) regular refinances, which replace an existing loan without extracting any equity; (2) cash-out refinances, which replace an existing loan with a *larger* loan, thereby extracting equity for the amount of the difference; and (3) new non-purchase originations, which directly extract equity for the amount of the new loan. In order to make this distinction, I construct a "debt history" for every property that records an estimate of the current amount of outstanding debt secured against the property at any point in time on up to two potential loans. Debt histories are constructed by

amortizing prior loan balances using the average interest rate at the time the loan was originated.[28] Given this history, when a new loan is observed, I am then able to determine whether that loan represents a purchase loan, cash-out refinance, new non-purchase origination, or regular refinance by comparing the size of the new loan to the estimated outstanding balance on the relevant existing loan (see Appendix A.2.3 for the details of this procedure). When a new refinance or purchase loan is observed, the old loan is replaced and the new loan serves as the basis for calculating remaining debt going forward.

Having categorized loans in this way, I then construct an annual panel that records for each property whether the current owner extracted equity in a particular year and if so, how much equity was extracted. I define total equity extraction in a given year as the sum of non-purchase originations and cash withdrawn through cash-out refinances during that year. Similarly, an owner is defined as having extracted equity in a given year if total equity extracted is greater than zero. For properties built prior to 1997, the panel covers the full sample period from 1997–2012; for properties built afterwards, the construction year is used as the first year of observation. Each observation is also uniquely associated with a particular "ownership-spell" for that property. Ownership spells are defined to include all years between ownership-changing transactions, where the first ownership spell starts in either 1997 or the year that the property was built. In Appendix A.2.3, I provide details validating the accuracy of this equity extraction measure against two measures provided at the aggregate level based on data from Equifax credit reports and the Freddie Mac Quarterly Cash-Out Refinance report. In both cases, my measure of equity extraction is shown to be highly correlated with national aggregates.

1.4.3. Sample Restrictions and Descriptive Statistics

Starting with the full sample of 286,484 properties, I impose several restrictions in order to arrive at my primary analysis sample. I first drop any property that could not be matched to a housing

---

[28]All loans are amortized using the average offered interest rate on a 30-year fixed rate mortgage in the month that the loan was originated. Monthly average offered interest rates are taken from the Freddie Mac Primary Mortgage Market Survey (PMMS). Since the DataQuick data do not distinguish between HELOCs and closed end liens, all loans are treated as fully amortizing with an initial principal balance equal to the origination amount, which for HELOCs, represents the maximum draw-down amount. See Appendix A.2.3 for the details of this procedure.

subdivision. This eliminates 31,603 properties located primarily in rural and outlying areas of the county where SDAT does not assign subdivision IDs. I further drop all properties located in subdivisions containing no MPDUs. This restriction eliminates 167,117 properties, many of which were located in densely populated areas consisting mostly of rental housing or in older subdivisions to which the regulation did not apply. Among subdivisions containing MPDUs, I further require that at least one MPDU expires during the DataQuick sample period. This eliminates 35,236 properties located in either older subdivisions containing only MPDUs that had already expired as of 1997 or in more recently developed subdivisions containing only MPDUs that had yet to expire as of 2012. Finally, I require that all MPDUs within a subdivision have non-missing expiration dates and that at least 95 percent of the MPDUs were matched to their corresponding DataQuick property ID in a way that required the property unit number to agree. The latter requirement is imposed because MPDUs are typically townhomes or condominiums, which, in addition to a standard street address, also have a unit number. The final sample contains 31,244 properties located in 69 subdivisions throughout the county.

Table 1.1 presents descriptive statistics for both the full sample of properties and the restricted sample used in the analysis. For the analysis sample, summary statistics are presented pooling across all properties as well as separately for non-MPDUs and MPDUs. In Panel A., the unit of observation is the individual property; in Panel B., it is the transaction; and in Panel C., it the property-year. All dollar amounts here and throughout the paper are converted to real 2012 dollars using the Consumer Price Index for All Urban Consumers (CPI-U). To limit the influence of extreme outliers, transaction prices are winsorized at the 0.5th and 99.5th percentiles in the full sample.

The differences between the full sample (columns 1–2) and the analysis sample (columns 3–4) are largely what would be expected given the nature of the sample restrictions imposed. Properties in the analysis sample are newer and larger than properties in the full sample, reflecting the fact that the MPDU regulations only apply to subdivisions constructed after 1974. Despite being newer, these properties transact at slightly lower prices than the average house in the county, again likely

23

reflecting the fact that many of the properties in the oldest and most expensive region of the county immediately bordering Washington, D.C. are located in subdivisions that do not contain MPDUs and are thus not included in the analysis sample. With regard to borrowing behavior, the average owner in the analysis sample is slightly more likely to extract equity in a given year relative to the average owner in the county but, conditional on extracting, typically borrows less.

Within the analysis sample, the differences between market-rate units (columns 5–6) and MPDUs (columns 7–8) are also largely in line with what would be expected given the nature of the MPDU program. MPDUs are smaller and substantially cheaper than non-MPDUs. This price difference is due to a combination of differences in housing characteristics and the price control itself, which mechanically lowers prices for part of the sample period even holding characteristics constant. Similarly, reflecting their lower incomes, owners of MPDUs are more likely to purchase their homes with FHA-insured loans at higher initial loan-to-value ratios relative to owners of market-rate units. Owners of MPDUs are also less likely to extract equity in a given year and, conditional on extracting, typically borrow less relative to owners of market-rate units. Some of the difference in equity extraction behavior is due to intrinsic differences in the preferences and characteristics of owners of different types of housing, while some of it is driven by the existence of the MPDU borrowing restrictions. By comparing how these overall average differences change around the time the price controls expire and controlling flexibly for aggregate trends and observable characteristics, my empirical strategy isolates the portion of the difference in prices and borrowing behavior that is driven by the MPDU program itself.

Economic and Demographic Representativeness

Given the income limits and eligibility requirements associated with the MPDU program, it is also interesting to compare how the economic and demographic characteristics of the homeowners in my analysis sample (particularly those who purchase price-controlled homes) compare to a more nationally representative sample. To that end, I match a subset of the transactions data to loan application data made publicly available through the Home Mortgage Disclosure Act (HMDA).

The HMDA data provide loan-level information on borrower income and race for nearly all home mortgage applications filed in the United States and serve as a useful gauge of the national representativeness of my sample along these dimensions.

The details of the match and several figures comparing the national distributions of income and race with the distributions for my analysis sample and the subset of transactions involving an MPDU home are provided in Appendix A.2.5. As shown in Appendix Figure A.7, the incomes of the households who purchase the price-controlled units are actually quite similar to the income of the typical homebuyer in the U.S. during my sample period. The income distribution among buyers of price-controlled homes is roughly centered around the median of the national distribution and spans a large portion of the interquartile range of that distribution. The racial breakdowns, shown in Figure A.8, are less similar. Relative to the national average, Montgomery County has an unusually high share of Asian households and slightly higher shares of Black and Hispanic households. The high Asian share likely reflects the industrial composition of the county, which is a major hub for the biotech industry, while the higher Black and Hispanic shares are most likely a result of the fact that the national sample contains many non-minority households located in the Midwest and other less diverse regions of the country. To the extent that the propensity to extract equity following a loosening of collateral constraints differs substantially by race, such differences may affect the external validity of my results. However, given the similarity in the income distributions, these differences would have to persist even conditional on income in order for the estimates derived from my sample to differ drastically from a more nationally representative sample.

1.5. Empirical Framework

1.5.1. Identification Strategy

I estimate the effect of expiring price controls on collateral values and borrowing behavior using a difference-in-differences research design that compares outcomes among market-rate units (the control group) and MPDUs (the treatment group) in the same housing development before and after the expiration of the price control. The key identifying assumption is that in the absence

of the expiring price control, the borrowing behavior and prices paid by owners of MPDUs and owners of market-rate units in the same subdivision would have evolved in parallel.

In the results section below, I present direct evidence to support the validity of the parallel trends assumption by showing that outcomes for MPDUs and non-MPDUs move together during the period prior to the expiration of the price control and that their trends only begin to diverge afterwards. This fact is both reassuring and perhaps unsurprising. While fixed differences in the characteristics of MPDU properties and their owners from those of market-rate units may lead to constant differences in prices and borrowing behavior, there is no particular reason to expect that the *evolution* of prices and borrowing over time should vary greatly across these two groups. Properties within the same subdivision are exposed to the same changes in local amenities, school quality, and crime and frequently belong to a common homeowner's association. As a result, changes in the willingness to pay of the marginal neighborhood entrant, and thus market prices, should evolve similarly for all properties in the development. Along the same lines, homeowners within a given subdivision all face the same changes in aggregate determinants of equity extraction, such as interest rates and credit standards, and should therefore be expected to display similar changes in borrowing behavior. While MPDU owners may be more likely to be subject to income shocks that could induce them to extract equity (Hurst and Stafford, 2004), the timing of such shocks would have to be highly correlated both with the expiration of the price control and across MPDU owners within a subdivision in order to generate systematic differences in trends.

1.5.2. Estimation

My baseline econometric model is a simple difference-in-differences regression fit at the individual property level. Specifically, I estimate regressions of the following form:

$$y_{ist} = \alpha_s + \delta_t + X_{it}'\gamma + \beta_1 \cdot MPDU_i + \beta_2 \cdot MPDU_i \times Post_{st} + \epsilon_{ist}, \qquad (1.10)$$

where y_{ist} is an outcome for property i in subdivision s in year t, α_s are subdivision fixed effects, δ_t are year fixed effects, X_{it} is a vector of possibly time-varying property characteristics, and ϵ_{ist}

is an error term assumed to be conditionally uncorrelated with unobserved determinants of y_{ist}. The dummy variable $MPDU_i$ is a treatment indicator that takes the value one if property i is an MPDU, while the $Post_{st}$ indicator takes the value one if year t falls on or after the year the first price control in subdivision s expires. I define the treatment date in this way to take into account both the fact that market-rate properties have no explicit expiration date and that controlled properties within the same subdivision may expire at different times as a result of construction lags and differences in initial purchase dates. Using the first expiration date is conservative and should only serve to attenuate the estimates since a small number of properties will be counted among the "treated" group before their controls actually expire.[29] The coefficient of interest is β_2, which measures the differential change in the outcome for MPDUs relative to non-MPDUs following the expiration of the price control, holding constant individual housing characteristics and aggregate differences in outcomes across subdivisions and over time. To account for serial correlation and subdivision specific random shocks, I cluster the standard errors at the subdivision level in all specifications.

One potential concern with this specification is that in addition to providing an increase in collateralized debt capacity, the expiration of the price control may also create an incentive for MPDU owners to sell their homes. As a result, differences in prices and borrowing behavior following the expiration of the price control may be driven by changes in the composition of MPDU properties that transact or changes in the characteristics of owners of MPDU homes and not necessarily the change in collateral values.[30] I address this concern directly by estimating specifications that also include property and ownership-spell fixed effects. Including property fixed effects in the price re-

---

[29]Over half of all MPDUs in my sample expire within two years of the first MPDU in their subdivision and roughly 75 percent expire within five years. These differences are most likely due to normal construction lags. Differences larger than five years likely come from one of two sources: (1) phased property development in larger subdivisions which may be built out over longer periods of time, and (2) MPDU owners reselling during the control period in the latter portion of the sample, when program rules dictated that price controls reset if the property is sold during the control period.

[30]In Appendix A.3.1, I present results showing that expiring price controls do in fact lead to an increase in housing turnover for previously controlled units of roughly three to five percentage points per year. While a three to five percentage point increase is not nearly large enough to lead to a total turnover of the MPDU housing stock during my sample period, it nonetheless justifies the use of specifications that include property or owner fixed effects due to the concern that part of the effect could be driven by differences in the transacted housing stock or changes to ownership over time.

gressions controls for changes to the set of houses that transact and identifies the effect of expiring price controls by comparing *within*-property changes in prices between MPDUs and market-rate units following the expiration of the price control. Similarly, including ownership-spell fixed effects in the equity extraction regressions controls for differences in the characteristics of owners and identifies the effect of expiring price controls by comparing *within*-owner changes in borrowing behavior between owners of MPDUs and market-rate units. In all cases, results from these specifications are not meaningfully different from those that do not include property or owner fixed effects.

As a more flexible alternative to (1.10), I also estimate specifications that allow the effect of the price control to differ by year relative to the first control period expiration. Specifically, let $\tau(s)$ denote the year the first MPDU in subdivision s expires. To capture the full time path of the effect of expiring price controls, I estimate specifications of the form

$$y_{ist} = \alpha_s + \delta_t + X'_{it}\gamma + \beta_1 \cdot MPDU_i + \sum_{\rho=-5}^{5}\left[\eta_\rho \cdot \mathbb{1}_{t-\tau(s)=\rho} + \beta_{2,\rho} \cdot MPDU_i \times \mathbb{1}_{t-\tau(s)=\rho}\right] + \epsilon_{ist},$$
$$(1.11)$$

where $\mathbb{1}_{t-\tau(s)=\rho}$ is a relative year dummy taking the value one if the current year falls ρ years after the expiration of the price control and zero otherwise. All other variables are as previously defined. The coefficients η_ρ and $\beta_{2,\rho}$ measure the baseline trend in the outcome for non-MPDUs and the differential trend for MPDUs, respectively, around the time the price control expires. I show results for up to five years preceding and following the expiration of the price control, grouping all years outside that window into the effects for relative years -5 and 5. Relative year -1 is always the omitted category so that the coefficients should be interpreted relative to the year prior to the first price control expiration within the subdivision. These coefficients are informative about both the timing of the effect of price control expirations and the validity of the parallel trends assumption. If MPDUs and non-MPDUs have common pre-trends, then the $\beta_{2,\rho}$ coefficients should be equal to zero for any $\rho < 0$.

1.6. Price Controls, Collateral Values, and Borrowing Behavior

This section presents the main estimates of the effect of expiring price controls on the transaction prices (i.e. collateral values) of previously controlled MPDUs and borrowing behavior among the owners of those properties. As an initial assessment of the validity of the parallel trends assumption, I begin by presenting simple graphical results for each of three main outcomes: (1) log transaction prices, (2) the annual probability of extracting equity, and (3) total equity extracted per year. In order to quantify the causal effects of interest, I then present a series of formal difference-in-differences estimates for each of the three outcomes. These estimates are subsequently combined to yield estimates of the marginal propensity to borrow out of increases in housing collateral. Finally, to provide additional evidence for the role of collateral constraints in governing the borrowing response to expiring price controls, I also examine heterogeneity in the response across the distribution of initial leverage. Unless otherwise specified, all results pertain to the set of transactions and property-years contained in the analysis sample described in Section 3.3.

1.6.1. Graphical Evidence

As a point of departure for the empirical analysis, Figure 1.5 plots calendar year-adjusted means for each of the three main outcomes. Means are plotted separately for MPDUs (blue circles) and non-MPDUs (orange squares) as a function of years relative to the first control period expiration within the relevant subdivision. In each panel, relative year zero represents the year the first MPDU within the subdivision expired. Means are shown for up to five years preceding and following the expiration of the price control, grouping all years outside that window into the means for relative years —5 and 5. The plotted means should be interpreted as the mean outcome among MPDUs and non-MPDUs in a given relative year, adjusted for aggregate county-wide trends affecting all properties.[31] For visual reference, the dashed lines plot a linear trend for each outcome, derived from the fitted values of a regression of the binned means on a linear term in relative year. For

---

[31]The means are adjusted for calendar year in order to remove the effect of the housing cycle, which would otherwise swamp the variation in the figure. Adjusted means were created by regressing the indicated outcome on a full set of calendar year fixed effects and averaging the residuals from that regression separately for MPDUs and non-MPDUs within relative year bins. To clarify the interpretation of the y-axis, the grand mean of each outcome was then added back in to both series.

MPDUs, only the pre-period means were used to construct the fitted values.

Consistent with the aggregate descriptive statistics presented in Table 1.1, in any given relative year, MPDU properties transact at lower prices relative to non-MPDUs and their owners are less likely to extract equity from their homes. However, for all three outcomes, the MPDU means diverge significantly from their pre-period trend starting in the year the first MPDU price control expires. There is no corresponding shift in the outcomes for non-MPDUs. As a result, a large portion of the gap in outcomes that exists during the imposition of the price control disappears once the price control expires. After 5 years, roughly half of the raw gap in prices and total equity extraction and over three quarters of the gap in the annual probability of extracting equity are eliminated. The remaining gaps reflect fixed differences in the characteristics of MPDU properties and their owners that would presumably exist even in the absence of the price control (many of which are controlled for in the analysis below). Importantly, the non-MPDU trend for each outcome is almost exactly parallel to the pre-period MPDU trend, providing strong support for the validity of the parallel trends assumption underlying the difference-in-differences estimates that follow.

1.6.2. The Effect of Expiring Price Controls on Collateral Values

To more precisely quantify the effect of expiring price controls on the collateral value of MPDU properties, Table 1.2 presents estimates from the pooled difference-in-differences specification given by equation (1.10) using log transaction prices as the outcome. The first column reports estimates from a baseline specification that includes only the MPDU main effect, the interaction of that effect with the Post indicator, fixed effects for both the year of observation and the age of the property in that year, and a series of time-invariant property characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome. Since MPDU homes are frequently built as condos and townhomes, I also fully interact the condo indicator with the year fixed effects, age fixed effects, and other time-invariant characteristics. This allows the aggregate trends, age profiles, and hedonic value of fixed property characteristics to freely vary with property type. The coefficient estimate on the MPDU main ef-

30

fect implies that during the imposition of the price control, MPDU properties sell at a discount of roughly 45 log points relative to observationally identical non-MPDU properties. This price gap is then completely eliminated following the expiration of the price control, as evidenced by the identical but opposite signed coefficient on the MPDU×Post indicator. The bottom panel also reports the implied percentage change and absolute dollar change associated with the estimated 45 log point increase.[32] Following the expiration of the price control, transaction prices at previously controlled MPDUs increase by roughly 57 percent. Applying that figure to the mean of the pre-period transaction price (in levels) among MPDUs implies an increase of roughly $93,000.

The remaining columns of the table add a series of control variables which increasingly restrict the nature of the comparison that is being used to identify the effect of the price control. In the second column, I add a set of fixed effects for each of the 69 subdivisions. This specification removes the influence of average differences in price levels across subdivisions and identifies the effect of the price control by comparing prices for observationally identical MPDUs and non-MPDUs within the same subdivision before and after the expiration of the price control. The third column not only allows for average differences in price *levels* but also allows the aggregate *trend* in prices to vary across subdivisions by interacting the subdivision fixed effects with a linear time trend. Finally, to address concerns related to differential turnover at MPDU properties after the price control expires, column four includes a full set of property fixed effects. In this specification, the time-invariant property characteristics and MPDU main effect drop out, and the effect of the price control is identified by comparing within-property changes in prices for properties that are and are not MPDUs.

The estimated effects are all highly significant and relatively stable across specifications, implying that expiring price controls lead to an increase in prices at previously controlled MPDUs that ranges from 35 to 50 log points. Converting these estimates into dollars implies that expiring price controls lead transaction prices to increase by roughly $66,000–$106,000. Due to the shared

---

[32]Following Kennedy (1981), I calculate the implied percentage increase associated with the coefficient estimate $\hat{\beta}_2$ and its standard error $\hat{\sigma}_{\beta_2}$ as $\%\Delta = 100 \times \left[\exp\left(\hat{\beta}_2 - \frac{1}{2}\hat{\sigma}^2_{\beta_2} \right) - 1 \right]$.

profit agreement, half of that increase belongs to the county while the owner retains the other half as equity. Assuming that banks are willing to lend against the full increase in equity, these estimates imply that the average MPDU owner experiences an increase in collateralized borrowing capacity of approximately $33,000–$53,000 upon the expiration of the price control.

To give a sense of the dynamics of the price effect, Figure 1.6 plots estimates from the more flexible difference-in-differences specification given by equation (1.11). These estimates are obtained from a regression that includes all of the same controls as the specification in column 3 of Table 1.2, but which allows the effect of the price control to vary separately for MPDUs and non-MPDUs by year relative to the first control period expiration. The series in orange squares plots the coefficient estimates on the ten relative year main effects, while the series in blue circles plots the sum of the relative year main effects and their interaction with the MPDU "treatment" dummy along with the 95 percent confidence interval for that sum. In both cases, relative year -1 is the omitted category, so that the two series can be interpreted as the trends for non-MPDUs (orange squares) and MPDUs (blue circles) relative to their respective values in the year prior to when the first MPDU in the subdivision went off of price control. Prices for MPDUs diverge sharply from their pre-period trend starting precisely in the year that the first price control expires. In contrast, the trend among market-rate units is completely flat. The price effect grows over time for MPDU properties, reflecting the fact that some MPDUs had yet to actually expire as of the date the first MPDU within their subdivision expired. Importantly, the trends are statistically indistinguishable in the period prior to the expiration of the price control and only begin to diverge in the year of expiration, thus lending further support to the validity of the parallel trends assumption required for identification in the difference-in-differences research design.

1.6.3. The Borrowing Response to Increases in Housing Collateral

The evidence presented in the previous section suggests that expiring price controls lead to large increases in the resale value of previously controlled MPDU homes. Assuming MPDU owners have the ability to borrow against up to half of that increase implies a similarly large increase in collateralized borrowing capacity. In this section, I explore how the borrowing behavior of

MPDU homeowners responds to the additional collateral released by expiring price controls.

Extensive Margin Borrowing Responses

I begin by considering the homeowners' extensive margin choice of whether or not to extract equity from their home. To do so, I turn to the annual property-level panel and estimate versions of the pooled difference-in-differences regression given by equation (1.10) using as the outcome an indicator for whether the property's owner extracted equity in a particular year. Table 1.3 presents the results from these regressions. In the first three columns, the control variables are the same as those used for estimating the price effect in Table 1.2 and are introduced in the same order across the columns. To account for the potential for differential turnover and changes in ownership at MPDU properties following the expiration of the price control, the fourth column includes fixed effects for each of the 57,333 unique ownership spells observed in the panel. These four specifications are all estimated using simple linear probability models. To explore the sensitivity of the results to alternative estimators, columns 5 and 6 report the marginal effects from probit and logit models estimated using the same controls contained in column 3.

Across specifications, the estimates are extremely precise and highly stable. The baseline specification in column 1, which contains only the MPDU main effect, the MPDU×Post interaction, year fixed effects, and property characteristics, indicates that during the imposition of the price control, owners of MPDU properties are on average substantially less likely to extract equity from their homes relative to owners of market-rate units but that the propensity to borrow increases differentially for MPDU owners following the expiration of the price control. The coefficient estimate on the MPDU×Post interaction term implies that expiring price controls lead to a 4.1 percentage point increase in the probability of extracting equity among owners of previously controlled units. This effect is large and represents an approximate 100 percent increase over the pre-period mean of 3.9 percentage points among MPDU owners reported in the bottom panel of the table. Adding subdivision fixed effects and their interaction with a linear time trend in columns 2 and 3 hardly changes the coefficient. Comparing the MPDU main effect with the interaction term

implies that expiring price controls close between 70–80 percent of the gap in equity extraction probabilities that exists between owners of MPDUs and non-MPDUs during the period of price control. Including ownership spell fixed effects in column 4, which restricts the comparison to take place using only within-owner variation in equity extraction, reduces the coefficient on the interaction term only slightly to 3.4 percentage points. Finally, estimating the marginal effects via probit or logit specifications in columns 5 and 6 also does not meaningfully change the magnitude of either the MPDU main effect or its interaction with the Post indicator. Taken together, the results presented in Table 1.3 suggest that in response to the increase in collateral values induced by the expiring price control, MPDU owners increase their annual probability of home equity extraction by roughly 4 percentage points, which corresponds to an increase of approximately 100 percent over their pre-period average propensity to borrow.

Figure 1.7A plots the dynamics of the effect of expiring price controls on the extensive margin probability of extracting equity. This figure is directly analogous to Figure 1.6 and was constructed from the coefficient estimates on the relative year main effects and their interaction with the MPDU "treatment" dummy as specified in equation (1.11). The regression from which these coefficient estimates are generated included all the same controls as the specification in column 3 of Table 1.3. The series in orange squares shows the trend in equity extraction among owners of non-MPDU properties, while the series in blue circles shows the trend and 95 percent confidence intervals for MPDU owners. Starting in the year the first MPDU price control expires, the MPDU trend diverges sharply from from its pre-period trend while the trend for non-MPDU owners remains smooth. Consistent with the parallel trends assumption, the equity extraction probabilities among owners of MPDUs and non-MPDUs move together in the period prior to the expiration of the price control and only diverge starting in the year of expiration. Furthermore, almost all of the increase in equity extraction among MPDU owners occurs in the year the first price control expires. This suggests that MPDU owners are responding directly to the increase in access to collateral induced by the expiring price control. The fact that the extraction probabilities remain elevated relative to their pre-period level among MPDUs beyond the first year further suggests that the removal of the price control provides owners of MPDUs with access to the same natural

increases in collateral made available to owners of non-MPDU properties through normal house price appreciation.

Combined Extensive and Intensive Margin Borrowing Responses

While the results in the previous subsection provide evidence that the likelihood of extracting equity responds significantly to the expiration of the MPDU price controls, they say nothing with respect to how the *amount* of equity extracted responds. Borrowers may respond to increases in collateral both on the extensive margin and by increasing the amount they borrow conditional on extracting equity. This section presents estimates of the combined effect of these two margins of adjustment on the annual amount of equity extracted among owners of MPDU homes.

Table 1.4 presents results from estimating the pooled difference-in-differences regression using the total amount of equity extracted in each year as the dependent variable instead of an indicator for equity extraction as before. In years when homeowners do not extract equity, this variable is set equal to zero; in years when they do extract equity, it set equal to the sum of all non-purchase originations and cash withdrawn through cash-out refinances during that year. The first four columns of the table present results from OLS regressions that are directly analogous to those presented for the extensive margin response. To take into account the fact that in many years homeowners do not extract any equity, the fifth column presents results based on a tobit specification where the equity extraction variable is treated as being censored from below at zero. This specification explicitly adjusts for the fact that the decision to extract equity may be made separately from the choice of how much to extract by estimating separate equations for the "participation" and "amount" decisions. To make the interpretation of the estimates consistent across columns, in the fifth column I report the marginal effects implied by the estimated tobit coefficients for the expected value of the censored outcome.

As with the extensive margin, the estimated response of total equity extracted is positive, statistically significant, and relatively stable across all specifications. The coefficient estimates on the MPDU×Post indicator from the OLS specifications in columns 1–4 imply that expiring price

controls lead to an increase in the average amount of equity extracted per year of roughly $2,300–$3,000. The tobit marginal effects are bigger and imply an increase of approximately $4,400. These effects are large relative to the $2,600 pre-period mean amount of equity extracted among MPDU owners. Comparing the MPDU×Post coefficient with the MPDU main effect suggests that expiring price controls close between 75–100 percent of the pre-period gap in equity extraction between owners of MPDUs and non-MPDUs. The dynamics of the effect are also shown in Figure 1.7B and mirror the results for the extensive margin reported in the same figure. The average amount of equity extracted jumps sharply among MPDU owners in the year the price control expires and remains high relative to its pre-period level for the remainder of the sample period. There is also no evidence of differential trends for MPDUs and non-MPDUs during the pre-period. While the estimates on the relative year effects for MPDU properties are less precise due to the additional variation introduced by including the intensive margin response, the conclusion remains the same. Expiring price controls lead to a substantial increase in home equity-based borrowing among owners of previously controlled units.

The Marginal Propensity to Borrow Out of Increases in Housing Collateral

A rough gauge of the economic magnitude of the borrowing responses I estimate can be provided by combining the estimates of the increase in collateral values implied by the transaction price regressions reported in Table 1.2 and the results on total equity extraction just discussed in Table 1.4. Specifically, the price results in Table 1.2 imply that expiring price controls lead to an increase in transaction prices at previously controlled MPDUs that ranges between $66,000–$106,000. Assuming that MPDU owners are able to borrow against up to half of the price increase, this implies an increase in pledgeable collateral of roughly $33,000–$53,000. The results in Table 1.4 imply that following the expiration of the price control, borrowers increase the average amount of equity they extract from their homes by roughly $2,300–$4,400 per year. Applying those estimates to the year the price control expires implies a marginal propensity to borrow out of increases in housing collateral that ranges from $0.04–$0.13.

As a point of reference, these figures can be compared to recent estimates from the literature on the *overall* effect of house prices changes on homeowner borrowing, which combine both collateral and wealth effects. For example, Haurin and Rosenthal (2006) estimate that a $1 increase in house prices leads to an increase of roughly $0.13 to $0.16 in total household debt. Disney and Gathergood (2011) estimate a smaller effect on total debt that ranges between $0.06–$0.10 and is similar to the results of Bhutta and Keys (2014), who focus explicitly on home-equity debt and provide estimates that imply a marginal propensity to borrow of roughly $0.07. Two important outliers are the estimates provided by Mian and Sufi (2011) and Mian and Sufi (2014), who focus exclusively on the most recent housing cycle and estimate marginal propensities to borrow of $0.25 and $0.19, respectively. While differences in methodology and estimation samples make it difficult to make a direct comparison between the estimates provided in this paper and those just discussed, the $0.04–$0.13 range provided above suggests that a significant portion of the effect of house prices on home equity-based borrowing is driven by collateral values rather than wealth effects.

1.6.4. Heterogeneity in the Borrowing Response by Initial LTV

In this section, I provide further evidence that the increase in home equity extraction among owners of MPDUs following the expiration of the price control is driven by the relaxation of previously binding collateral constraints by examining heterogeneity in the magnitude of the borrowing response across the distribution of initial leverage. If collateral constraints are driving the response, then we might expect those whose borrowing capacity was most limited prior to the expiration of the price control to respond more aggressively. While a borrower's initial leverage is endogenous and may be correlated with other unobservable factors determining equity extraction, it is also a relatively direct measure of collateralized borrowing capacity. Thus, evidence that borrowers with higher initial leverage respond more aggressively to the expiring price control would be consistent with a role for binding collateral constraints.

To test whether borrowers with higher initial leverage are more responsive, I restrict attention to the set of properties that are observed to transact at least once during the sample period and to the set of ownership spells that begin with a transaction. This restriction is imposed so that I

can accurately measure the initial loan-to-value ratio (LTV) associated with each ownership spell. Ownership spells are then grouped into four categories based on their initial LTV: (1) less than or equal to 70%, (2) between 70% and 80%, (3) between 80% and 95%, and (4) greater than 95%. Using these groups, I estimate the following regression:

$$y_{ijst} = \gamma_j + \delta_t + \alpha_s \cdot t + X'_{it}\gamma + \sum_{k=1}^{4} \beta_k \cdot MPDU_i \times Post_{st} \times LTV_{jk} + \epsilon_{ijst}, \qquad (1.12)$$

where y_{ijst} is an equity extraction outcome measured at time t and associated with ownership spell j of property i located in subdivision s, γ_j are ownership spell fixed effects, δ_t are year fixed effects, $\alpha_s \cdot t$ is a subdivision specific linear time trend, and X_{it} is a vector of time varying property characteristics. The primary variables of interest are the interaction terms involving the LTV_{jk} variables, which are a set of dummy variables indicating which of the four LTV groups the ownership spell belongs to. Because the specification includes ownership spell fixed effects, all time-invariant characteristics associated with either the property or the ownership spell drop out so that the vector X_{it} includes only property age and its interaction with the condo dummy and there are no main effects for the MPDU dummy or the LTV group indicators.

The β_k coefficients measure how the effect of the expiring price control varies across the distribution of initial leverage by comparing within-owner changes in borrowing behavior following the expiration of the price control across borrowers with different initial LTVs. Figure 1.8 plots these coefficient estimates along with their 95 percent confidence intervals for both the extensive margin equity extraction indicator (shown in blue bars and measured along the right axis) and the total amount of equity extracted per year (shown in orange bars and measured along the left axis). In both cases, the estimated effects for the lowest LTV group are small and statistically indistinguishable from zero while the effects for the higher LTV groups are all statistically significant and increase monotonically in initial leverage.[33] That is, MPDU owners whose initial debt is high rel-

---

[33]While the standard errors are relatively wide, one sided hypothesis tests for the difference between the highest LTV group and the lowest LTV group fail to reject the null hypothesis that the effects are larger among those in the higher group at the five percent level for both outcomes. Similarly, one sided test for whether the effects are larger among those in the two highest groups relative to those in the two lowest groups also fails to reject the null at the ten percent level for the extensive margin response and at the five percent level for the total equity extraction measure.

ative to the controlled price and whose collateralized borrowing capacity is therefore most limited during the imposition of the price control are precisely the set of borrowers who are most likely to respond to its elimination by extracting equity from their homes.

1.6.5. Additional Robustness Checks

Placebo Tests

As a further test of the parallel trends assumption underlying the main difference-in-differences estimates provided in Section 1.6.3, I also conduct a series of placebo tests for the effect of the price control on borrowing behavior. Each placebo estimate is generated by randomly assigning a false first MPDU expiration date to each of the 69 subdivisions. Using those false dates, I then replicate the pooled difference-in-differences estimate for both the extensive margin probability of equity extraction and the total amount of equity extracted per year using the specification that includes all of the property characteristics as well as the subdivision fixed effects and their interaction with a linear time trend. To prevent the placebo estimate from being influenced by any jump in the outcome at the true expiration date, I only use data from either the pre-period or the post-period depending on whether the false date falls before or after the true first expiration date for the relevant subdivision. This exercise is repeated 1,000 times and the distribution of the resulting coefficients for each outcome is plotted in Figure 1.9. The true estimate is also shown in the figure using a vertically dashed line. The true estimates are taken from column 3 of Table 1.3 for the extensive margin response and from column 3 of Table 1.4 for the total equity extraction response. As is clear from the figure, the true estimates are far larger than any of the placebo estimates, and the distribution of placebo estimates for both outcomes is centered around zero. This suggests that the results I find are unlikely to have been generated by pure chance and lends further validity to the identifying assumption of parallel trends.

Matching Estimates

Another potential concern with the main difference-in-differences estimates provided in Section 1.6.2 and Section 1.6.3 is that they rely on standard OLS estimation, which can be sensitive to differences in the distribution of covariates across "treatment" and "control" groups and relies heavily on extrapolation in areas where the covariates do not overlap (Imbens, 2004). In Appendix A.3.2, I report the results from an alternative approach to estimating the effect of the expiring price control using a semi-parametric propensity score matching estimator (cf. Heckman et al., 1997, 1998). This approach, which is described in detail in the appendix, alleviates the concern over covariate imbalance by restricting attention to a set of properties with overlapping characteristics and constructing the counterfactual outcome for each MPDU property using a locally weighted average of the outcomes among the non-MPDU properties whose characteristics are most similar. The results from this approach are reported in Appendix Table A.5 and yield estimates for the effect of the expiring price control on transaction prices, the annual probability of equity extraction, and the total amount of equity extracted per year that are all qualitatively similar to those reported above.

1.7. Evidence on the Uses of Extracted Funds

The results in the previous section provide strong evidence that homeowners respond to increases in the collateral value of their homes by extracting equity; however, the aggregate impact of this behavior depends on how the borrowed money is used. In particular, if homeowners simply use the extracted funds to pay off other existing debt or reinvest them into more liquid assets, then the aggregate effects of rising collateral values would not be as large as if the money were used to fund consumption or investment expenditures. While my data do not allow me to provide a full account of the uses of extracted funds, in this section I present two pieces of evidence from the housing market that suggest that at least some portion of the borrowed money is used to fund consumption or home improvement expenditures.

1.7.1. Evidence from Home Improvement Permits

Focusing first on home improvement expenditures, I show that expiring price controls are associated with a disproportionate increase in the likelihood of applying for home improvement permits among owners of previously controlled units. Given the concomitant increase in equity extraction documented in the previous section, this suggests that at least some portion of the extracted money was used to fund new residential investments. Of course, this result also raises the concern that the observed increase in equity extraction may not be driven by access to new collateral but by the fact that expiring price controls could increase the owner's incentives to invest in debt-financed home improvements. While disincentives for residential investment have been shown to be important in the context of rent control (Autor et al., 2014), I argue that they are unlikely to be the driving force behind the increase in equity extraction in this context for three reasons. First, the formula used to determine the controlled resale price takes into account any documented home improvements and adjusts the resale price upward dollar-for-dollar on a cost basis. Because of this, the MPDU price controls generate little disincentive for investment during the control period. Second, for owners who plan to stay in the home beyond the end of the control period, the expiration of the price control has no effect on the incentive to invest. These owners know at the time they make the investment that they will eventually receive half of its full market value, regardless of whether that investment is made before or after the price control expires. Finally, if the only factor driving the observed increase in equity extraction was a change in the demand for debt-financed home improvements, then such an effect should presumably manifest itself equally across all MPDU owners. However, as shown in the previous section, the increase in equity extraction is concentrated primarily among the set of homeowners with high initial leverage, for whom collateral constraints are presumably more important. For these reasons, it seems likely that the direction of causality runs from equity extractions (induced by increased access to collateral) to home improvements and not the other way around.

I use data from the Montgomery County Department of Permitting Services to estimate the effect of expiring price controls on residential investment behavior. This data, which is described in

further detail in Appendix A.2.6, contains address level information on all building and home improvement permit applications filed since 2000 for all parts of the county except for the cities of Gaithersburg and Rockville. I match the permit applications to the DataQuick property file using the same approach used to match the list of MPDU addresses. Having matched the data, I then construct an annual panel that records for each property located outside of Gaithersburg or Rockville whether a permit application was filed for that property in a particular year. The permits data includes applications for both new construction and improvements. To avoid confusing new construction with home improvements, I only include property-year observations that are at least two years after the year the property was built. The panel thus runs from 2000–2012, unless the property was built during that time period, in which case the data beings in the second year after the property was built.

Using this panel, I estimate versions of the pooled difference-in-differences regression given by equation (1.10) where the outcome is now an indicator for whether a home improvement permit application was filed for a property in a particular year. Table 1.5 reports the coefficient estimates for the MPDU×Post interaction term, which measures the differential increase in the likelihood of applying for permitted residential investment among owners of MPDU properties following the expiration of the price control. Across the columns, the control variables and specifications are the same as those used for estimating the extensive margin equity extraction response and are introduced in the same order. While the smaller sample size leads to a modest loss of precision, all of the estimated effects are positive and significant at the five percent level. For the OLS specifications, the point estimates imply that expiring price controls lead to an increase in the probability of filing a home improvement permit application of roughly 1 percentage point while the probit and logit specifications yield slightly lower estimates of roughly 0.6–0.7 percentage points. These effects are large relative to the pre-period mean of 1 percentage point among MPDU owners reported in the bottom panel and suggest that borrowers likely used some portion of the equity they extracted to fund new residential investment expenditures.

1.7.2. Evidence from Foreclosures

The second piece of evidence I provide regarding the uses of extracted funds draws inferences based on the ex-post performance of the loan. If equity extraction is merely a means for portfolio diversification or paying off existing debt, then one would not necessarily expect that the act of extracting equity itself would expose borrowers to additional risk or raise their probability of mortgage default and subsequent foreclosure. On the other hand, if borrowers use some of the extracted funds to pay for current consumption and investment expenditures, then their total leverage would increase, potentially putting them at higher risk of default and foreclosure (Bhutta and Keys, 2014). In this section, I provide evidence that equity extractions induced by the increase in collateral values at the time the price controls expire are more likely to end in foreclosure relative to equity extractions driven by other motives, which suggests that they are also more likely to be used for the purposes of funding current expenditures.

While the DataQuick data does not allow me to track the time at which a particular loan becomes delinquent or enters foreclosure, it does contain an indicator for whether an ownership transfer occurred as a result of a foreclosure sale or bank repossession. Using this information, I am able to determine for every loan observed in the loan-level data whether that loan was followed by a subsequent foreclosure. To measure differences in foreclosure rates associated with individual instances of equity extraction, I restrict attention to loans in the non-purchase loans dataset that are coded as equity extractions and estimate versions of the following regression specified at the loan level

$$Foreclosure_{ijst} = \alpha_s + \delta_t + X'_{ijt}\gamma + \beta_1 \cdot MPDU_j + \beta_2 \cdot MPDU_j \times Post_{st} + \epsilon_{ijst}, \quad (1.13)$$

where $Foreclosure_{ijst}$ is an indicator denoting whether equity extraction i associated with property j located in subdivision s and originated at time t was followed by a foreclosure within up to three years after its origination. Since the focus is on three-year foreclosure rates, I only include equity extractions that occur between 1997–2009 to ensure that I can observe up to three years

of potential foreclosure information for every loan. In addition to the standard set of property characteristics, the vector X_{ijt} also includes dummy variables indicating whether the loan was FHA-insured or had an adjustable interest rate. The coefficient of interest is β_2, which measures the differential increase in three-year foreclosure rates associated with equity extractions secured against MPDU properties following the expiration of the price control relative to the change in foreclosure rates associated with equity extractions secured against non-MPDU properties. A positive value for β_2 suggests that equity extractions that occur in response to the increased collateral made available by the expiring price control are at higher risk of foreclosure relative to equity extractions motivated by other factors.

Table 1.6 reports estimates of this coefficient from various versions of equation (1.13). All specifications include the standard set of property characteristics as well as the dummies for FHA and adjustable rate mortgages and fixed effects for the year of origination and the age of the property in that year. The second column adds subdivision fixed effects, which are further interacted with a linear time trend in column 3. Columns 4 and 5 report probit and logit marginal effects using the same specification as in column 3. In all cases, the estimated effect on the MPDU×Post interaction term is positive and precisely estimated. The estimates imply that the three-year foreclosure rate associated with equity extractions secured against MPDU properties increased by roughly 1.5–2 percentage points relative to equity extractions secured against non-MPDU properties following the expiration of the price control. This effect represents between 70–90 percent of the overall average three-year foreclosure rate among equity extractions secured against MPDU properties. Equity extractions induced by expiring price controls are thus substantially more risky than those motivated by other reasons, which suggests that they are also more likely to have been used for the purposes of funding current expenditures rather than simply paying off existing debt or portfolio diversification.

1.8. Conclusion

For many households, houses are both the largest asset they own and the most readily available source of pledgeable collateral against which they can borrow. Changes in the value of housing

thus have the potential to lead to significant changes in the desire and ability of individual home-owners to borrow. The macroeconomic consequences of house price-induced changes in individual borrowing behavior depend crucially on whether those changes are driven by wealth effects, which in aggregate may be offset by opposing changes among renters and others who are "short" housing, or by the relaxation of binding collateral constraints, which do not have the same offsetting effects. Empirical analyses of the effect of house price fluctuations on homeowner borrowing often struggle to distinguish between these two channels, as it is difficult to find instances in which changes to the value of a home do not represent both a direct increase in a household's net worth and an indirect expansion of their collateralized borrowing capacity.

This paper addresses these challenges and isolates the role of collateral values by exploiting a unique feature of an inclusionary zoning policy in Montgomery County, Maryland that imposes temporary price controls on owner-occupied housing units. Because the duration and stringency of these price controls are set by formula and known in advance at the time of purchase, their expiration has no effect on the owner's expected lifetime wealth. Changes in borrowing behavior among owners of controlled units at the time of expiration can thus be directly attributed to the effect of the price control on the collateral value of the home. Using this fact, I show that changes in the collateral value of housing have important effects on homeowner borrowing behavior. Specifically, following the expiration of the price controls, the probability of home equity extraction increases differentially among owners of previously controlled units by roughly 4 percentage points relative to owners of observationally identical non-controlled units in the same housing development. Comparing the increase in equity extraction to the increase in available collateral implied by the change in prices at the time of expiration yields an estimate of the marginal propensity to borrow out of increases in housing collateral of approximately $0.04–$0.13. These estimates are roughly within the range of existing estimates from the literature of the *total* effect of house price increases on homeowner borrowing and suggest that collateral constraints are an important factor driving that relationship. The magnitude of the borrowing response is also monotonically increasing in the homeowner's initial leverage, providing further evidence that collateral constraints are the dominant force leading owners of controlled units to respond to the expiration of the price

control by borrowing against their homes. Finally, evidence from home improvement permit applications and subsequent loan performance suggests that at least some portion of the extracted funds were used to finance current consumption or investment expenditures and not simply used as a means for portfolio diversification or paying down existing debt. These results corroborate existing evidence on the importance of collateral constraints based on indirect proxy measures and have implications both for understanding the microeconomic mechanisms driving the relationship between house prices and homeowner borrowing and the macroeconomic consequences of that relationship.
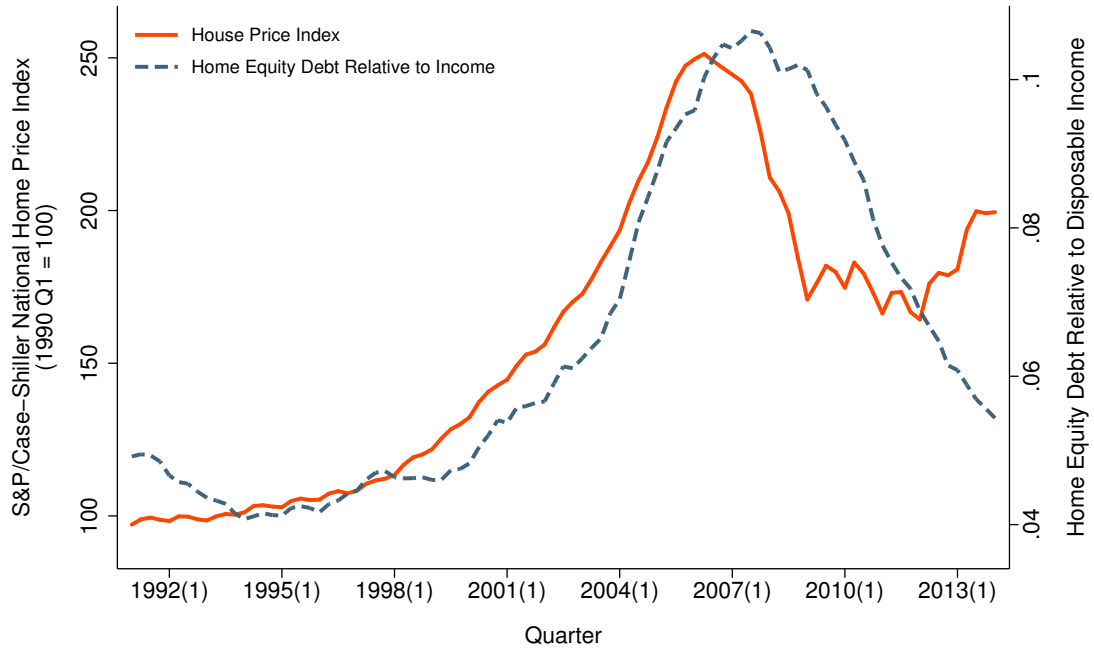
FIG. 1.1.—House Prices and Home Equity Debt Relative to Disposable Personal Income. This figure plots aggregate trends in U.S. house prices and home equity debt relative to disposable personal income at a quarterly frequency over the period 1990–2014. Aggregate data on home equity debt come from the Federal Reserve Flow of Funds (Z.1 Release, Series FL893065125.Q). Disposable personal income data come from the National Income and Product Accounts (BEA Account Code A067RC1). The home equity debt series is only available from the Flow of Funds beginning in the fourth quarter of 1990. The house price index is normalized to 100 in the first quarter of 1990, and the underlying data come from the S&P/Case-Shiller National Home Price Index.

47

FIG. 1.2.—Spatial Distribution of MPDU Properties in Two Example Subdivisions. This figure shows the location of several MPDU properties in two representative subdivisions. MPDU properties are marked with an orange circle. All unmarked homes are market-rate units. The two shaded areas identify the subdivision boundaries.

FIG. 1.3.—Number of Expiring Price Controls by Year. This figure plots the trend in price control expiration dates at a yearly frequency for all owner-occupied MPDU properties in the county. Each dot plots the number of properties whose price controls expired or will expire in the indicated year. The shaded grey area marks the period of time for which information on housing transactions and home equity-based borrowing is availble from DataQuick.
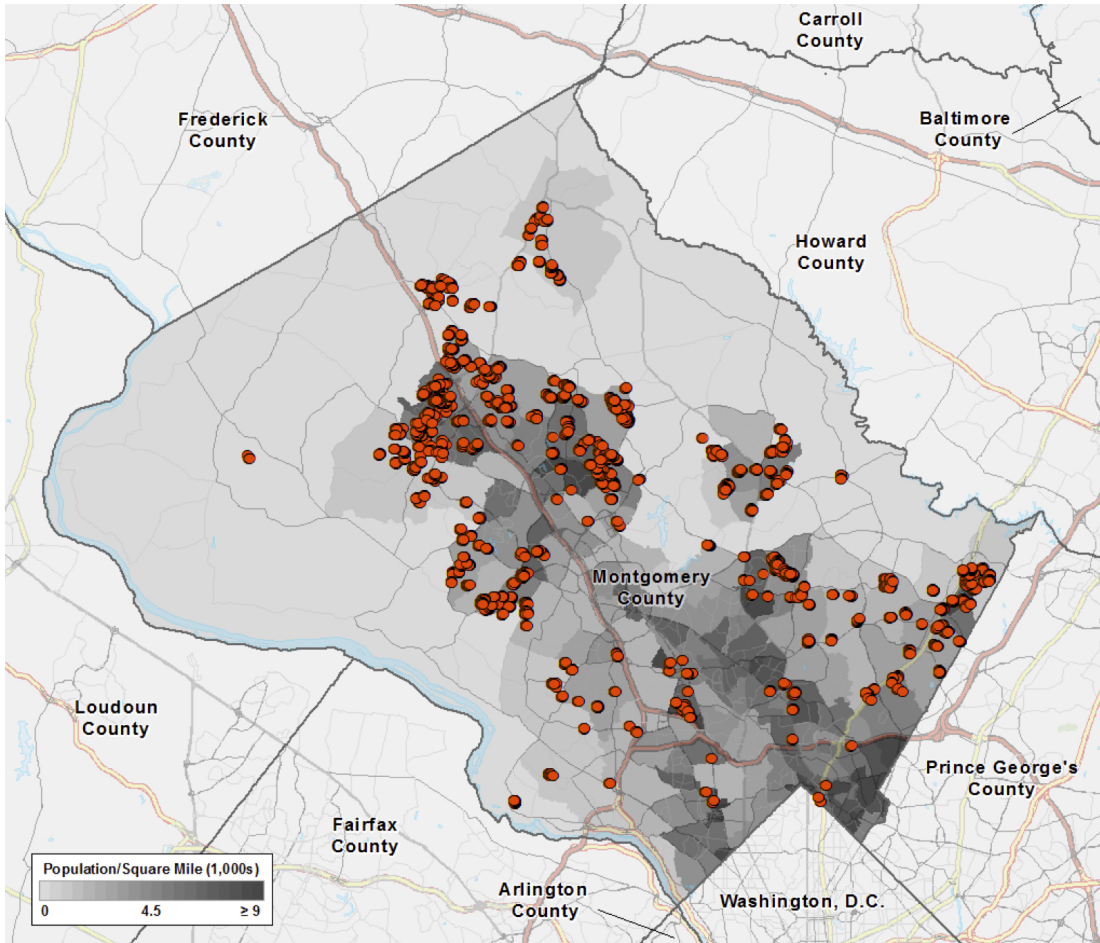
FIG. 1.4.—Geographic Distribution of MPDU Properties within Montgomery County, Maryland. This figure shows the location of all MPDU properties that were successfully matched to a property in the DataQuick assesment file (N=7,404). MPDU properties are marked with an orange circle. Census tracts within Montgomery County are shaded according to their population density as reported in the 2010 American Community Survey.

Panel A. Log Transaction Price

Panel B. Probability of Extracting Equity

Panel C. Total Equity Extracted ($1,000s)
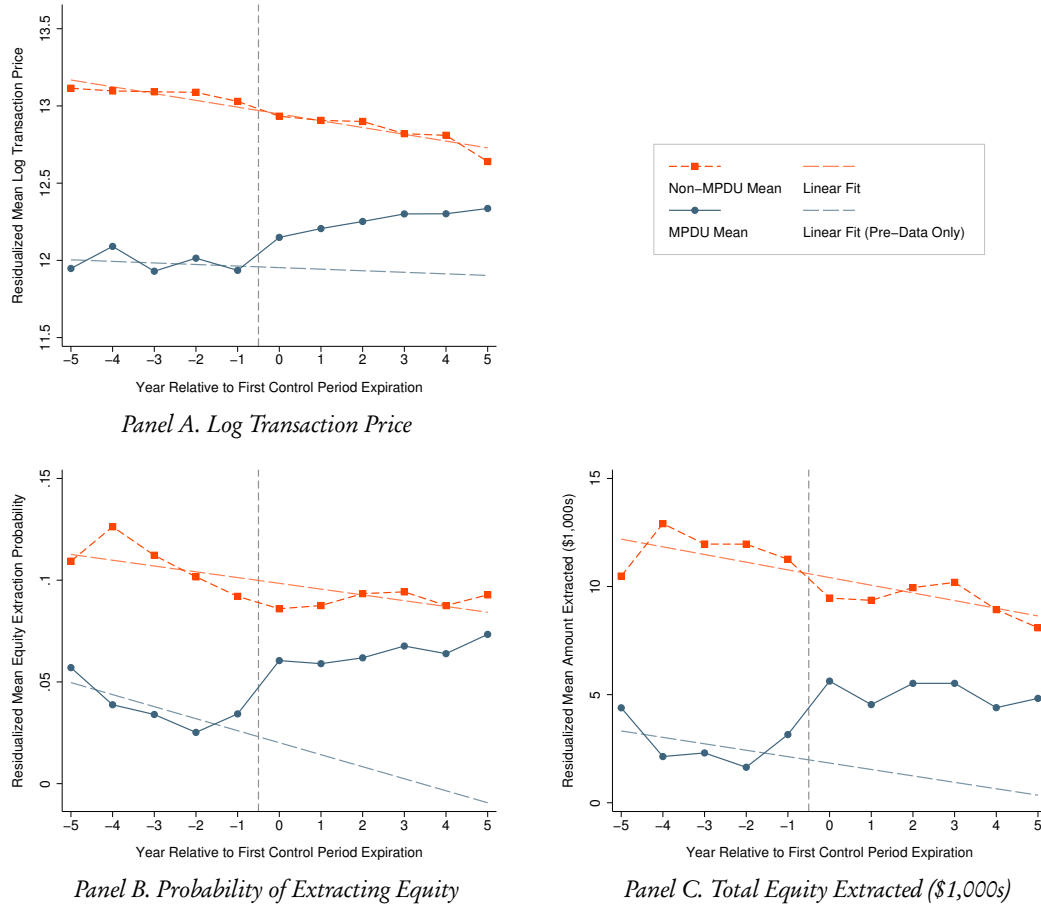
FIG. 1.5.—Assessing the Parallel Trends Assumption. This figure plots calendar year-adjusted means for each of the three primary outcomes—log transaction prices, the annual probability of extracting equity, and total equity extracted per year—separately for MPDUs (blue circles) and non-MPDUs (orange squares) as a function of years relative to the first control period expiration within the relevant subdivision. Relative year zero represents the year the first MPDU within the subdivision expired. Means are shown for up to five years preceding and following the expiration of the price control, grouping all years outside that window into the means for relative years −5 and 5. Adjusted means were created by regressing the indicated outcome on a full set of calendar year fixed effects and averaging the residuals from that regression separately for MPDUs and non-MPDUs within relative year bins. To clarify the interpretation of the y-axis, the grand mean of the outcome was then added back in to both series. The dashed lines plot the fit from a regression of the binned means on a linear term in relative year. For MPDUs, only the pre-period means were used to construct the fitted values.
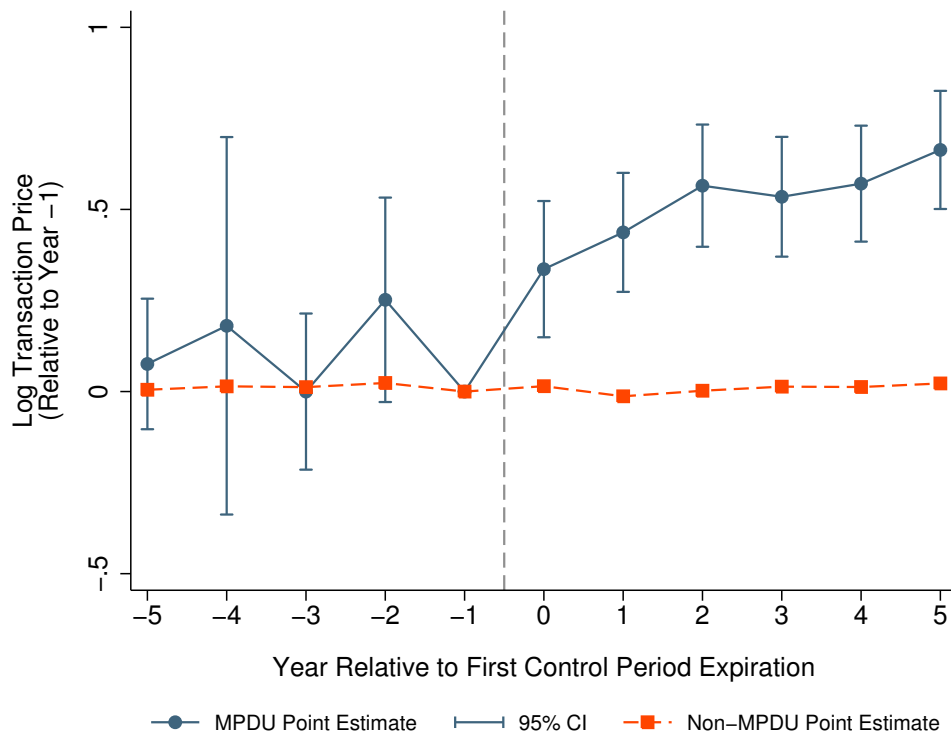
FIG. 1.6.—Dynamic Effects of Expiring Price Controls on Transaction Prices of MPDU Properties. This figure reports estimates of the effect of expiring price controls on the transaction prices of MPDU properties derived from a flexible difference-in-differences regression that allows the effect to vary by year relative to the expiration of the price control. Estimates were constructed by regressing the log of the transaction price on an indicator for whether the associated property is an MPDU and the interaction of the MPDU indicator with a series of dummy variables indicating whether the year of observation falls in a given relative year as measured from the year the first MPDU in the relevant subdivision expired. Relative year zero denotes the year the first price control in the subdivision expired. Relative year −1 is the omitted category so that all estimates should be interpreted as relative to the year prior to expiration. Results are shown for five years preceding and following the expiration of the price control, with all years outside that window grouped into the effects for relative years −5 and 5. The series in orange squares plots the coefficient estimates on the relative year main effects, which represent the trend in log prices among non-MPDU properties. The series in blue circles plots the estimate and 95 percent confidence interval for the sum of the relative year main effects and the interaction of those effects with the MPDU indicator, representing the trend among MPDU properties. The 95 percent confidence intervals are based on standard errors which were clustered at the subdivision level. The regression also included year fixed effects, subdivision fixed effects and their interaction with a linear time trend and a set of property characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms, stories, and property age, as well as an indicator for whether the property is a condo or townhome and the interaction of that indicator with the year fixed effects and all of the other property characteristics.

Panel A. Probability of Extracting Equity   *Panel B. Total Equity Extracted ($1,000s)*

Fig. 1.7.—Dynamic Effects of Expiring Price Controls on the Borrowing Behavior of MPDU Homeowners. This figure reports estimates of the effect of expiring price controls on the borrowing behavior of MPDU homeowners derived from a flexible difference-in-differences regression that allows the effect to vary by year relative to the expiration of the price control. Estimates were constructed by regressing an indicator for whether the homeowner extracted equity in a particular year (Panel A.) and the total amount of equity extracted per year (Panel B.) on an indicator for whether the associated property is an MPDU and the interaction of the MPDU indicator with a series of dummy variables indicating whether the year of observation falls in a given relative year as measured from the year the first MPDU in the relevant subdivision expired. Relative year zero denotes the year the first price control in the subdivision expired. Relative year −1 is the omitted category so that all estimates should be interpreted as relative to the year prior to expiration. Results are shown for five years preceding and following the expiration of the price control, with all years outside that window grouped into the effects for relative years −5 and 5. The series in orange squares plots the coefficient estimates on the relative year main effects, which represent the trend in annual equity extraction probabilities (Panel A.) and average equity extracted per year (Panel B.) among non-MPDU properties. The series in blue circles plots the estimate and 95 percent confidence interval for the sum of the relative year main effects and the interaction of those effects with the MPDU indicator, representing the trend among MPDU properties. The 95 percent confidence intervals are based on standard errors that were clustered at the subdivision level. The regressions also included year fixed effects, subdivision fixed effects and their interaction with a linear time trend and a set of property characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms, stories, and property age, as well as an indicator for whether the property is a condo or townhome and the interaction of that indicator with the year fixed effects and all of the other property characteristics.

FIG. 1.8.—Heterogeneity Across the Distribution of Initial Leverage in the Effect of Expiring Price Controls on the Borrowing Behavior of MPDU Homeowners. This figure reports estimates of the effect of expiring price controls on the borrowing behavior of MPDU homeowners derived from a difference-in-differences regression that allows the effect to vary according the the homeowner's initial LTV. Estimates for the extensive margin probability of extracting equity are shown in blue bars and measured along the right axis while estimates for the total amount of equity extracted per year are shown in orange bars and measured along the left axis. The 95 percent confidence intervals for each estimate are also shown and are based on standard errors that were clustered at the subdivision level. The height of each bar corresponds to the coefficient estimate on the triple interaction term between an indicator for whether the property is an MPDU, an indicator for whether the year of observation falls on or after the year the first price control in the relevant subdivision expired, and an indicator for whether the initial LTV for the ownership spell fell with in the range indicated on the x-axis. The regressions also included fixed effects for the ownership spell, the year of observation, and the age of the property in that year, as well as subdivision specific linear time trends and the interaction between the age fixed effects and an indicator for whether the property is a condo or townhome. To be able to accurately measure initial leverage, the sample was restricted to the set of properties that were observed to transact at least once and to the set of ownership spells that began with a transaction (N=211,249).

54

Panel A. Probability of Extracting Equity

Panel B. Total Equity Extracted ($1,000s)
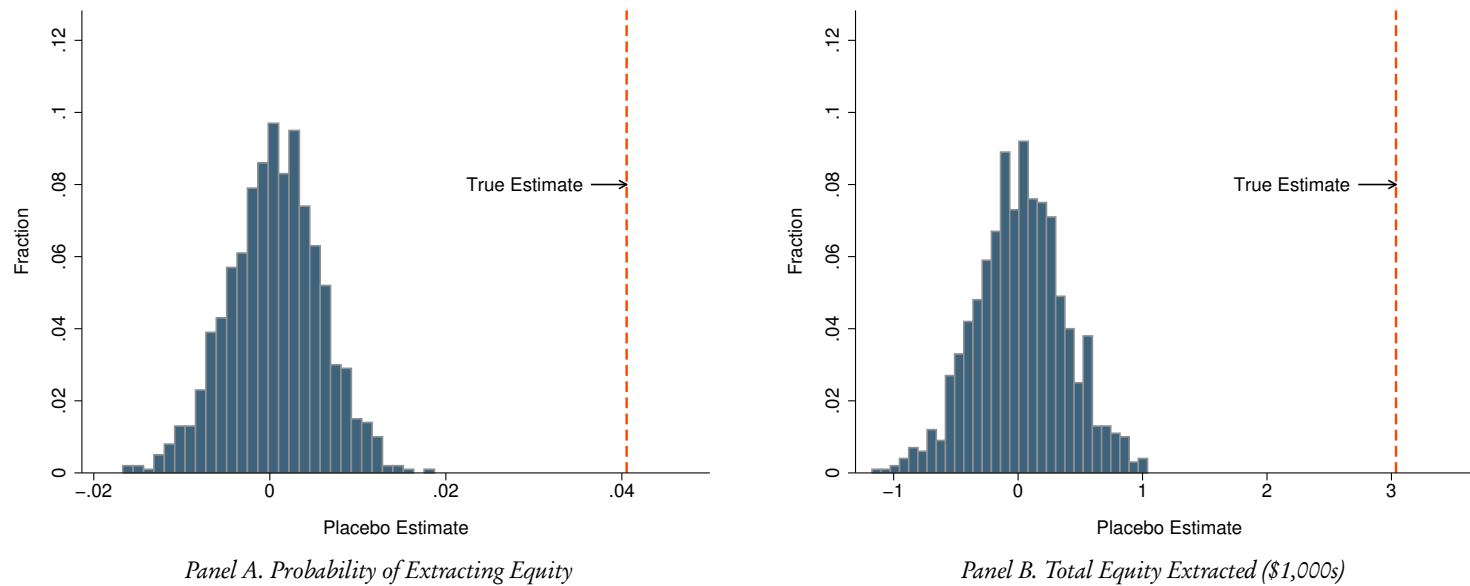
FIG. 1.9.—Placebo Tests of the Effect of Expiring Price Controls on the Borrowing Behavior of MPDU Homeowners. This figure reports results from a series of placebo tests of the effect of expiring price controls on the probability of extracting equity (Panel A.) and the total amount of equity extracted per year (Panel B.) among owners of MPDU properties. Each panel plots the distribution of 1,000 placebo estimates for the indicated outcome. The vertically dashed lines show the true estimates, which were taken from column 3 of Table 1.3 for Panel A. and column 3 of Table 1.4 for Panel B. Each placebo estimate was created by randomly assigning a false first MPDU expiration date to each of the 69 subdivisions and generating a difference-in-differences estimate using that false date. To prevent the placebo estimate from being influenced by any jump in the outcome at the true expiration date, only data from either the pre-period or the post-period was used depending on whether the false date fell before or after the true first expiration date within the relevant subdivision. In addition to the MPDU main effect and its interaction with the Post indicator, the regressions used to generate the placebo estimates include year fixed effects, subdivision fixed effects and their interaction with a linear time trend and a set of property characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms, stories, and property age, as well as an indicator for whether the property is a condo or townhome and the interaction of that indicator with the year fixed effects and all of the other property characteristics.

TABLE 1.1

SUMMARY STATISTICS FOR PROPERTIES, TRANSACTIONS, AND ANNUAL MEASURES OF EQUITY EXTRACTION

	Full Sample		Analysis Sample					
	All Properties		All Properties		Non-MPDUs		MPDUs	
Panel A. Property Characteristics								
Square Footage (1000's)	1.90	(4.50)	1.98	(1.01)	2.07	(1.02)	1.18	(0.22)
Number of Bathrooms	2.62	(1.17)	3.00	(1.05)	3.08	(1.04)	2.26	(0.77)
Number of Stories	1.65	(0.53)	1.86	(0.42)	1.85	(0.41)	1.95	(0.53)
Age (Years)	37.80	(20.28)	24.73	(9.86)	25.10	(10.13)	21.18	(5.66)
Number of Properties	286,484		31,244		28,278		2,966	
Panel B. Transaction/Buyer Characteristics								
Transaction Price ($1000's)	444.09	(330.85)	418.09	(290.70)	432.81	(294.80)	237.37	(140.34)
Loan Amount ($1000's)	351.18	(673.69)	330.10	(200.84)	339.53	(202.00)	210.58	(137.71)
Loan-to-Value Ratio	0.85	(0.57)	0.86	(0.61)	0.85	(0.58)	0.94	(0.90)
Fraction FHA Insured	0.14	—	0.16	—	0.15	—	0.27	—
Number of Properties	156,206		19,152		17,719		1,433	
Number of Transactions	233,879		30,209		27,934		2,275	
Panel C. Equity Extraction Measures								
Amount Extracted ($1000's)	7.63	(40.97)	8.45	(40.75)	8.85	(42.08)	4.63	(24.37)
Amount Extracted $\mid > 0$ ($1000's)	101.08	(113.07)	92.48	(101.96)	94.14	(103.91)	69.93	(66.31)
Probability of Extracting Equity	0.08	—	0.09	—	0.09	—	0.07	—
Number of Properties	286,484		31,244		28,278		2,966	
Number of Ownership Spells	493,386		57,333		52,390		4,943	
Number of Property-Years	4,383,768		483,805		437,977		45,828	

NOTE.—This table presents descriptive statistics for both the full sample of single-family residential properties with non-missing housing characteristics contained in the DataQuick assessment file (columns 1–2) and the restricted sample used in the analysis (columns 3–8). All table entries represent sample means or, in parentheses, standard deviations. For the analysis sample, summary statistics are presented pooling across all properties as well as separately for non-MPDUs and MPDUs. In Panel A., the unit of analysis is the individual property; in Panel B., it is the transaction; and in Panel C., it the property-year. All dollar amounts are converted to real 2012 dollars using the Consumer Price Index for All Urban Consumers (CPI-U). Transaction prices are winsorized at the 0.5th and 99.5th percentiles in the full sample.

TABLE 1.2
THE EFFECT OF EXPIRING PRICE CONTROLS ON
THE TRANSACTION PRICES OF MPDU PROPERTIES

	(1)	(2)	(3)	(4)
MPDU	-0.455***	-0.616***	-0.620***	
	(0.070)	(0.067)	(0.067)	
MPDU × Post	0.455***	0.499***	0.505***	0.349**
	(0.077)	(0.084)	(0.083)	(0.139)
Property Characteristics	X	X	X	
Year and Age FEs	X	X	X	X
Subdivision FEs		X	X	
Subdivision Trend			X	X
Property FEs				X
Implied %Δ	57%	64%	65%	40%
Implied $Δ (1,000s)	$93	$104	$106	$66
R-squared	0.76	0.81	0.81	0.94
Number of Observations	30,209	30,209	30,209	30,209

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on transaction prices for MPDU properties. Each column reports a separate regression estimated at the transaction level where the dependent variable is the log of the transaction price. Coefficients are reported for the "treatment" dummy, denoting whether the property is an MPDU, and the interaction of that dummy with an indicator for whether the year of observation falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the 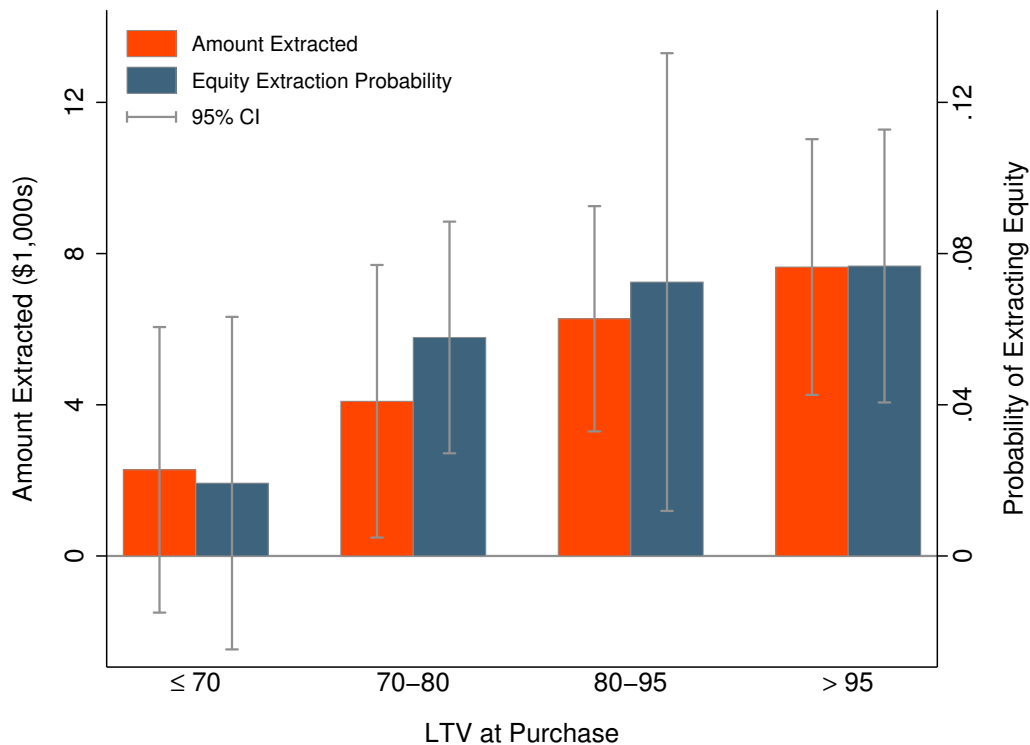year fixed effects and all of the other property characteristics including property age. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. The first row of the bottom panel reports the implied percentage increase in prices associated with the coefficient estimate on the MPDU×Post indicator reported in the second row of the table. The implied dollar increase is calculated by applying that percentage increase to the mean price (in levels) among MPDU properties in the period prior to the expiration of the price control. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 1.3

THE EFFECT OF EXPIRING PRICE CONTROLS ON THE ANNUAL
PROBABILITY OF EXTRACTING EQUITY AMONG OWNERS OF MPDU PROPERTIES

	OLS				Probit	Logit
	(1)	(2)	(3)	(4)	(5)	(6)
MPDU	-0.057***	-0.050***	-0.050***		-0.053***	-0.054***
	(0.007)	(0.008)	(0.008)		(0.006)	(0.006)
MPDU × Post	0.041***	0.040***	0.041***	0.034***	0.044***	0.042***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)
Property Characteristics	X	X	X		X	X
Year and Age FEs	X	X	X	X	X	X
Subdivision FEs		X	X		X	X
Subdivision Trend			X	X	X	X
Ownership Spell FEs				X		
Pre-Expiration MPDU Mean	0.039	0.039	0.039	0.039	0.039	0.039
Number of Observations	483,805	483,805	483,805	483,805	483,805	483,805

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on the annual probability of extracting equity among owners of MPDU properties. Each column reports a separate regression estimated at the property-year level where the dependent variable is an indicator for whether the property owner extracted equity from the home in a particular year. Coefficients are reported for the "treatment" dummy, denoting whether the property is an MPDU, and the interaction of that dummy with an indicator for whether the year of observation falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. Columns 1–4 report coefficient estimates from linear probability models, while columns 5–6 report marginal effects from probit and logit specifications. The mean of the dependent variable among MPDU properties in the period prior to the first price control expiration is reported in the second to last row. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 1.4

THE EFFECT OF EXPIRING PRICE CONTROLS ON TOTAL EQUITY
EXTRACTED PER YEAR (IN $1,000S) AMONG OWNERS OF MPDU PROPERTIES

	OLS				Tobit
	(1)	(2)	(3)	(4)	(5)
MPDU	-2.572***	-3.514***	-3.607***		-5.939***
	(0.755)	(0.664)	(0.693)		(0.632)
MPDU × Post	2.319***	2.926***	3.037***	2.794***	4.407***
	(0.715)	(0.664)	(0.700)	(0.880)	(0.667)
Property Characteristics	X	X	X		X
Year and Age FEs	X	X	X	X	X
Subdivision FEs		X	X		X
Subdivision Trend			X	X	X
Ownership Spell FEs				X	
Pre-Expiration MPDU Mean	2.662	2.662	2.662	2.662	2.662
Number of Observations	483,805	483,805	483,805	483,805	483,805

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on the annual amount of equity extracted among owners of MPDU properties. Each column reports a separate regression estimated at the property-year level where the dependent variable is the amount of equity (in $1,000s) that the property owner extracted from the home in a particular year. Coefficients are reported for the "treatment" dummy, denoting whether the property is an MPDU, and the interaction of that dummy with an indicator for whether the year of observation falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. Columns 1–4 report coefficient estimates from OLS regressions, while column 5 reports the marginal effects for the expected amount of equity extraction (censored and uncensored, treating censored as zero) from a tobit specification. The mean of the dependent variable among MPDU properties in the period prior to the first price control expiration is reported in the second to last row. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 1.5

THE EFFECT OF EXPIRING PRICE CONTROLS ON THE ANNUAL PROBABILITY
OF PERMITTED RESIDENTIAL INVESTMENT AMONG OWNERS OF MPDU PROPERTIES

	OLS				Probit	Logit
	(1)	(2)	(3)	(4)	(5)	(6)
MPDU × Post	0.0089**	0.0086**	0.0092**	0.0136**	0.0071**	0.0063**
	(0.0034)	(0.0037)	(0.0037)	(0.0055)	(0.0031)	(0.0031)
Property Characteristics	X	X	X		X	X
Year and Age FEs	X	X	X	X	X	X
Subdivision FEs		X	X		X	X
Subdivision Trend			X	X	X	X
Ownership Spell FEs				X		
Pre-Expiration MPDU Mean	0.0106	0.0106	0.0106	0.0106	0.0106	0.0106
Number of Observations	385,192	385,192	385,192	385,192	385,192	385,192

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on the annual probability of permitted residential investment among owners of MPDU properties. Each column reports a separate regression estimated at the property-year level where the dependent variable is an indicator for whether the property owner filed an application for a home improvement permit in a particular year. The table reports the coefficient on the interaction term between the "treatment" dummy, denoting whether the property is an MPDU, and an indicator for whether the year of observation falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. Columns 1–4 report coefficient estimates from linear probability models, while columns 5–6 report marginal effects from probit and logit specifications. The mean of the dependent variable among MPDU properties in the period prior to the first price control expiration is reported in the second to last row. The sample excludes properties located in the cities of Gaithersburg and Rockville, where permit application data is not available, and property-year observations occurring prior to 2000, the first year that permit applications are observed. To avoid mistaking new construction for home improvements, it also excludes any property-year observations occurring less than two years after the property was built. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 1.6

THE EFFECT OF EXPIRING PRICE CONTROLS ON THE THREE-YEAR
FORECLOSURE RATE AMONG EQUITY EXTRACTIONS SECURED AGAINST MPDU PROPERTIES

	OLS			Probit	Logit
	(1)	(2)	(3)	(4)	(5)
MPDU × Post	0.017***	0.015***	0.015***	0.020***	0.020***
	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)
Property Characteristics	X	X	X	X	X
Loan Characteristics	X	X	X	X	X
Year and Age FEs	X	X	X	X	X
Subdivision FEs		X	X	X	X
Subdivision Trend			X	X	X
Dep. Var. MPDU Mean	0.022	0.022	0.022	0.022	0.022
Number of Observations	45,719	45,719	45,719	45,719	45,719

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on the three-year foreclosure rate among equity extractions secured against MPDU properties. Each column reports a separate regression estimated at the loan level where the dependent variable is an indicator for whether the loan was followed by a foreclosure within three years of origination. The sample includes only non-purchase loans coded as equity extractions that were originated during the period 1997–2009. The table reports the coefficient on the interaction term between the "treatment" dummy, denoting whether the property is an MPDU, and an indicator for whether the year of origination falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of origination and the age of the property in that year as well as property and loan characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. The loan characteristics include an indicator for whether the loan was FHA-insured and whether it had an adjustable interest rate. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. Columns 1–3 report coefficient estimates from linear probability models, while columns 4–5 report marginal effects from probit and logit specifications. The mean of the dependent variable among MPDU properties is reported in the second to last row. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

CHAPTER 2 : The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit

2.1. Introduction

Buyers face a bewildering array of financing options when purchasing a home. Should they pay cash or take out a mortgage? If the latter, should it have a fixed rate or an adjustable rate? How large a down payment should they make? Given that housing makes up the lion's share of most owners' portfolios, these and related questions are fundamental to their financial well-being. Yet there is little research that credibly identifies how households respond to changes in the many parameters of this problem. In this paper, we focus on one element of the problem—the choice of how much debt to incur—in order to provide novel and credible estimates of the interest rate elasticity of mortgage demand.

The magnitude of this elasticity has important implications for policy-relevant questions in several areas of economics. For example, given that mortgages constitute the majority of total household debt, the elasticity plays a significant role in governing the degree to which monetary policy affects aggregate consumption and savings behavior (Hall, 1988; Mishkin, 1995; Browning and Lusardi, 1996). In public finance, the elasticity is also important for understanding the effect of the home mortgage interest deduction on both government tax revenue and household consumption (Poterba, 1984; Poterba and Sinai, 2008, 2011). Similarly, the elasticity also has implications for the effects of government intervention in the secondary mortgage market, where federal policy directly influences mortgage rates through the purchase activity of the government-sponsored enterprises (GSEs), Fannie Mae and Freddie Mac (Sherlund, 2008; Adelino et al., 2012; Kaufman, 2012). This final consideration has become particularly salient recently in light of the ongoing debate over the future of the GSEs in the wake of the 2007-2008 financial crisis.

Despite these potentially important policy implications, there are essentially no existing causal estimates of the extent to which individual loan sizes respond to changes in interest rates.[1] This

---

[1]One exception is in recent work by Fuster and Zafar (2014), who survey households about their mortgage choices

is due in large part to data limitations and lack of plausibly exogenous variation in interest rates, which which have led prior research in this area to focus on other aspects of mortgage choice. The literature estimating interest rate elasticities of other smaller components of consumer credit—such as credit card, auto, and micro-finance debt—has been more fruitful, thanks to the availability of detailed microdata and variation in interest rates arising from either direct randomization or quasi-experimental policy changes (Gross and Souleles, 2002; Alessie et al., 2005; Attanasio et al., 2008; Karlan and Zinman, 2008, 2014).[2] In the spirit of these studies, we estimate the interest rate elasticity of mortgage demand using microdata on over 2.7 million mortgages and an identification strategy that leverages "bunching" at nonlinearities in household budget constraints.

We identify the effect of interest rates on borrower behavior by exploiting a regulatory requirement imposed on the GSEs that generates exogenous variation in the relationship between loan size and interest rates. Specifically, the GSEs are only allowed to purchase loans for dollar amounts that fall below the conforming loan limit (CLL), a nominal cap set by their regulator each year. Because loans purchased by the GSEs are backed by an implicit government guarantee, interest rates on loans above this limit ("jumbo loans") are typically higher than rates on comparable loans just below the limit.

While identifying the precise magnitude of this interest rate differential is challenging due to borrower sorting around the limit, some insight can still be gleaned from examining the raw data.[3] For example, Figure 2.1A plots the interest rate as a function of the difference between the loan amount and the conforming limit for all fixed-rate mortgages in our analysis sample that were originated in 2006.[4] There is a clear discontinuity precisely at the limit, with average interest rates

---

under randomized hypothetical interest rate scenarios and find results similar to what we report below. In other related work, Martins and Villanueva (2006) estimate how the extensive margin probability of obtaining a mortgage responds to an interest rate subsidy for low-income households in Portugal but do not report direct estimates of the effects on loan size. Similarly, several others including Follain and Dunsky (1997), Ling and McGill (1998), Dunsky and Follain (2000), and Jappelli and Pistaferri (2007) have estimated how mortgage debt responds to changes in the rate at which interest expenses can be deducted from personal income but do not focus explicitly on effective interest rates themselves.

[2]See Zinman (2015) for a review of the empirical literature on demand elasticities in consumer credit markets.

[3]Many papers have attempted to overcome this challenge using a variety of empirical methods. See, for example, Hendershott and Shilling (1989), Passmore et al. (2002), Passmore et al. (2005), Sherlund (2008) and Kaufman (2012). We address these issues in detail below in Section 2.4.

[4]Each dot is the average interest rate within a given $5,000 bin relative to the limit. The dashed red lines are the predicted values from a regression fit using the binned data, allowing for changes in the slope and intercept at the limit.

on loans just above the limit being approximately 20 basis points higher than those on loans just below.

The difference in interest rates between jumbo and conforming loans creates a substantial "notch" in the intertemporal budget constraint of households deciding how much mortgage debt to incur. This notch induces some borrowers who would otherwise take out loans above the conforming limit to instead bunch right at the limit. This behavior is confirmed by Figure 2.1B, which shows the fraction of all loans in our analysis sample which fall into any given $5,000 bin relative to the conforming limit in effect at the date of origination.[5] Consistent with the notion that borrowers bunch at the conforming limit, the figure shows a sharp spike in the fraction of loans originated in the bin immediately below the limit, which is accompanied by a sizable region of missing mass immediately to the right of the limit.

The intuition behind our empirical strategy is to combine reasonable estimates of the jumbo-conforming spread with a measure of the excess mass of individuals who bunch at the conforming limit to back out estimates of the interest rate (semi-)elasticity of demand for mortgage debt. Recent papers in public finance have developed methods for estimating behavioral responses to non-linear incentives in similar settings (Saez, 2010; Chetty et al., 2011; Kleven and Waseem, 2013).[6] We adapt these methods to the case of mortgage choice in the face of a notched interest rate schedule. To the best of our knowledge, ours is the first application of these methods to the mortgage market or to a consumer credit market of any kind.

Our preferred specifications indicate that the average size of a borrower's first (fixed-rate) mortgage declines by between 2 and 3 percent for each 1 percentage point rise in the first mortgage rate, holding constant all prices and interest rates that do not change at the conforming limit. Because both the bunching estimates and the jumbo-conforming spread estimates vary depending on the

---

See Section 2.3 for details on sample construction. The year 2006 is chosen for illustrative purposes only. We estimate the jumbo spread using all available loans below in Section 2.4.

[5]During our sample period the conforming limit varied from around $215,000 in 1997 to it's peak in 2006 and 2007 of approximately $420,000. Appendix Figure A.11 shows the full path of the limit in real an nominal terms during this period.

[6]Other recent applications of these and similar methods include Sallee and Slemrod (2010); Manoli and Weber (2011); Best and Kleven (2015); Chetty et al. (2013); Gelber et al. (2013) and Kopczuk and Munroe (2014).

assumptions used in estimation, we also provide alternative estimates under a range of different scenarios. These estimates imply a decline of between 1.5 and 5 percent for a 1 percentage point increase in the mortgage rate.

While the mortgage demand elasticity is of innate interest, its interpretation depends in part on the channels through which borrowers adjust their first mortgage balance. Our second main contribution is to provide suggestive evidence on this margin. Borrowers can reduce the initial balance of their first mortgage in at least three ways. First, they can make a larger down payment on the same house at the same price. Second, they can take out a second mortgage to cover the loan balance in excess of the conforming limit. Third, they can lower the price of the house they buy, either by negotiating with the seller or by choosing a less expensive house. We show that about one-third of bunching borrowers take out second mortgages, which suggests that the reduction in *total* mortgage debt in response to a 1 percentage point rise in the first mortgage interest rate is between 1.5 and 2 percent.

To gauge the economic magnitude of the effects we estimate, we apply them to recently proposed increases to the fee that the GSEs charge lenders to cover the costs associated with guaranteeing investor returns on their mortgage-backed securities. We estimate that the proposed fee increases would reduce the total dollar volume of fixed-rate conforming mortgage originations by approximately one-fifth of one percent. When we apply our elasticity to similar increases in fees that have occurred in the recent past, we estimate an effect on the order of one-half of one percent.

The remainder of the paper is organized as follows: Section 2.2 presents our conceptual framework. In Section 2.3 and Section 2.4, we discuss our data and empirical research design. We then present our main results in Section 2.5-Section 2.6. Section 2.7 applies these results to changes in the GSE guarantee fees, and Section 2.8 concludes by discussing avenues for future research.

2.2. Theoretical Framework

We begin by considering a simple two-period model of mortgage choice.[7] Although highly stylized, this model highlights the most relevant features of our empirical environment and generates useful predictions for household behavior in the presence of a nonlinear mortgage interest rate schedule. The model is similar in spirit to those in the recent literature in public finance studying behavioral responses to nonlinear incentives in other contexts. For example, Saez (2010), Chetty et al. (2011), Chetty et al. (2013), and Gelber et al. (2013) study labor supply and earnings responses to kinked income tax and social security benefit schedules. Similar models have also been developed to study behavioral responses in applications more analogous to ours, where the budget constraint features a notch as opposed to a kink. Our analysis draws heavily on the framework developed by Kleven and Waseem (2013) and Best and Kleven (2015) who study behavioral responses to notched income and real estate transfer tax schedules, respectively. Kopczuk and Munroe (2014) have also used this framework to study real estate transfer taxes and others have used it to study fuel economy regulation (Sallee and Slemrod, 2010) and retirement incentives (Manoli and Weber, 2011). Ours is the first application to the mortgage market or to a credit market of any kind.

2.2.1. Baseline Case: Linear Interest Rate Schedule

Households live for two periods. Since our primary focus is on the intensive margin choice of how much debt to incur conditional on purchasing a home, we shut down housing choice by assuming that each household must purchase one unit of housing services in the first period at an exogenous per-unit price of p.[8] Households can finance their housing purchase with a mortgage, m, which may not exceed the total value of the house. The baseline interest rate on the mortgage is given by r and does not depend on the mortgage amount. In the second period, housing is liquidated, the mortgage is paid off, and households consume all of their remaining wealth.

The household's problem is to maximize lifetime utility by choosing consumption in each period,

---

[7]The underlying theory is similar to that in Brueckner (1994), among other papers.

[8]In the appendix, we relax the assumption that households cannot choose the quantity of housing services to consume. While the intuition is the same, extending the model in this way prevents us from deriving a closed-form solution, which makes it less useful for motivating the empirical work.

denoted by c_1 and c_2.[9] In general, the household solves:

$$\max_{c_1,c_2}\{U(c_1,c_2) = u(c_1) + \delta u(c_2)\} \tag{2.1}$$

$$\text{s.t. } c_1 + p = y + m \tag{2.2}$$

$$c_2 = p - (1+r)m \tag{2.3}$$

$$0 \leq m \leq p, \tag{2.4}$$

where $\delta \in (0,1)$ is the discount factor and y is first period income. Solving equation (2.2) for c_1 and substituting this, along with equation (2.3), into equation (2.1) allows us to rewrite the household's problem in terms of mortgage debt,

$$V = \max_{m}\{u(y + m - p) + \delta u(p - (1+r)m)\}, \tag{2.5}$$

subject now only to the borrowing constraint (2.4).

To proceed, we make several simplifying assumptions. First, we assume that household preferences are given by the constant elasticity function $u(c) = \frac{1}{1-\xi}c^{1-\xi}$.[10] Second, heterogeneity in the model is driven by the discount factor, which is assumed to be distributed smoothly in the population according to the density function $f(\delta)$. For illustrative purposes, we assume that y, p, and ξ are constant across households; however, this assumption is not crucial and we discuss below how relaxing it affects the interpretation of our results.[11] Finally, we assume that households end up at an interior solution with a positive mortgage amount and a loan-to-value ratio of less that 100 percent—that is, constraint (2.4) does not bind.

---

[9]Since we impose the exogenous requirement that households consume one unit of housing services, we suppress the argument for housing consumption and express the household's problem as a choice over non-housing consumption only.

[10]This functional form allows us to derive a closed-form solution, but all of the basic results hold with more general utility functions.

[11]We model heterogeneity primarily through the discount factor in order to facilitate the graphical discussion of bunching behavior below. While household income and the intertemporal elasticity of substitution are also potentially important sources of heterogeneity in our model, these factors only affect mortgage choice through shifts in the lifetime budget constraint and through differences in the curvature of indifference curves. Differences in discount factors, on the other hand, affect mortgage choice by linearly shifting indifference curves *along* the lifetime budget constraint. This allows us to graphically represent how a notched interest rate schedule induces bunching behavior in a simple budget constraint diagram.

Under these assumptions, we can solve explicitly for mortgage demand:

$$m^* = \frac{p - (\delta(1+r))^{1/\xi}(y-p)}{(\delta(1+r))^{1/\xi} + (1+r)}. \tag{2.6}$$

Because ξ, y, and p are assumed to be constant across households, this relationship provides a one-to-one mapping between a household's value of δ, and its optimal mortgage choice when faced with the baseline interest rate schedule.[12] Given the assumption of a smooth distribution for δ, this mapping will induce a smooth baseline distribution of mortgage amounts, which we denote using the density function $g_0(m)$.

2.2.2. Notched Interest Rate Schedule

We now consider the effect of introducing a notch in the baseline interest rate schedule at the conforming loan amount $\bar{m}$. Loans above this limit are subject to a higher interest rate, leading to the new schedule $r(m) = r + \Delta r \cdot \mathbb{1}(m > \bar{m})$.[13] Here, Δr is the difference in interest rates between jumbo and conforming loans and $\mathbb{1}(m > \bar{m})$ is an indicator for jumbo loan status. Combining equations (2.2) and (2.3) yields the lifetime budget constraint

$$C = y - m \cdot [r + \Delta r \cdot \mathbb{1}(m > \bar{m})], \tag{2.7}$$

where $C = c_1 + c_2$ is lifetime consumption. This budget constraint is plotted in Figure 2.2A along with indifference curves for two representative households.

The notch in the budget constraint induces some households to bunch at the conforming loan limit. In Figure 2.2A, household L is the household with the lowest baseline mortgage amount—the largest δ—who locates at the conforming limit in the presence of a notch. This household is unaffected by the change in rates and takes out a loan of size $\bar{m}$ regardless of whether the notch exists. Household H is the household with the highest pre-notch mortgage amount—the smallest

---

[12]Technically, for this mapping to be one-to-one it must be true that $y > \frac{r}{1+r}p$. If this condition holds then m^* is strictly decreasing in δ. This is likely to be the case for any reasonable values of r and p.

[13]The institutional details of the conforming loan limit and the channels through which it leads to higher interest rates on jumbo loans are provided in the appendix.

δ—that locates at the conforming limit when the notch exists. When faced with a linear interest rate schedule, this household would choose a mortgage of size $\bar{m} + \Delta\bar{m}$. With the notch, however, the household is indifferent between locating at $\bar{m}$ and the best interior point beyond the conforming limit, m^I. Any household with a baseline mortgage amount in the interval $(\bar{m}, \bar{m} + \Delta\bar{m}]$ will bunch at the conforming loan amount, $\bar{m}$. Furthermore, no household will choose to locate between $\bar{m}$ and m^I in the notch scenario.

The density when a notch exists, $g_1(m)$, will therefore be characterized by both a mass of households locating precisely at the conforming limit as well as a missing mass of households immediately to the right of the limit. The effect of the notch on the mortgage size distribution is shown in the density diagram in Figure 2.2B. The solid black line shows the density of loan amounts in the presence of the notch and the heavy dashed red line to the right of the notch shows the counterfactual density that would exist in the absence of the conforming loan limit.

Because households can be uniquely indexed by their position in the pre-notch mortgage size distribution, the number of households bunching at the conforming limit is given by:

$$B = \int_{\bar{m}}^{\bar{m}+\Delta\bar{m}} g_0(m)dm \approx g_0(\bar{m})\Delta\bar{m}, \tag{2.8}$$

where the approximation assumes that the counterfactual no-notch distribution is constant on the bunching interval $(\bar{m}, \bar{m} + \Delta\bar{m}]$.[14]

The expression in equation (2.8) is the primary motivation for our empirical strategy. Given estimates of the amount of bunching, $\hat{B}$, and the counterfactual density at the conforming loan limit, $\hat{g}_0(\bar{m})$, we can solve for $\Delta\bar{m}$, the behavioral response to the interest rate difference generated by the conforming limit. This behavioral response represents the reduction in loan size of the marginal bunching individual. Scaling this response by an appropriate measure of the change in the effective interest rate yields an estimate of the interest rate elasticity of mortgage demand.[15]

---

[14]This approximation merely simplifies the discussion. In the empirical application we allow for curvature in the counterfactual distribution.

[15]Much of the structure in the model above is not needed for this result to hold. All we require is that households can be uniquely indexed by their choice of mortgage size in the pre-notch scenario and that the counterfactual distribution

2.2.3. *Heterogeneous Intertemporal Elasticities, Incomes, and Prices*

The derivation of equation (2.8) rests upon the assumption that ξ, y, and p are constant across households. In that case, it is possible to back out the exact change in mortgage amount for the marginal bunching individual. When intertemporal elasticities, incomes, and prices are allowed to vary across households, the amount of bunching instead identifies the *average* response among the marginal bunching individuals associated with each intertemporal elasticity, income, and price level. To see this, let the joint density of discount factors, intertemporal elasticities, incomes, and prices be given by $\bar{f}(\delta, \xi, y, p)$, where $y \in (0, \bar{y}]$, $\xi \in (0, \bar{\xi}]$, and $p \in (0, \bar{p}]$ for some upper bounds, $\bar{y}$, $\bar{\xi}$, and $\bar{p}$. For a fixed (ξ, y, p) triple, the bunching interval is determined in exactly the same way as in the baseline model. Denote this interval $(\bar{m}, \bar{m} + \Delta \bar{m}_{\xi,y,p})$, where $\Delta \bar{m}_{\xi,y,p}$ is the behavioral response of the marginal bunching individual among those with intertemporal elasticity $1/\xi$, income y, and who face price p. Further, let $\bar{g}_0(m, \xi, y, p)$ denote the joint density of mortgage sizes, intertemporal elasticities, incomes, and prices in the pre-notch scenario and $g_0(m) \equiv \int_\xi \int_y \int_p \bar{g}_0(m, \xi, y, p)\, dp\, dy\, d\xi$ the unconditional mortgage size density. The amount of bunching can then be expressed as

$$B = \int_\xi \int_y \int_p \int_{\bar{m}}^{\bar{m}+\Delta\bar{m}_{\xi,y,p}} \bar{g}_0(m, \xi, y, p)\, dm\, dp\, dy\, d\xi \approx g_0(\bar{m}) E\left[\Delta \bar{m}_{\xi,y,p}\right]. \qquad (2.9)$$

In this case, estimates of bunching and the counterfactual mortgage size distribution near the conforming limit allow us to back out the average change in mortgage amounts due to the interest rate difference generated by the conforming loan limit.[16]

2.3. Data

To conduct our empirical analysis, we use data on loan sizes and interest rates from two main sources. The first is a proprietary data set of housing transactions from DataQuick (DQ), a private vendor which collects the universe of deed transfers and property assessment records from munic-

---

of mortgage sizes be smooth. Any model for which these conditions hold would generate equation (2.8).

[16]Kleven and Waseem (2013) show a directly analogous result in the context of earnings responses to notched income tax schedules.

ipalities across the U.S. These data serve as our primary source of information on loan size. The second data source consists of loan-level records collected by Lender Processing Services (LPS) and contains extensive information on interest rates, borrower characteristics, and loan terms, which we use to estimate the jumbo-conforming spread. A brief description of each data source and our sample selection procedures is given below.

2.3.1. DataQuick

Each record in the DQ data set represents a single transaction and contains information on the price, location, and physical characteristics of the house, as well as the loan amounts on up to three loans used to finance the purchase. We restrict the sample to include only transactions of single-family homes with positive first loan amounts that took place within metropolitan statistical areas (MSAs) in California between 1997 and 2007. We use data from California because that is where the information from DataQuick is most reliable, particularly for identifying when multiple loans were used to finance a purchase. In addition, because average house prices in California are higher than in other states, we expect that the differences between the typical transaction and one financed with a loan near the conforming limit will be less stark in California than in other parts of the country.

We limit our time frame to the period between 1997 and 2007 for several reasons. First, the LPS data that we use to estimate the jumbo-conforming spread are most comprehensive from the mid-1990s on. Second, we want to ensure that the conforming limit was being set in a consistent way across all years in the sample. Until 2007, a single conforming limit was set annually according to a formula and was imposed uniformly across all of the lower 48 states. However, after the GSEs were taken into government conservatorship in 2008, the standards for determining the conforming limit were changed in several ways, including a provision that allows it to vary across different metropolitan areas.

The final reason we avoid using post-2007 data is that there were significant changes to the structure of the mortgage market during the financial crisis that could potentially confound our analysis.

For example, the jumbo securitization market almost completely dried up during this period, which lead to a sharp reduction in the number of jumbo loans originated and a large rise in the jumbo-conforming spread (Fuster and Vickery, 2013). We limit our sample period to years before 2007 in order to avoid conflating the reduction in supply of jumbo loans during the housing bust with the demand-side response to the conforming limit that we are most interested in. Finally, we drop all loans originated from October through December, since banks may hold such loans in their portfolios until the conforming limit changes in January (Fuster and Vickery, 2013).[17] These restrictions leave us with a primary estimation sample of approximately 2.7 million transactions across 26 MSAs. Summary statistics for this sample and the sub-sample with first loan amounts within $50,000 of the conforming limit are reported in Appendix Table A.6.

2.3.2. LPS

The primary disadvantage of the DQ data set for studying mortgages is that it does not record interest rates and lacks important information on borrower characteristics, such as credit scores and debt-to-income ratios. Consequently, we turn to data from LPS to estimate the jumbo-conforming spread, as well as interest rates on second mortgages taken out at closing. The LPS data are at the loan level and run from 1997 to the present, covering approximately two-thirds of the residential mortgage market.[18] The data contain extensive information on mortgage terms and borrower characteristics, as well as geographic identifiers down to the zip code level. We focus on first mortgage originations for home purchases and apply the same set of restrictions described above for the DQ data, in particular the limitations to California and the first nine months of each year between 1997 and 2007. Descriptive statistics for the full LPS sample the sub-sample of loans with first loan amounts within $50,000 of the conforming limit are reported in Appendix Table A.7.

---

[17] We also drop extreme outliers in appraisal value or LTV ratio.

[18] Although data are available from earlier years, they are less comprehensive and the loans have higher average "seasoning," meaning that it takes longer after origination for them to appear in the data set (Fuster and Vickery, 2013). If loans that are quickly prepaid or foreclosed on never appear, seasoned data may be less representative of the universe of loans.

2.4. Empirical Methodology

2.4.1. Estimating the Behavioral Response to the Conforming Limit

In Section 2.2, we showed that the behavioral response to the conforming loan limit can be derived from estimates of the amount of bunching and the counterfactual mass at the limit. To estimate these quantities, we follow the approach taken by Kleven and Waseem (2013).

Since we are primarily interested in estimating the behavioral response in percentage terms, we first take logarithms of the loan amounts. We then center each loan in our data set at the (log) conforming limit in the year that the loan was originated. A value of zero thus represents a loan size exactly equal to the conforming limit while all other values represent (approximate) percentage deviations from the conforming limit. We group these normalized loan amounts into bins centered at the values m_j, with $j = -J, \ldots, L, \ldots, 0, \ldots, U, \ldots, J$, and count the number of loans in each bin, n_j. To obtain estimates of bunching and the counterfactual loan size distribution, we define an excluded region around the conforming limit, $[m_L, m_U]$, such that $m_L < 0 < m_U$ and fit the following regression to the count of loans in each bin

$$n_j = \sum_{i=0}^{p} \beta_i (m_j)^i + \sum_{k=L}^{U} \gamma_k \mathbb{1}\left(m_k = m_j\right) + \epsilon_j. \tag{2.10}$$

The first term on the right hand side is a p-th degree polynomial in loan size and the second term is a set of dummy variables for each bin in the excluded region. Our estimate of the counterfactual distribution is given by the predicted values of this regression omitting the effect of the dummies in the excluded region. That is, letting $\hat{n}_j$ denote the estimated counterfactual number of loans in bin j, we can write

$$\hat{n}_j = \sum_{i=0}^{p} \hat{\beta}_i (m_j)^i. \tag{2.11}$$

Bunching is then estimated as the difference between the observed and counterfactual bin counts

in the excluded region at and to the left of the conforming loan limit,

$$\hat{B} = \sum_{j=L}^{0} \left(n_j - \hat{n}_j \right) = \sum_{j=L}^{0} \hat{\gamma}_j, \tag{2.12}$$

while the amount of missing mass due to bunching is $\hat{M} = \sum_{j>0}^{U} \left(n_j - \hat{n}_j \right) = \sum_{j>0}^{U} \hat{\gamma}_j$.

The parameter of primary interest is $\Delta \hat{\bar{m}}$, the empirical analogue of $\Delta \bar{m}$ from equation (2.8). This parameter represents the average behavioral response of the marginal bunching individual measured as a percentage deviation from the conforming limit. Following the theory, we calculate it as

$$\Delta \hat{\bar{m}} = \frac{\hat{B}}{\hat{g}_0(\bar{m})}, \tag{2.13}$$

where $\hat{g}_0(\bar{m}) = \sum_{j=L}^{0} \left(\hat{n}_j \right) / \left| \frac{m_0 - m_L}{L} \right|$ is the estimated counterfactual density of loans in the excluded region at and to the left of the conforming loan limit. Intuitively, if the ratio of bunched to counterfactual loans is large, the existence of the limit has a large effect on the behavior determining the observed distribution of loan amounts. We calculate standard errors for all estimated parameters using a bootstrap procedure, as in Chetty et al. (2011).[19]

There are two key identifying assumptions necessary for equation (2.13) to provide a valid estimate of the behavioral response to the conforming limit. The first is that the counterfactual loan size distribution that would exist in the absence of the limit is smooth. That is, any spike in the loan size distribution at the conforming limit must be solely attributable to the existence of the limit and not some other factor. We test for violations of this assumption below by examining how the distribution of loan sizes changes when the conforming limit moves from one year to another. The second assumption is that households can be uniquely indexed by their counterfactual choice of mortgage size in the absence of the limit—that is, there is a well-defined marginal buncher. While this assumption is fundamentally untestable, most reasonable models of mortgage choice would

---

[19]At each iteration (k) of the bootstrap loop we draw with replacement from the estimated errors, $\hat{\epsilon}_j$, in equation (2.10) to generate a new set of bin counts, n_j^k. We then re-estimate bunching using these new counts. Our estimate of the standard error for $\Delta \hat{\bar{m}}$ is the standard deviation of the estimated $\Delta \hat{\bar{m}}^k$s. The same procedure produces standard errors for all the other bunching parameters that we report.

imply such a result.

In order to estimate the components of equation (2.13), there are several free parameters that we must choose: the bin width ($\left|\frac{m_0 - m_L}{L}\right|$), the order of the polynomial (p), and the location of the lower and upper limits of the excluded region (m_L and m_U). Following Kleven and Waseem (2013), we choose the upper limit to minimize the difference between bunching ($\hat{B}$) and missing mass to the right of the notch in the excluded region ($\hat{M}$).[20] For the other three parameters, our preferred specification uses 1-percent bins, a 13th-degree polynomial, and sets $m_L = 0.025$. We prefer this specification because, among the parameter configurations we considered, it yields the smallest difference between $\hat{B}$ and $\hat{M}$ in the sample that pools across all years and loan types.[21]

2.4.2. Estimating the Jumbo-Conforming Spread

In order to convert our estimate of borrowers' responsiveness to the conforming loan limit into an elasticity, we also need to estimate the magnitude of the change in rates that borrowers face. This exercise is complicated by the fact that there is a large class of borrowers who, as we demonstrate, bunch precisely at the conforming limit. These borrowers may have unobserved characteristics that are correlated with interest rates and that might bias an estimate of the jumbo-conforming spread based on a simple comparison of observed mortgage rates. However, this concern is not as grave as it may first appear. In particular, we are aided greatly by the fact that mortgage rates are typically determined based on a well defined set of borrower and loan characteristics that are all readily observable in the LPS data. To the extent that we are able to fully control for these characteristics, our estimates of the jumbo-conforming spread should be relatively close to the true interest rate differential facing the average borrower in our sample.

---

[20]This is done using the following iterative procedure: First, initialize m_U at a small amount (m_U^0) near the limit and estimate bunching ($\hat{B}^0$), missing mass ($\hat{M}^0$), and the difference between the two, ($\hat{B}^0 - \hat{M}^0$). Next, increase m_U by a small amount to m_U^1 and calculate the difference $\hat{B}^1 - \hat{M}^1$. We repeat this process until $\hat{B}^k - \hat{M}^k > \hat{B}^{k-1} - \hat{M}^{k-1}$, at which point we stop and take m_U^{k-1} to be the upper limit of the excluded region.

[21]It is worth noting, however, that the estimated missing mass from the right of the limit need not be exactly equal to the number of bunched loans since the procedure we use to estimate bunching ignores both extensive margin responses and the leftward shift of the distribution outside of the excluded region generated by intensive margin responses among those who do not bunch. If these types of responses have a sizable effect on the observed loan size distribution, then choosing parameters to minimize the difference between bunching and missing mass could lead to bias in the estimated behavioral response. To account for this, we explore robustness to various choices of the underlying parameters, which often yield estimates of $\hat{B}$ that are smaller than $\hat{M}$ but give very similar estimates of $\Delta\hat{m}$ to our preferred specification.

With this in mind, our main approach to estimating the jumbo-conforming spread follows that of Sherlund (2008), who exploits the sharp discontinuity at the conforming loan limit while also controlling semiparametrically for all other relevant determinants of interest rates. Of course, in a finite sample, it is not possible to control completely for all observed determinants of interest rates and there may be some unobserved characteristics which our controls are unable to capture. To account for this, we also estimate models which use a discontinuous function of the value of the home as an instrumental variable (IV) for jumbo loan status, as described in detail below.

Unlike Sherlund (2008), who uses an analogue to local linear regression, we incorporate the semi-parametrics in standard ordinary least squares regressions. We do this both to reduce the computational burden and to allow for a straightforward comparison with the IV estimates. In particular, we estimate variants of the following equation

$$
\begin{aligned}
r_{i,t} = {} & \alpha_{z(i),t} + \beta J_{i,t} + f^{J=0}\big(m_{i,t}\big) + f^{J=1}\big(m_{i,t}\big) + s^{LTV}\big(LTV_{i,t}\big) + s^{DTI}\big(DTI_{i,t}\big) \\
& + s^{FICO}\big(FICO_{i,t}\big) + PMI_{i,t} + PP_{i,t} + g\big(TERM_{i,t}\big) + \epsilon_{i,t},
\end{aligned}
\tag{2.14}
$$

where $r_{i,t}$ is the interest rate on loan i originated at time t, α is a zip-code by time fixed effect, and J is a dummy variable for whether the loan amount exceeds the conforming limit. In the spirit of a regression discontinuity design, we interact J with cubic polynomials in the size of the mortgage separately on either side of the conforming limit ($f^{J=0}\big(m_{i,t}\big)$ and $f^{J=1}\big(m_{i,t}\big)$) in order to control for any underlying continuous relationship between loan size and interest rates. In addition, we include splines in the loan-to-value ratio (LTV), debt-to-income ratio (DTI), and credit score ($FICO$) as well as fixed effects for whether the borrower took out private mortgage insurance (PMI) and if the mortgage had a prepayment penalty (PP). Finally, we also control flexibly for the length of the mortgage ($TERM$).[22] The coefficient of interest is β, which provides a valid estimate of the jumbo-conforming spread under the assumption that we have successfully controlled for borrower selection around the limit.

---

[22]The exact specifications are described in the results section below.

Instrumenting for Jumbo Loan Status

If there are other unobserved determinants of interest rates which are also correlated with jumbo loan status, then estimates of β based on equation (2.14) will produce biased estimates of the true jumbo-conforming spread. To gauge the extent to which this may be affecting our results, we also estimate a version of equation (2.14) in which we instrument for jumbo loan status using a discontinuous function of the value of the home, following Kaufman (2012).[23]

When making a loan, lenders typically require an independent appraisal as a check that the agreed transaction price accurately reflects the value of the home. The "value" in the denominator of the LTV ratio is then set as the lesser of the appraised value and the transaction price. Many buyers purchase a home with an LTV of exactly 80 percent, both because it is a longstanding norm and because exceeding 80 percent typically requires purchasing private mortgage insurance. Consequently, if the value of the home is just under the conforming loan limit multiplied by 1.25, then a buyer is substantially less likely to take out a jumbo loan than if the value is just over this limit. Figure 2.3A confirms this, showing that the fraction of loans that are jumbos jumps discontinuously as the value crosses 125 percent of the conforming limit.

This fact suggests an approach in which we instrument for $J_{i,t}$ in equation (2.14) with whether the value of the home falls above or below $1.25\,\bar{m}$, focusing in particular on a narrow range around the discontinuity. The key to the exogeneity of this instrument is that, compared with their actual loan amount, borrowers are less able to finely control the transaction price and have essentially no control over the exact outcome of their appraisal. As a result, some borrowers who do not want to or are not able to deviate from the 80 percent norm may be induced into or out of jumbo-loan status based on factors which are not entirely within their control. As supporting evidence for this assumption, Figure 2.3B shows that the distribution of house values around 125 percent of the conforming limit is quite smooth, in contrast with the distribution of loan amounts around the conforming limit itself, as in Figure 2.1B.[24]

---

[23] Adelino et al. (2012) and Fuster and Vickery (2013) employ similar strategies to look at the effects of the conforming limit on house prices and on mortgage supply, respectively.

[24] Kaufman (2012), who uses the same LPS data that we do, motivates this instrument only in terms of a borrower's

This IV approach is not a panacea, however. As Kaufman (2012) notes, it identifies a local average treatment effect among borrowers who choose to increase their first mortgage balance in order to keep their LTV constant despite their home value being just above 125 percent of the conforming limit. But in this paper, we are interested in estimating the average elasticity among the *entire* population of borrowers with counterfactual loan amounts above the limit. If there is heterogeneity in the jumbo-conforming spread, then those facing the lowest spread will be the most likely to take out a larger loan in response to a higher home value. Consequently, it is likely that the IV estimates provide a lower bound on the average spread in the population. Given the clear difficulty of estimating the "true" jumbo-conforming spread in the full population of borrowers, our preferred approach is to estimate the spread using both techniques and present a range of plausible elasticities.

2.5. Bunching and Jumbo-Conforming Spread Estimates

The next three sections present our primary empirical results. We begin in this section by presenting graphical evidence documenting bunching at the conforming loan limit as well as formal estimates of bunching and the behavioral response to the jumbo-conforming spread. We then present a series of estimates of the magnitude of the jumbo-conforming spread which we combine with the bunching estimates in Section 2.6 to calculate elasticities. In Section 2.6.3, we conclude with a discussion on the ways in which borrowers appear to be adjusting their loan sizes.

2.5.1. Bunching at the Conforming Limit

Results for all Borrowers

As a starting point for our empirical analysis, Figure 2.4 plots both the observed (log) loan size distribution and the counterfactual distribution estimated from the bunching procedure using all available loans in the DQ sample. Although our estimation is carried out in the full sample, in

---

inability to control the appraisal amount. His motivation is a bit problematic, because the variable that both he and we use is actually the minimum of the appraisal amount and the transaction price, as noted above. In principle, this variable could be contaminated if buyers can manipulate transaction prices around 125 percent of the conforming limit in response to differences in the jumbo-conforming spread, perhaps through negotiation with sellers. In practice, however, Figure 2.3B strongly suggests that there is no such selection and that the exclusion restriction is valid.

this (and subsequent) figures we have narrowed our focus to the range of loans which fall within 50 percent of the conforming limit. The x-axis shows the difference between the log loan amount and the log conforming limit in the year the loan was originated, so that zero is the limit itself and each bin represents roughly a 1 percent incremental deviation from the limit. The y-axis on the right indicates the number of loans in each bin, while the y-axis on the left indicates the fraction of all loans represented by that number.

The connected black line plots the histogram of (log) loan size, which exhibits a sharp peak at the limit. This bin contains approximately 100,000 loans, representing 4 percent of the entire sample, which is roughly four times as many loans as in the bin immediately to the left. The black line also shows a clear deficit of loans to the right of the limit, with the first bin containing only about half as many loans as the bin immediately to the left of the limit. The heavy dashed red line shows the fitted polynomial that we take as our counterfactual loan size distribution. The vertical dashed gray lines represent the lower (m_L) and upper (m_H) limits of the excluded region, as defined in Section 2.4.1.

The estimated number of loans bunching at the limit is reported in the figure and is calculated as the sum of the differences between the black and red lines in each bin in the excluded region at and to the left of the conforming limit. As the plot makes clear, bunching is remarkably sharp; almost all of the approximately 84,000 "extra" loans in this region are in the bin that contains the limit itself. Our estimate of $\Delta \hat{m}$, the behavioral response to the conforming limit, is also reported in the figure. It implies that the average marginal bunching borrower reduces his loan balance by roughly 3.8 percent.

The first column of Table 2.1 repeats these estimates along with their standard errors and several other parameters estimated during the bunching procedure. As another way of gauging the magnitude of the response, the third row of Table 2.1 reports a measure of the "excess mass" at the conforming limit. We calculate this as the ratio of the number of loans bunching at the limit to the number of loans which would have been there in its absence. The estimate implies that there are roughly 3.78 times more loans at the conforming limit than would have otherwise been expected.

All of these parameters are precisely estimated.

In the last row of Table 2.1, we also report the upper limit of the excluded region used in estimation (m_H). If there were no extensive margin responses (borrowers leaving the market entirely) then this number would provide an estimate of the largest percent reduction in mortgage size among bunching individuals. That is, no individual with a counterfactual loan size more than m_H percent larger than the conforming limit would be induced to bunch. Given extensive margin responses, it is possible that our estimate of m_H differs from the true cutoff value. Nonetheless, it provides a useful gauge of the magnitude of behavioral responses among those who reduce their mortgage sizes the most. The estimate implies an upper bound on behavioral responses of roughly 16 percent, meaning that nearly all of the borrowers bunching at the conforming limit would have had mortgages that were less than 16 percent larger than the limit had it not existed.

While the bunching evidence reported in Figure 2.4 is quite stark, it is interesting to note that there are still a number of borrowers who locate just above the conforming limit despite the large marginal interest rate associated with doing so. There are three potential explanations for this behavior: inelastic preferences, heterogeneous costs of adjusting first mortgage balances, or borrower-level differences in the magnitude of the jumbo-conforming spread that lead some borrowers to face no change in price at the conforming limit. In our analysis of the magnitude of the jumbo-conforming spread we were able to largely rule out the third explanation. We find no evidence of heterogeneity in the interest rate spread across borrower credit scores or the distribution of zip-code level median income, median house prices, or price-to-income ratios. This suggests that the borrowers who locate in the region just above the conforming limit are likely doing so as a result of either inelastic preferences or adjustment costs.

In the context of income taxes, Kleven and Waseem (2013) are able to distinguish between the role of preferences and adjustment costs using the fact that a notch can sometimes create a dominated region in which no wage earner, regardless of tax elasticity, would choose to locate in the absence of adjustment costs. By counting the number of wage earners observed in the dominated region they are able to back out an estimate of adjustment costs. Unfortunately, we cannot perform a similar

exercise here because there is no such dominated region in our setting. In the terminology of Kleven and Waseem (2013), the jumbo-conforming spread creates a "downward notch" where, for any finite loan amount above the limit, there exists a first mortgage demand elasticity sufficiently close to zero such that some borrower would be willing to take out that loan. It is therefore not possible to estimate the magnitude of any potential adjustment costs in this setting. Because of this, it is important to note that the elasticities we estimate are necessarily "reduced form," in the sense that they incorporate the effect of adjustment costs and are not driven entirely by the intertemporal elasticity of substitution alone.

Finally, as noted above, one of the key assumptions necessary for bunching to identify behavioral responses is that the counterfactual loan size distribution be smooth. While there is no way to directly test this assumption, one way to evaluate its plausibility is to examine what happens to the empirical distribution of loan sizes as the conforming limit moves from one year to the next. If the conforming limit is the only thing causing bunching, then bunching should track the movement of the conforming limit, and the distribution should be smooth at previous and future conforming limits. Figure 2.5 plots the empirical loan size distribution separately for the years 2000, 2002, and 2004.[25] At any given nominal loan amount, the distribution appears to be smooth except in the year for which that loan amount serves as the conforming limit. While not definitive, this result strongly supports the counterfactual smoothness assumption.

Fixed versus Adjustable Rate Mortgages

In addition to looking at the effect of the conforming limit on overall loan size, recent work by both Fuster and Vickery (2013) and Kaufman (2012) draws attention to several stylized facts that make it particularly interesting to investigate heterogeneity in the response by *type* of loan.[26] In particular, these authors document a sizable and sharp decline in the share of fixed-rate mortgages (FRMs) relative to adjustable-rate mortgages (ARMs) precisely at the conforming limit. We repli-

---

[25]We avoid plotting consecutive years to preserve the clarity of the picture, but the results hold at any horizon.

[26]In the appendix, we also examine heterogeneity by classifying borrowers by race and income. We find that low-income or minority borrowers are substantially less prone to bunching at the limit than high-income or non-minority borrowers.

cate this stylized fact in Figure 2.6 using our own sample of loans from DataQuick. Using the same 1 percent bins as before, this figure plots the share of loans that are FRMs as a function of loan size relative to the conforming limit. To the left of the limit, the FRM share declines gradually as the loan amount increases, reaching roughly 55 percent just below the limit. It then spikes to about 75 percent at the limit before falling to 20 percent immediately to the right. Beyond the conforming limit, the share then rises, eventually reaching a plateau of about 35 percent.

This drop in the FRM share is not a coincidence. Fixed-rate mortgages are generally estimated to have a larger jumbo-conforming spread relative to ARMs due to the fact that their returns are much more vulnerable to interest rate risk.[27] Since the FRM share well to the right of the limit is substantially lower the FRM share to the left of the limit, a quick glance at Figure 2.6 might suggest an extensive margin response. That is, in response to the higher jumbo spread for FRMs, some jumbo borrowers may choose to substitute toward ARMs.

In contrast, we argue that the change in the FRM share at the conforming limit occurs because more FRM borrowers than ARM borrowers bunch at the limit, without switching loan type. To show this, we separately estimate bunching for both fixed-rate and adjustable-rate mortgages. If the drop in FRM share at the limit is driven primarily by borrowers switching to ARMs, then we should expect to see both a downward shift in the observed distribution of FRMs relative to its counterfactual immediately to the right of the limit and a concomitant *upward* shift in the ARM distribution.

Figure 2.7A and Figure 2.7B show the results from this exercise for FRMs and ARMs, respectively. The standard errors and additional bunching parameters are also reported separately for each loan type in columns 2 and 3 of Table 2.1. While Figure 2.7A shows a substantial downward shift in the FRM distribution to the right of the limit, Figure 2.7B shows no corresponding upward shift in the ARM distribution. In fact, much like in the figures for the FRM and combined samples, the ARM

---

[27]We replicate this well-documented difference in spreads using our own sample of loans from LPS in Section 2.5.2. Since jumbo loans are harder to unload onto the secondary mortgage market, originators will demand a higher interest rate on jumbo FRMs relative to jumbo ARMs in order to compensate them for the additional risk they bear by having to hold the loans in portfolio.

distribution features a region of missing mass immediately to the right of the limit. Moreover, in our preferred specification, the missing mass for each type of loan is roughly equal to the mass of that type of loan bunching at the limit.[28]

While we do not believe that our results invalidate any of the conclusions drawn by Fuster and Vickery (2013) or Kaufman (2012), they do illuminate the fact that perhaps the most intuitive channel for the drop in the FRM share above the limit—substitution between FRMs and ARMs—is not the correct one. With this in mind, for the remainder of the paper we will present estimates for FRMs and ARMs separately.

2.5.2. Jumbo-Conforming Spread

To convert the behavioral responses estimated from bunching into elasticities, we next need to obtain an estimate of the interest rate differential at the limit. Table 2.2 presents estimates of the jumbo-conforming spread, following the strategies discussed in Section 2.4.2. We estimate the spread using OLS and IV for fixed- and adjustable-rate mortgages separately, with four different specifications each. All of the specifications include controls for the distance to the conforming limit (linear, quadratic, and cubic terms) interacted with the jumbo loan indicator variable, as well as controls for the loan-to-value (LTV) ratio, debt-to-income (DTI) ratio, missing DTI ratio, FICO credit score, missing FICO score, whether the loan includes private mortgage insurance (PMI), and whether the loan has a prepayment penalty. They also include zip-code by month fixed effects and fixed effects for standard loan lengths, such as 15, 30, and 40 years, as well as a linear term to capture the effects of nonstandard lengths.

Across the columns of the table, the four specifications are: (1) a baseline, using all available data, with linear controls for the LTV and DTI ratios and the FICO score; (2) the same specification replacing the linear controls for LTV, DTI and FICO with more flexible cubic B-splines; (3) the same specification as in (2) but with a sample limited to loans within $50,000 of the conforming

---

[28]Of course, since we only observe average responses, it is still possible that some borrowers choose ARMs over FRMs because of the limit, particularly if there is heterogeneity in the costs of ARMs and FRMs within the population. But since the figures do not suggest any noticeable *aggregate* response, an offsetting group of borrowers would have to be choosing FRMs over ARMs because of the limit.

limit; and (4) the same specification in (2) but with a sample limited to loans within $10,000 of the limit.

For FRMs, applying least squares yields estimates of the jumbo-conforming spread that are tightly clustered around 17 to 18 basis points and precisely estimated, regardless of the specification. These estimates are similar to Sherlund's (2008) estimate of 22 basis points, despite our use of a simpler estimation technique and a different data set covering a smaller geographic area and a shorter time horizon.[29]

As discussed above, the OLS specifications do not control for borrower selection on unobservables around the conforming limit, which is a particular concern given the substantial bunching that we have just highlighted. To address this potential selection issue, the row labeled "IV" presents estimates of the spread when we instrument for the jumbo indicator with an indicator for whether the value of the house exceeds 125 percent of the the conforming limit. As would be expected, these estimates are somewhat less precise than the OLS results but are still significantly different than zero at conventional levels. The estimates run from 10 to 13 basis points and are uniformly smaller than the OLS results, possibly reflecting either borrower selection or the fact that the IV approach estimates a local average treatment effect among a population of borrowers who may face a lower spread.[30] Reassuringly, these estimates are quite similar to Kaufman's (2012) estimate of 10 basis points, using essentially the same technique but a different sample of loans.[31]

Our estimates of the jumbo-conforming spread for initial ARM rates, presented in the lower half of Table 2.2, are considerably more noisy and merit further study. The OLS estimates are uniformly negative—that is, jumbo loans have *lower* rates than conforming loans—and the negative coefficient gets bigger as the sample becomes more focused on loans near the limit. It is quite unlikely that this is an accurate representation of pricing of ARMs above and below the conforming limit. Indeed,

---

[29]Specifically, Sherlund (2008) uses data from the Monthly Interest Rate Survey that cover the entire U.S. from 1993 to 2007. We use LPS data from California only from 1997 to 2007.

[30]Regressing the jumbo indicator on the indicator for whether the value exceeds 125 percent of the conforming limit, including all of the other covariates as controls—that is, the first stage of the IV—yields a coefficient of around 0.2, with slight variation depending on the specification. The instrument is quite strong; in all cases, the standard errors on this coefficient are tiny and the F-statistics very large.

[31]Kaufman (2012) uses LPS data that cover the entire U.S. from 2003 to 2007.

we saw in the previous section that some ARM borrowers are bunching at the limit, albeit fewer than in the FRM market. Some aspect of the loans must be leading those borrowers to prefer loans at the limit to loans above it. One possibility is that the negative coefficients result from selection. The "IV" row uses the same instrument as in the FRM results. Here we see coefficients that range from about zero, in the first two columns, to 7 and 4 basis points, in columns 3 and 4. The standard errors are too large to rule out a negative spread, although the point estimates are comparable to other estimates in the literature that are also close to zero.[32]

It is somewhat surprising that addressing borrower selection around the limit using the IV specification produces more positive coefficients in the ARM case and less positive coefficients in the FRM case. Given the observed bunching at the limit among ARM borrowers, it is also surprising to find no strong support for a positive jumbo-conforming spread among ARM loans. While we control for many relevant aspects of mortgage contracts, such as prepayment penalties and private mortgage insurance, it is possible that there are other unobserved factors that are more relevant for ARMs than for FRMs and which could be biasing these results. Consequently, the remaining discussion in this paper, including the calculation of elasticities, focuses on the results for FRMs.

2.6. Elasticities

2.6.1. Calculating Elasticities with a Notched Budget Constraint

As discussed above, the higher mortgage rate for loans above the conforming limit creates a "notch" in which the *average* price jumps discontinuously, rather than a "kink" in which the *marginal* price changes discontinuously but the average price is continuous. That is, borrowers must pay the higher interest rate on the entire balance of the loan, not just the balance in excess of the limit. As a result, it is not appropriate to calculate an elasticity using our estimate of the jumbo-conforming spread as the denominator. Instead, we follow an approach similar to the "reduced-form approximation" suggested by Kleven and Waseem (2013). The idea is to construct a measure

---

[32]Kaufman (2012) uses the IV approach to estimate an effect of jumbo status on initial ARM rates of about negative 5 basis points. Fuster and Vickery (2013) use data from surveys of loan officers, which should in principle hold borrower characteristics constant, and find effects on ARM rates of 0 to 10 basis points between mid-2006 and mid-2007.

of the implicit marginal cost facing the marginal bunching borrower as a result of the conforming limit. We can measure this cost either in terms of the monthly payment or the interest rate itself.

Focusing first on the monthly payment, let $P(m, r)$ denote the monthly payment on a 30-year mortgage of amount m at fixed interest rate r.[33] The total increase in payment generated by the conforming limit for a loan of size $m > \bar{m}$, is equal to

$$P(m, \hat{r} + \Delta\hat{r}) - P(\bar{m}, \hat{r}), \tag{2.15}$$

where $\bar{m}$ is the conforming limit, $\hat{r}$ is the estimated conforming rate, and $\Delta\hat{r}$ is the estimated jump in the interest rate at the limit (the jumbo-conforming spread).

The marginal increase in payment per dollar of the loan, averaged over the distance from the limit to m, is then

$$P^*(m) = \frac{P(m, \hat{r} + \Delta\hat{r}) - P(\bar{m}, \hat{r})}{m - \bar{m}}. \tag{2.16}$$

Letting $P'(\bar{m}) = \frac{P(\bar{m}, \hat{r})}{\bar{m}}$ be the (constant) marginal payment per dollar of the loan for loans made at the conforming interest rate $\hat{r}$, we can then define

$$\epsilon^P = \frac{\Delta\hat{m}}{\log\left(P^*\left(\bar{m} + \Delta\hat{m}\right)\right) - \log\left(P'(\bar{m})\right)} \tag{2.17}$$

as the elasticity of mortgage demand with respect to the implicit marginal increase in monthly payment implied by our estimate of the behavioral response, $\Delta\hat{m}$.

While the monthly payment is an intuitive measure of the cost of a mortgage, the correct theoretical price per dollar of the loan is the underlying interest rate itself. For $m > \bar{m}$, define $r^*(m)$ such that

$$(m - \bar{m}) \cdot r^*(m) = m \cdot (\hat{r} + \Delta\hat{r}) - \bar{m} \cdot \hat{r}. \tag{2.18}$$

---

[33]We use a standard formula to compute the monthly payment: $P(m, r) = m \frac{\frac{r}{12}\left(\frac{r}{12} + 1\right)^{360}}{\left(\frac{r}{12} + 1\right)^{360} - 1}$.

This $r^*(m)$ is the implicit interest rate on the loan amount in excess of the conforming limit $(m - \bar{m})$, taking into account the jump in the overall rate. Solving explicitly for $r^*(m)$ yields

$$r^*(m) = \hat{r} + \Delta\hat{r} + \Delta\hat{r} \cdot \frac{\bar{m}}{m - \bar{m}}. \tag{2.19}$$

Equation (2.19) makes clear that $r^*(m)$ is equal to the jumbo rate $(\hat{r} + \Delta\hat{r})$ plus a term that is increasing in the jumbo-conforming spread $(\Delta\hat{r})$ and decreasing in the size of the loan relative to the conforming limit $(m - \bar{m})$.[34] For loans just above the limit this additional term is very large, reflecting the fact that the higher interest rate on jumbo loans is applied to the full balance of the loan.[35]

As with the monthly payment, given an estimate of $r^*(m)$, we can then calculate the (semi-)elasticity of (first) mortgage demand implied by our estimate of $\Delta\hat{m}$,

$$\epsilon_s^r = \frac{\Delta\hat{m}}{r^*\left(\bar{m} + \Delta\hat{m}\right) - \hat{r}}. \tag{2.20}$$

As before, $\Delta\hat{m}$ is estimated in logs, so it represents the approximate percentage change in mortgage demand induced by the conforming limit, while the denominator measures the level change in interest rates.[36]

The expressions given by equations (2.17) and (2.20) represent two different measures of the magnitude of borrower responses to interest rates. The two measures are not equivalent, both because of the nonlinearity of the denominators and because the monthly payment is a nonlinear function of the interest rate. Both measures are potentially useful. While the rate itself is the correct theoretical cost of borrowing, it is likely that many borrowers use the monthly payment to compare

---

[34]For a loan that is $\Delta\hat{m}$ above the limit in approximate percentage terms, equation (2.19) simplifies to $r^*\left(\bar{m} + \Delta\hat{m}\right) = \hat{r} + \Delta\hat{r} + \frac{\Delta\hat{r}}{\Delta\hat{m}}$.

[35]For example, in 2006 the conforming limit was \$417,000 and average interest rate on loans made just below that was 6.38 percent. Given our OLS estimate of the jumbo-conforming spread of roughly 17 basis points, this implies an average rate of 6.55 percent for a loan made just above the limit. Thus, the marginal interest rate on the last \$1,000 of a \$418,000 loan originated in 2006 was approximately $r^*(418,000) = 6.55 + 0.20 \times 417 = 89.95$ percent!

[36]We present our interest rate estimates as semi-elasticities because it is a bit more intuitive to consider changes in interest rates in basis or percentage points.

different loans or to compare owning to the monthly cost of renting (Attanasio et al., 2008).

2.6.2. *Estimated Elasticities of First Mortgage Demand*

The first two columns of Table 2.3 report the semi-elasticities we calculate for a range of estimates of $\Delta\hat{\hat{m}}$ and the jumbo-conforming spread, $\Delta\hat{r}$. The semi-elasticities and associated standard errors, calculated using the delta method, are shown in the lower-right portion of the table. Each semi-elasticity is calculated from the estimate of $\Delta\hat{r}$ reported at the top of that column and the estimate of $\Delta\hat{\hat{m}}$ at the beginning of that row.

Our preferred estimate of bunching for FRMs from Table 2.1 (0.063) is shown in the middle row. The other two estimates (0.052 and 0.083) are the smallest and largest estimates of $\Delta\hat{\hat{m}}$ across a range of different options for the three parameters chosen *ex ante*: the bin width, the polynomial order, and the lower limit of the excluded region.[37] They provide reasonable bounds on the variation in the elasticity implied by these parameters. The jumbo-conforming spread estimates are taken from column 3 of Table 2.2 and correspond to the OLS (column 1) and IV (column 2) estimates, respectively.

The estimated semi-elasticities range from -0.015 to -0.053, with our preferred estimates in the middle row at -0.022 and -0.031. The associated standard errors are relatively small, although those using the noisier IV estimate of the jumbo-conforming spread are larger than those using the OLS estimate. The semi-elasticities can be interpreted as the percentage change in the balance of a first mortgage demanded in response to a 1 basis point increase in the interest rate. As an example, our preferred estimates imply that an increase in the mortgage rate from 5 percent to 6 percent—100 basis points—would lead to a decline in first mortgage demand of 2 to 3 percent.

Columns 3 and 4 report the analogous elasticities with respect to the marginal monthly payment for the same set of estimates of $\Delta\hat{\hat{m}}$ and equivalent estimates of the jumbo-confirming spread, estimated in logs.[38] The preferred estimates in the middle row are -0.27 using the OLS estimate of

---

[37] We considered bin widths of 0.01, 0.025 and 0.05; polynomials of order 7, 9, 11, and 13; and lower limits of 0.025, 0.05, 0.075 and 0.1.

[38] The log jumbo-conforming spread estimates are provided in Appendix Table A.9.

the spread and -0.35 using the IV estimate, indicating that a one percent increase in payment leads to about a third of a percent decline in mortgage demand.

2.6.3. Accounting for Second Mortgages

The elasticities reported in Table 2.3 tell us how much borrowers reduce their first mortgage balance in response to the jumbo-conforming spread, but they do not tell us how the borrowers adjust. There are three primary channels through which a borrower can reduce the size of her first mortgage, each of which have different implications for the interpretation of our main results. First, a borrower could simply bring more cash to the table, making a larger down payment and taking out a smaller loan.[39] Second, she could take out an additional mortgage for the amount of debt desired in excess of the conforming limit. Finally, she could spend less on housing, which (holding leverage constant) would lead her to take out a smaller mortgage.[40]

To measure the extent to which borrowers are using second mortgages to lower their first mortgage balance, Figure 2.8 plots the number of transactions financed using second loans as a function of the associated fixed-rate first mortgage value relative to the conforming limit. Consistent with the notion that many of the bunching borrowers take out second mortgages, there is a sharp spike in the number of transactions which are financed with two loans precisely at the limit. The plot suggests that roughly 25,000 more second loans were taken out in the bin at the conforming limit relative to the bin just below it, which provides a reasonable counterfactual. Combining this with our estimate from Table 2.1 that about 70,000 FRM borrowers bunch at the limit suggests that roughly 35 percent of FRM borrowers who bunch do so by taking out a second mortgage. These borrowers are presumably shifting debt from their first mortgage onto their second, holding combined LTV roughly constant while reducing their first-mortgage LTV.

The remaining 65 percent of "excess" borrowers must be either putting up more cash, or spending less on housing than they otherwise would. In the appendix, we attempt to further distin-

---

[39]"Putting up more cash" could be accomplished in many ways, including taking money out of savings, reducing current consumption, or taking out non-mortgage debt. We do not observe any data that would allow us to distinguish between these cases.

[40]This could happen either through substitution to a lower-quality home or through direct price negotiation with the seller.

guish between these two margins using back-of-the-envelop calculations involving the observed first mortgage LTVs and combined LTVs at the conforming limit. These calculation suggest that the 65 percent who do not take out second mortgages are likely putting up cash rather than buying cheaper houses. Regardless of how they adjust, however, the roughly two-thirds of borrowers who bunch at the conforming limit without the use of a second mortgage must be reducing their *total* mortgage debt.

Our bunching measure provides estimates of the mortgage demand elasticity for the marginal bunching borrower, while the estimate of second mortgage use is for the average bunching borrower. Since the fraction of all bunching borrowers who take out a second mortgage is not necessarily the same as the fraction of marginal bunching borrowers who do so, we cannot definitively solve for the response of total mortgage debt among the latter group. If we assume, however, that the use of second mortgages is the same among the average and marginal bunchers, then we can scale down our elasticity estimates by a factor of one-third to provide a rough estimate of the effect of a change in rates on total mortgage debt. Specifically, multiplying our preferred first mortgage semi-elasticity estimates of -0.022 and -0.031 by two thirds yields total debt semi-elasticities of about -0.015 and -0.021. That is, a one percentage point increase in rates should reduce total mortgage debt by between 1.5 and 2 percent.

2.6.4. Economic Magnitudes and Interpretation

Are these estimates large or small? In an absolute sense, they appear to be quite small. For example, during our sample period the average interest rate on a loan taken out just below the conforming limit was approximately 6.5 percent.[41] Our semi-elasticity estimates imply that a one percentage point increase from that baseline (a 15 percent increase), leads to a reduction in total mortgage debt of only 1.5–2 percent. This corresponds to an elasticity of only about 0.1–0.13.

Our finding of a relatively small intensive-margin behavioral response to changes in interest rates

---

[41]The mean interest rate across all fixed-rate mortgages falling within $10,000 below the conforming limit in the LPS data was 6.53 percent. Similarly, when we predict interest rates for everyone in our sample if they were to have taken out a loan just below the conforming limit using the regression estimates we use to estimate the jumbo-conforming spread, the mean predicted rate is approximately 6.6 percent.

is inline with several other findings from the literature. For example, Fuster and Zafar (2014) survey households about their mortgage choices under randomized hypothetical interest rate scenarios and find that their chosen downpayment fractions respond in a way that would imply even smaller intensive-margin semi-elasticities ranging between 0.6 and 1.8. Several studies of behavioral responses to interest rate changes in other consumer credit markets have found similarly small effects. For example, Attanasio et al. (2008) and Karlan and Zinman (2008) document small responses to interest rate changes for auto loans and microcredit, respectively.

These small responses could be explained by the presence of other binding constraints that make households less responsive to interest rates than they otherwise might be. For example, since both mortgages and auto loans are tied to the purchase of a large durable and frequently require a down payment, one explanation for the small responses we document here and those documented for auto loans could be the shadow cost of equity financing (i.e. downpayment constraints). Another possible explanation is the time frame over which households are able to adjust. Our results only apply to the relatively short-term decision of how much debt to incur at origination. If households are able to re-finance out of the higher jumbo rates after origination then the implied long-run behavioral response to the jumbo-conforming spread may be larger. This would be consistent with results from the literature suggesting that demand responses to interest rate changes are larger when measured over longer horizons (Karlan and Zinman, 2014) or when there are more available substitutes (Gross and Souleles, 2002).

2.7. Policy Application: GSE Guarantee Fee Increases

As noted in the introduction, the magnitude of borrower responses to changes in mortgage interest rates has important implications in several domains of economic policy. In this section, we illustrate how our elasticity estimates could be used to gauge the potential effects of one recently proposed policy change.

In February 2012, the Federal Housing Finance Agency (FHFA) published its "Strategic Plan for Enterprise Conservatorships," outlining the steps that the agency plans to take to fulfill its legal

obligations as conservator for the GSEs. As part of this plan, the FHFA established a goal of gradually reducing the dominant role that the GSEs currently play in the mortgage market. One of the primary proposed mechanisms for achieving this goal is a series of increases in the GSEs' guarantee fee (or "g-fee") up to the level that "one might expect to see if mortgage credit risk was borne solely by private capital" (FHFA, 2012).

The g-fee is the amount that Fannie Mae and Freddie Mac charge mortgage lenders in order to cover the costs associated with meeting credit obligations to investors in GSE mortgage-backed securities (MBS). It is typically collected in two components: an upfront fee assessed as a fraction of the balance of the loan at origination and a recurring annual fee equal to a fraction of the outstanding principal balance remaining at the end of each year.[42] The fee has risen several times in recent years, both by Congressional mandate and as part of the first steps in implementing FHFA's strategic plan.

In December 2013 the FHFA announced plans to increase the recurring fee by an additional 10 basis points for all loans and to reduce the up-front fee by 25 basis points for loans originated in all but four states.[43] The FHFA estimated that the combined effect of these increases and decreases would generate an overall average increase in the effective annual g-fee of roughly 11 basis points (FHFA, 2013).[44] A report by the FHFA Inspector General noted that "Significant guarantee fee increases, under some scenarios, could result in higher mortgage borrowing costs and dampen both consumer demand for housing and private sector interest in mortgage credit risk."(FHFA OIG, 2013).[45]

Our estimates of the interest rate elasticity of mortgage demand can be used to gauge the potential

---

[42]In 2012, g-fees from single-family mortgages generated roughly $12.5 billion in revenue for the GSEs, up 12% from the previous year (FHFA OIG, 2013).

[43]The four states for which the up-front fee would not be reduced are Connecticut, Florida, New Jersey, and New York. The fee would remain higher in these states to compensate for the greater costs associated with lengthy foreclosure timelines there.

[44]FHFA reports this 11 basis point number as the combined overall effect of an approximate 14 basis point increase in the g-fee for 30-year mortgages and a 4 basis point increase for 15-year mortgages, on average across all Freddie and Fannie loans.

[45]The future of this plan remains uncertain, particularly after it was delayed by incoming director Mel Watt in January 2014.

magnitude of any reductions in mortgage borrowing resulting from the proposed g-fee increases. To carry out this calculation we make three simplifying assumptions. First, we assume that the full 11 basis point increase in the g-fee would be passed through to borrowers in the form of higher interest rates on conforming mortgages. Second, we assume that this increase in rates for conforming mortgages would not have any general equilibrium effects on interest rates for non-conforming loans. Finally, we assume that our estimates of the total mortgage demand (semi)-elasticity of 1.5 to 2 apply at all points of the mortgage size distribution, not just at the conforming limit. Under these assumptions, our estimates imply that the proposed increase in the g-fee would reduce the total dollar volume of fixed-rate conforming mortgage originations by roughly 0.17 to 0.22 percent relative to what it otherwise would have been.

Under the same set of assumptions, we can also provide an estimate of the cumulative effects of g-fee increases to date. Between 2006 and the first quarter of 2013, the g-fee rose from approximately 20 basis points to 50 basis points, with much of the rise occurring in two waves in 2012 (FHFA OIG, 2013). Multiplying the 30 basis point differential by our elasticity implies a reduction in the dollar volume of fixed-rate conforming mortgage originations of 0.45 to 0.60 percent.[46]

2.8. Conclusion

In this paper, we use techniques for estimating behavioral responses from bunching at budget constraint nonlinearities in order to estimate the effects of the conforming loan limit on first mortgage demand. We combine these estimates with estimates of the jumbo-conforming spread to calculate the interest rate (semi-)elasticity of mortgage demand. Our estimates imply that size of a borrowers first mortgage falls by between 2 and 3 percent in response to a 1 percentage point increase in the interest rate. Accounting for the third of bunching borrowers who take out second mortgages suggests that *total* mortgage demand falls by between 1.5 and 2 percent. Applying these elasticity estimates to recently proposed increases in GSE guarantee fees implies a reduction in fixed-rate

---

[46]While our elasticity is well suited to examining the direct intensive margin effects of these relatively small changes in fees, we emphasize that we cannot draw any inference on more general questions, such as how high the fees would need to rise to draw private capital back into this part of the MBS market, or what the effects of the fee increases would be on extensive margin responses.

conforming mortgage originations of approximately one-fifth of one percent. The implied cumulative effect of similar increases that have occurred over the past several years has been to reduce originations by approximately one-half of one percent.

We conclude by pointing to three potentially useful avenues for future research. First, our estimates are necessarily limited by their context. A large number of salient factors, especially the presence of adjustment costs and the availability of second mortgages, affect how borrowers respond to the limit and, in turn, our estimates of the demand elasticity. A better understanding of the importance of these factors for our estimates is required before they can be applied to more general policy questions. Second, our estimates abstract away from the potential effect of interest rates on the extensive margin choice of whether to purchase a home or not. A full accounting of the effect of interest rates on mortgage demand would need to incorporate this margin as well. Finally, although we show in the appendix that there seem to be differential responses of minority versus non-minority and high-income versus low-income borrowers, data limitations prevent us from being able to fully investigate this heterogeneity. Painting a fuller picture of heterogeneity in elasticities and adjustment costs may be as important as pinning down the overall average elasticity of demand.

Panel A. Mean Interest Rate by Loan Size (2006)   *Panel B. Loan Size Density*
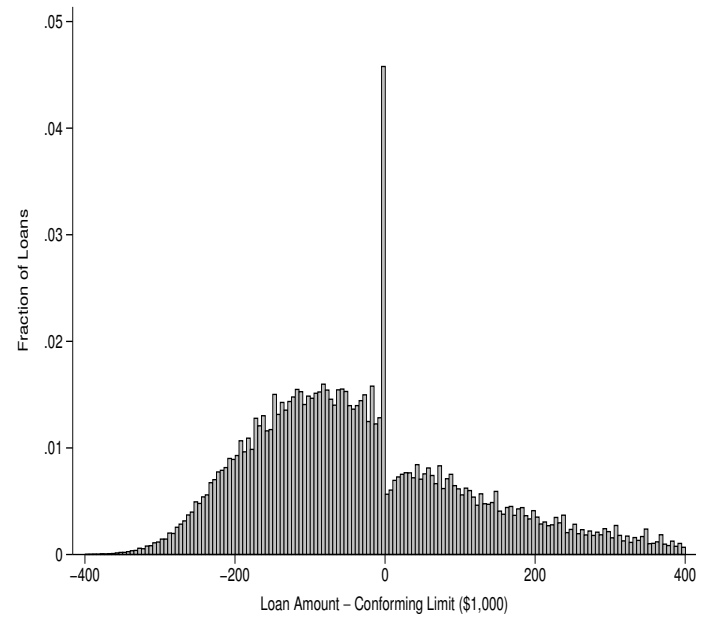
F<small>IG</small>. 2.1.—Panel (a) plots the mean interest rate as a function of the loan amount relative to the conforming limit for fixed-rate mortgages originated in 2006. Each dot represents the mean interest rate within a given $5,000 bin relative to the limit. The dashed red lines are predicted values from a regression fit to the binned data, allowing for changes in the slope and intercept at the conforming limit. Sample includes all loans in the LPS fixed-rate sample that fall within $100,000 of the conforming limit. Panel (b) plots the fraction of all loans that are in any given $5,000 bin relative to the conforming limit. Data are pooled across years and the sample includes all transactions in the primary DataQuick sample that fall within $400,000 of the conforming limit. In both panels each loan is centered at the conforming limit in effect at the date of origination so that a value of zero represents a loan at exactly the conforming limit.

Panel A. Budget Constraint Diagram

Panel B. Mortgage Size Density

FIG. 2.2.—Behavioral Response to Conforming Loan Limit. This figure shows the effect of the conforming loan limit (CLL) on individual behavior (Panel (a)) and the aggregate loan size density (Panel (b)). The CLL generates a notch in the intertemporal budget constraint, which is characterized by a jump at the limit equal to the jumbo-conforming spread times the CLL ($\Delta r \bar{m}$) and a change in slope to the right of the limit (Δr). The notch leads all borrowers with counterfactual loan sizes between $\bar{m}$ and $\bar{m} + \Delta \bar{m}$ to bunch at the limit. This behavior generates a discontinuity in the loan size density at the conforming limit, characterized by both a spike in the density of loans at the limit and a region of missing mass immediately to the right. The width of the region of missing mass is determined by the shape of the indifference curve of the marginal bunching individual, who is indifferent between locating at the CLL and the best interior point to the right of the limit (m^I).

Panel A. Jumbo Fraction

Panel B. Home Value Density

FIG. 2.3.—Jumbo Loan Fraction and Home Value Density Around 125 Percent of the Conforming Loan Limit. Panel (a) plots the fraction of loans in each bin with loan amounts greater than the conforming limit (jumbo loans). Panel (b) plots the fraction of all loans with values that are in any given $5,000 bin relative to 125 percent of the conforming limit. Value is measured as the minimum of the appraisal amount and the contract price. Data are pooled across years and each loan is centered at 125 percent of the conforming limit in effect at the date of origination, so that a value of zero represents a loan at exactly 125 percent of the conforming limit. Sample includes all transactions in the primary LPS sample that fall within $400,000 of the conforming limit. See text for details on sample construction.
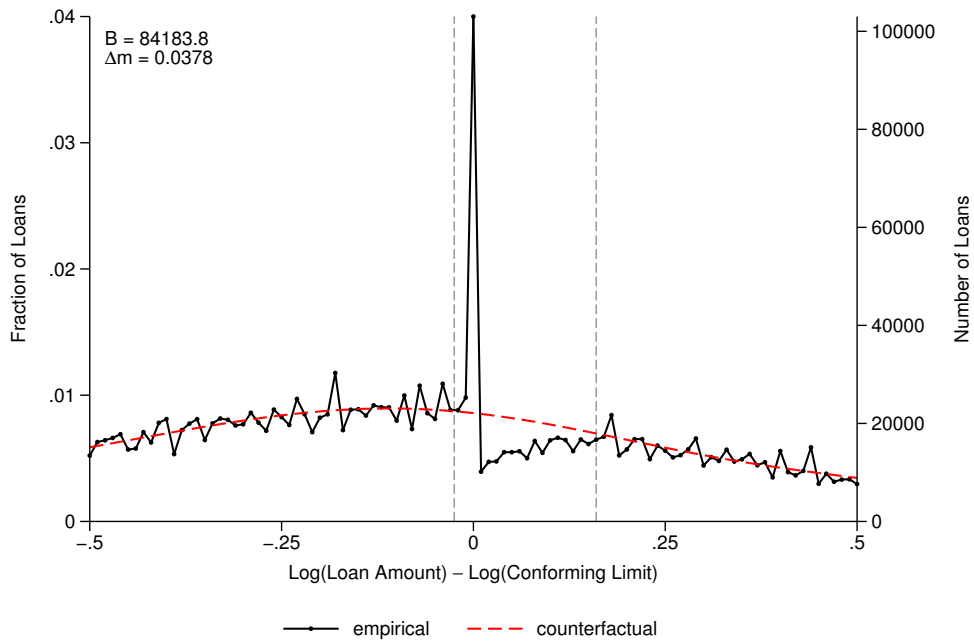
FIG. 2.4.—Bunching at the Conforming Limit, All Loans. This figure plots the empirical and counterfactual density of (log) loan size relative to the conforming limit for all loans. Estimation was carried out in the full sample of DQ loans, but the figure shows only loans within 50 percent of the conforming limit. The connected black line is the empirical density. Each dot represents the count (fraction) of loans in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed red line is the estimated counterfactual density obtained by fitting a 13th degree polynomial to the bin counts, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The figure also reports the estimated number of loans bunching at the limit (B) and the average behavioral response among marginal bunching individuals (Δm), calculated as described in Section 2.4.1.

FIG. 2.5.—Bunching Adjusting Over Time. This figure plots the empirical density of loan sizes separately for three different years. Each dot represents the fraction of loans originated in a given $10,000 bin in the indicated year. The vertically dashed lines mark the conforming limit that was in place in the indicated year. All dollar amounts are nominal.
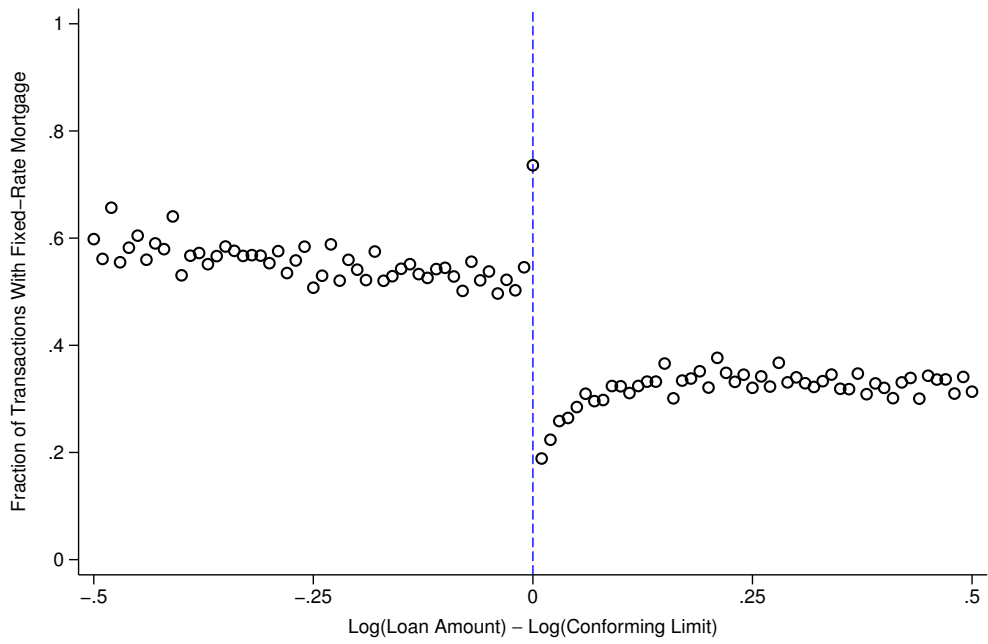
FIG. 2.6.—Fixed-Rate Mortgage Share Relative to the Conforming Limit. This figure plots the share of all transactions with a fixed-rate mortgage (FRM) as a function of loan size relative to the conforming limit. Each dot represents the fraction of FRMs in a given 1-percent bin relative to the limit in effect at the time of origination. FRM shares are calculated using the first mortgage associated with the transaction.

Panel A. Fixed-Rate Mortgages   *Panel B. Adjustable-Rate Mortgages*

FIG. 2.7.—Bunching at the Conforming Limit, FRMs and ARMs. These two panels plot the empirical and counterfactual densities of (log) loan size relative to the conforming limit, for FRMs and ARMs separately. Estimation was carried out in the full sample of DQ loans with fixed- and adjustable interest rates, respectively, but the figure shows only loans within 50 percent of the conforming limit. The connected black lines are the empirical densities. Each dot represents the count (fraction) of loans in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed red lines are the estimated counterfactual densities obtained by fitting 13th degree polynomials to the bin counts, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The panels also report the estimated number of loans bunching at the limit (B) and the average behavioral response among marginal bunching individuals (Δm), calculated as described in Section 2.4.1.

FIG. 2.8.—Number of Second Mortgages by First Mortgage Amount. This figure plots the number of transactions financed using two loans as a function of the first loan amount relative to the conforming limit. Each dot represents the number of transactions in a given 1-percent bin relative to the limit in effect at the time of origination. Sample includes only transactions with a fixed-rate first mortgage.

TABLE 2.1

BUNCHING ESTIMATES BY LOAN TYPE

	(1) Combined	(2) FRMs	(3) ARMs
Bunched Loans ($\hat{B}$)	84183.8	68306.3	16174.3
	(2687.0)	(1561.6)	(1446.7)
Behavioral Response ($\Delta\hat{m}$)	0.0378	0.0627	0.0144
	(0.0018)	(0.0025)	(0.0014)
Excess Mass ($\hat{B}/\sum_{j=L}^{0}\hat{n}_j$)	3.781	6.266	1.436
	(0.175)	(0.253)	(0.141)
Upper Limit (m_H)	0.160	0.160	0.080
	(0.028)	(0.021)	(0.012)

NOTE.—Each column reports the estimated number of loans bunching at the conforming limit ($\hat{B}$), the average (log) shift in mortgage balance in response to the conforming limit among marginal bunching individuals ($\Delta\hat{m}$), the excess mass at the conforming limit ($\hat{B}/\sum_{j=L}^{0}\hat{n}_j$), and the upper limit of the excluded region used in estimation (m_L). Estimates are reported separately for the combined sample of all loans (column 1), fixed-rate mortgages only (column 2), and adjustable-rate mortgages only (column 3). Standard errors (in parentheses) were calculated using the bootstrap procedure described in section 2.4.1.

TABLE 2.2
JUMBO-CONFORMING SPREAD ESTIMATES, PERCENTAGE POINTS

	(1) Baseline	(2) Splines	(3) Within $50k of CLL	(4) Within $10k of CLL
	Fixed-Rate Mortgages			
OLS	0.179	0.182	0.171	0.172
	(0.002)	(0.002)	(0.014)	(0.038)
IV	0.107	0.121	0.118	0.126
	(0.010)	(0.009)	(0.028)	(0.056)
Observations	1,061,738	1,061,738	263,641	87,617
	Adjustable-Rate Mortgages			
OLS	-0.076	-0.090	-0.299	-0.362
	(0.009)	(0.009)	(0.037)	(0.083)
IV	-0.001	0.004	0.074	0.040
	(0.020)	(0.019)	(0.054)	(0.109)
Observations	692,233	692,233	157,779	39,542

NOTE.—Standard errors in parentheses. Estimates of jumbo-conforming spread using OLS and IV with an indicator for home value being greater than $1.25\bar{m}$ used as an instrument for jumbo-status, as described in the text. Controls include distance to CLL (cubic), LTV ratio, DTI ratio, missing LTV and DTI ratios, FICO score, missing FICO score, PMI, prepayment penalty, and mortgage term, as well as month by zip-code fixed effects. Column 1 includes linear effects of LTV and DTI ratios. Column 2 includes cubic B-splines in LTV and DTI ratios, as well as FICO score. Columns 3 and 4 limit the sample to loans near the CLL.

TABLE 2.3
INTEREST RATE AND MONTHLY PAYMENT ELASTICITIES OF MORTGAGE DEMAND, FRMs ONLY

	(1)	(2)	(3)	(4)
	Interest Rate Semi-Elasticity (ϵ_s^r)		Payment Elasticity (ϵ^P)	
	$\Delta \hat{r}$		$\Delta \log \hat{r}$	
$\Delta \hat{\tilde{m}}$	0.171	0.118	0.023	0.017
	(0.014)	(0.028)	(0.002)	(0.005)
0.052	-0.015	-0.022	-0.188	-0.246
(0.003)	(0.002)	(0.005)	(0.022)	(0.064)
0.063	-0.022	-0.031	-0.267	-0.351
(0.003)	(0.002)	(0.008)	(0.028)	(0.091)
0.083	-0.037	-0.053	-0.445	-0.589
(0.004)	(0.004)	(0.013)	(0.051)	(0.156)

NOTE.—Table reports estimates and standard errors (in parentheses) of the interest rate semi-elasticity and monthly payment elasticity of mortgage demand for a range of different jumbo-conforming spreads and behavioral responses estimated from bunching. Each cell reports the elasticity implied by the estimated (log) behavioral response ($\Delta \hat{\tilde{m}}$), and corresponding jumbo-conforming spread estimated in percentage points ($\Delta \hat{r}$) and in logs ($\Delta \log \hat{r}$), respectively. Standard errors for the bunching estimates were calculated using the bootstrap procedure described in section 2.4.1. Standard errors for the elasticities were calculated using the delta method.

CHAPTER 3 : The Role of Contagion in the Last American Housing Cycle

3.1. Introduction

One of the striking features of the recent U.S. housing cycle is its heterogeneity across markets. Both the magnitudes and timing of price swings varied greatly across metropolitan areas (Sinai, 2012). Figure 3.1 plots the geography and timing of the start of housing booms at the metropolitan area level from 1993 to 2009 based on estimates reported in Ferreira and Gyourko (2011).[1] The top left panel marks the 15 primarily rust belt and interior markets that never boomed. The other panels show that the remaining markets boomed at very different times over a nearly decade-long period from 1997-2005 and that the timing of these booms was non-random. The housing boom spread from what were initially highly concentrated areas on the two coasts, with the earliest booms beginning between 1997-1999 in California and the mid-New England region. On the west coast, housing booms eventually spread inland towards central California and to neighboring states to the east and north. On the east coast, housing booms spread to other markets in New England and then to neighboring regions, eventually reaching the majority of Florida markets by 2004 and 2005. These patterns are suggestive of spillover effects that disseminate positive housing price shocks from one market to another.

In this paper we investigate whether such spillovers were an important element of the last American housing cycle and also directly test which mechanisms may have contributed to them. In the financial economics literature, a spillover of this type is often referred to as contagion when it is found following a negative shock to one or more countries or markets.[2] While we focus on spillovers from a positive shock in much of the paper, we use the terms spillover and contagion interchangeably in order to emphasize the close intellectual linkage of our work with the analysis

---

[1] They define the beginning of a metropolitan area's housing boom by the quarter in which there is a structural break (a discrete positive jump in this case) in the market's house price appreciation rate. This methodology and the rationale behind it are discussed more fully below in Section 3.2.

[2] See Forbes (2012) for an excellent recent review of that literature, and Dungey et al. (2005) for a more technical analysis of the challenges involved in convincingly estimating contagion or spillover effects. Previous work on financial market contagion includes studies of the 1987 U.S. stock market crash (King and Wadhwani, 1990; Lee and Kim, 1993), the 1994 Mexican peso crisis (Calvo and Reinhart, 1996), and the Hong Kong stock market and Asian currency crisis of 1997 (Corsetti et al., 2005).

of contagion in financial economics. We define contagion as the price correlation across space between two different housing markets following a shock to one market that is above and beyond that which can be justified by common aggregate trends.[3]

The nature of the housing market and richness of our data allow us to address several empirical concerns that plague previous contagion-related research. One example involves the determination of the relevant period(s) in which to study spillovers. For example, a non-*ad hoc* procedure for identifying the timing of a shock is preferred to an arbitrary choice of time period 'after the fact.' This typically is not feasible in most studies of stock market or currency crises because there is little or no variation across countries in the onset of those events. Fortunately, this was not the case during the most recent housing cycle which saw substantial variation across markets (Ferreira and Gyourko, 2011; Sinai, 2012). In addition, we are able to appeal to existing theory as a guide in helping to determine how to time the beginning of a local boom. As is described more fully in the next section, implications of Glaeser et al.'s (2012) dynamic version of the classic model of spatial equilibrium in urban economics lead us to date the beginning of a given market's boom by whether and when there was a structural break in that area's price appreciation rate. These estimates were shown in Figure 3.1 and provide us with substantial variation in the timing of the start of local market booms that can be used in testing for the existence of contagion.

A second advantage is provided by the use of a voluminous micro-level data on U.S. housing transactions. We have over 23 million observations on individual home sales in 99 metropolitan areas dating back to the early 1990s in most cases.[4] This data enables us to address specification search bias of the type identified by Leamer (1978), which arises when the same data is used to identify both the timing of a shock and the magnitude of the volatility during that period. Our strategy uses randomly split samples to separately identify the timing of booms and the magnitude of price

---

[3]There is no single, agreed-upon definition of contagion, but our definition is similar in spirit to many used in the financial economics literature. See, for example, Forbes (2012), which emphasizes the distinction between contagion and interdependence, with the latter term reflecting when events in one country affect others in all states of the world, not just after severe negative events.

[4]The property transaction data is collected by Dataquick or by intermediaries from county assessor's offices and captures the universe of sales.

volatility in those periods.[5] Most contagion studies in financial economics are not able to deal with this issue because they use a single aggregate stock index in each country. Doing so increases the likelihood of falsely concluding that there are more and bigger booms (or crisis periods in the financial economics literature) than truly exist.

Third, the richness of our data and the variation in the timing of booms across markets also helps us deal with omitted variable biases in several ways. We use the time line of a neighbor's boom as our source of variation in the data to identify contagion effects. Our baseline specification involves regressing a focal market's price changes on a series of indicators reflecting whether the relevant neighboring market is booming and how proximate a given period is in time to the start of that market's boom.[6] The added degrees of freedom afforded by the multiple, non-contemporaneous booms we observe also allow us to control for omitted factors that might reflect common economic shocks. We do so through the inclusion of time by census division dummies, lagged price changes, as well as a host of local fundamentals. We also are able to address the Forbes and Rigobon (2002) critique regarding heteroskedasticity, whereby increased volatility in the 'crisis period' (when the boom starts, in our context) generates upward bias in correlation coefficients across markets. We adjust for the volatility in prices being higher than normal when the boom starts by directly controlling for the time line of the focal market's boom. Finally, in an alternative specification, we also address the potential for reverse causality with an instrumental variable approach using further lags of close neighbor's price changes.

Our main conclusion is that contagion played a statistically and economically significant role in the development of the most recent housing boom. The elasticity of prices in a typical focal market with respect to those in the closest metropolitan area in the year following the beginning of the boom in that neighboring area ranges from 0.10 to 0.27.[7] Empirically, the upper end of our elas-

---

[5]This is the same strategy followed by Card et al. (2008) in their study of tipping points in residential segregation models.

[6]As expected, this is very important empirically. Naively regressing 'price on price' yields contagion estimates that are 3-5 times larger than the results we report below from our preferred specification.

[7]We also report evidence of on-going contagion effects as the boom builds, but those results are potentially confounded by feedback effects. This is not a concern at the start of local booms because the data show no pre-trends, with markets appearing to be on their equilibrium paths before the initial jump in price growth when the booms commenced.

ticity range implies that from one-fourth to one-third of the average jump in price growth at the start of a typical local boom was due to contagion effects. This average impact is driven entirely by the physically closest neighbor and is only detected if the nearest neighbor had a statistically significant housing boom. There is no evidence of spillovers on prices arising from more geographically distant markets. In addition, this impact does not vary materially with the number of miles between the focal market and its nearest neighbor. As a robustness check, we investigated whether there is evidence of contagion on the extensive margin. Hazard models show that the probability of a boom beginning this quarter is indeed influenced by close metros that boomed in the previous quarter.

We also investigated heterogeneity in the price contagion along a number of other non-distance related dimensions. For example, one might expect contagion to be stronger when the transmission is from larger neighbor to a smaller focal market. That is precisely what we find. The magnitude of the contagion effect also varies with the elasticity of supply of the focal metropolitan area. Specifically, there is no evidence of contagion in inelastically supplied markets, but the estimated impact in the most elastically supplied markets is double the average effect discussed above.

What mechanisms could explain the observed price contagion effect? Fundamental factors including local income, migration flows between focal and neighboring markets, lending behavior in both markets, and speculative behavior in the focal market, do not show much variation after the beginning of the boom of the nearest neighbor. We also include a simple form of expectations of these local market fundamentals in our main econometric model to see if this materially affects the magnitude of the estimated contagion effect. It does not. This indicates that at least some of the contagion we estimate may be due to forces not related to the fundamentals analyzed in this paper. This has potentially important implications for policy makers. To the extent that the spread of a housing boom is even partially due to non-fundamental forces (e.g., some type of irrational exuberance or otherwise mistaken perceptions of the influence of a neighboring market), it may be worth rethinking the advisability of policy makers not responding to a boom.

Despite finding important spillovers during the run up of the U.S. housing boom, we report mixed

evidence that contagion played a role in the bust. Estimated elasticities are zero at the beginning of the bust, but they increase to about 0.15 by the third year of the housing bust of the nearest neighbor. We also cannot detect any impact of contagion on the extensive margin during the bust. This is perhaps not all that surprising since the timing of the bust across MSAs is heavily concentrated in an 18 month period during 2006 and 2007, while the buildup of the housing boom took almost a decade. This highlights difficulties in detecting spillovers during economic or financial crashes/booms that quickly spread across countries, regions, or firms.[8]

While our research is motivated by prior research in financial economics noted above, it is also part of a growing body of work on the most recent housing cycle. One important strand of that research tries to understand whether the most recent cycle was a bubble.[9] Another more voluminous body of papers analyzes the bust and its consequences.[10] There are also some prior studies of price spillovers in the housing market.[11] In this paper we focus on one particular facet of the cycle, the role of contagion, but contribute to the literature by: a) looking at contagion over the full time span of the cycle, b) improving empirical identification in many respects, c) estimating heterogeneity in the contagion elasticity, and d) investigating the potential causes of

---

[8]In addition, we do not have the added advantage of relying on a prediction from external theory to date the beginning of the bust. In this part of the analysis, we follow the tradition in the contagion literature and date the bust's beginning in an *ad hoc* manner based on the assumption that it begins in the quarter after nominal price levels peaked in a given metropolitan area.

[9]Shiller (2005) provides perhaps the most famous characterization of the boom as a non-rational event. Others recently have estimated rational expectations general equilibrium models to try to explain the national aggregate price data (Favilukis et al., 2011) or the serial correlation and volatility of prices and quantities within and across metropolitan areas (Glaeser et al., 2012). Related work includes Arce and Lopez-Salido (2011), Burnside et al. (2011), Lai and van Order (2010), and Wheaton and Nechayev (2008).

[10]Much of this research focuses on the subprime sector (Bajari et al., 2008; Danis and Pennington-Cross, 2008; Demyanyk and van Hemert, 2011; Gerardi et al., 2007; Goetzmann et al., 2009; Mayer and Pence, 2008; Haughwout et al., 2010), mortgage securitization (Bubb and Kaufman, 2009; Keys et al., 2010), the default/foreclosure crisis (Adelino et al., 2009; Campbell et al., 2011; Foote et al., 2008; Gerardii et al., 2008; Mayer et al., 2009; Mian and Sufi, 2009; Mian et al., 2010; Piskorski et al., 2009) or the role of government regulation (Avery and Brevoort, 2010; Bhutta, 2009; Ho and Pennington-Cross, 2008)

[11]Some early examples include Clapp et al. (1995) and Dolde and Tirtiroglu (1997), who use data on towns in Connecticut and San Francisco (respectively) to test for the existence of cross-market linkages in price movements. More recently, Holly et al. (2010, 2011) use data on U.S. states and U.K. regions to study the spatial and temporal diffusion of changes in house prices during the most recent cycle. Fuss et al. (2011) and Cotter et al. (2011) use publicly available, metropolitan area-level house price series to test for the existence of some form of contagion during the most recent housing cycle. However, all studies use of the same aggregate market-level price data to determine both the timing of the crisis period and to measure the magnitude of volatility changes during that period makes their estimates susceptible to specification search bias. In addition, the timing of the shock is usually defined in an ad-hoc way and there are questions about how that research deals with omitted factors.

contagion across housing markets.

The plan of the paper is as follows. The next section motivates our use of an urban economic approach to analyzing contagion in housing markets and discusses our method for dating the beginning of the boom. Section 3.3 then describes the various data sources employed and the variables created. Section 3.4 discusses the different types of specifications estimated, reports results, and explores potential mechanisms that might explain the contagion effects. In Section 3.5 we estimate alternative specifications, look at the extensive margin of contagion, and whether we can find contagion at the housing bust. There is a brief conclusion.

3.2. An Urban Economic Motivation and Definition of Contagion

3.2.1. Timing of Local Housing Booms

Any analysis of possible contagion effects in the spreading of the recent housing boom first requires knowledge of the timing of the beginning of the boom in different markets. Our data on the timeline of local booms come from estimates discussed above and reported in Ferreira and Gyourko (2011). In that work, the start of local booms is determined by when there was a structural break in each area's price appreciation rate series. The justification for that strategy is based on implications of the dynamic spatial equilibrium model developed in Glaeser et al. (2012). In particular, that model implies that in steady state each local market will exhibit constant and continuous growth paths for house prices, new construction and population.[12]

Empirically, this means that we should see house prices in a given market growing at a constant rate unless there is a shock to local productivity, amenities, or expectations, in which case we

---

[12]Glaeser et al. (2012) introduce dynamics into Rosen's (1979) and Roback's (1982) classic static model of spatial equilibrium. In this compensating differential framework, house prices (P_i) are the entry fee paid to access the wages (W_i, which reflect productivity) and amenities (A_i) of labor market area i. Their model is closed with an assumption that there is some elastically supplied reference market area which is always open to another household. The utility level available in the reference market is given by U^*, and establishes the lower bound on utility provided in any market. In the long run, perfect mobility ensures that U^* is achieved in all markets, so that in equilibrium, no one has an incentive to move to another place which offers higher utility. A very simple, linear version of this framework would imply that $U^* = W_i + A_i - P_i$, so that $dP_i = dW_i + dA_i$ in equilibrium. The steady state rate of price appreciation need not be zero. Secular trends in house prices can come from an underlying trend in housing demand as long as the market is not in perfectly elastic supply. It can also arise from trends in physical construction costs under certain conditions.

would then observe a discrete jump in the appreciation rate for that market. The data are generally consistent with this predicted pattern. As a particularly stark example, Figure 3.2 plots annualized house price appreciation rates over time for the Las Vegas market. This graph shows that fast-growing market to be appreciating at a high, but roughly constant rate for many years before house price growth escalates sharply at the beginning of its boom. Informally, our approach defines the beginning of the housing boom in a local market as the point at which house price growth rates exhibit this type of discrete jump.

To formalize this idea, we start with the following reduced form model of house price growth in MSA i at time t:

$$PG_{i,t} = d_{i,t} + \epsilon_{i,t}, \quad t = 1, \ldots, T. \tag{3.1}$$

Glaeser et al. (2012) implies that $d_{i,t} = d_{i,0}$ for all t if the market is on its steady-state growth path. However, if there is a shock to local productivity, amenities, or expectations at time t then the price growth rate will exhibit a discrete jump in that period. The beginning of a local housing boom can thus be identified by testing for the existence of one or more structural breaks in the parameter $d_{i,t}$. To carry out this test we follow well-established methods in the time series literature for estimating structural breaks.

Borrowing heavily from Estrella's (2003) notation, the null hypothesis is that $d_{i,t}$ is constant for the entire sample period

$$H_0 : d_{i,t} = d_{i,0}, \quad t = 1, \ldots, T.$$

The alternative is that $d_{i,t}$ changes at some proportion, $0 < \pi_i < 1$, of the sample which marks the beginning of a housing boom in market i. Specifically the alternative hypothesis is

$$H_1 : d_{i,t} = \begin{cases} d_{i,1}(\pi_i) & t = 1, \ldots, \pi_i T \\ d_{i,1}(\pi_i) & t = \pi_i T + 1, \ldots, T. \end{cases}$$

For any given π_i, it is straightforward to carry out this hypothesis test. However, things are

slightly more complicated when π_i is unknown and the determination of its value is the primary interest, as the case here.

To see how we estimate the value of π_i and assess its statistical significance, let $\Pi_i = [\pi_{i,1}, \pi_{i,2}]$ be a closed interval in $(0, 1)$ and let S_i be the set of all observations from $t = int(\pi_{i,1}T)$ to $t = int(\pi_{i,2}T)$, where $int(\cdot)$ denotes rounding to the nearest integer. The estimated break point is the value t^* from the set S_i that maximizes the likelihood ratio statistic from a test of H_1 against H_0.[13] That is, for every $t \in S_i$ we construct the likelihood ratio statistic corresponding to a test of H_1 against H_0 for that value of t, and we take the t that produces the largest test-statistic as our estimated break point for MSA i.

Assessing the statistical significance of this breakpoint estimate requires knowing the distribution of the supremum of the likelihood ratio statistic as calculated from among the values in S_i. Let $\xi_i = sup_{S_i} LR$ denote this supremum. Andrews (1993) shows that this distribution can be written as

$$P(\xi_i > c) = P\left(sup_{\pi_i \in \Pi_i} Q_1(\pi_i) > c\right) = P\left(sup_{1 < s < \lambda_i} \frac{||B_1(s)||}{s^{1/2}} > c^{1/2}\right) \tag{3.2}$$

where $||B_1(s)||$ is the Bessel process of order 1, $\lambda_i = \pi_{i,2}(1 - \pi_{i,1})/\pi_{i,1}(1 - \pi_{i,2})$, and

$$Q_1(\pi_i) = \frac{(B_1(\pi_i) - \pi_i B_1(1))'(B_1(\pi_i) - \pi_i B_1(1))}{\pi_i(1 - \pi_i)}.$$

Direct calculation of the probability in (3.2) is non-trivial and prior research has relied on approximations that are typically based on simulation or curve-fitting methods (Andrews, 1993; Hansen, 1997). However, Estrella (2003) provides a numerical procedure for calculating exact p-values that does not rely on these types of approximations. We use this method to calculate p-values for the estimated break point, π_i, for each MSA in the sample.

Note that this method does not provide an unbiased estimate of the magnitude of the change in price growth rates at the breakpoint, $d_{i,2}$. Under the null hypothesis that there is no break

---

[13] We use the terms supremum and maximum interchangeably in this exposition. Technically, all of the results are in terms of the supremum of the likelihood ratio statistic.

point, the estimate of $d_{i,2}$ has a nonstandard distribution and OLS estimates of its magnitude will be upwardly biased in absolute value. This problem can lead to an increased chance of falsely concluding that $d_{i,2} \neq 0$ and is a form of specification search bias arising from the fact that the same data is being used to estimate both the timing and the magnitude of the structural break.

Several approaches for adjusting the estimate of the magnitude of structural break have been suggested and are typically based on simulations of the distribution of $d_{i,2}$ under the null hypothesis of no break point (Andrews, 1993; Hansen, 2000). Our approach to correcting the estimates of $d_{i,t}$ follows the recently suggested method used by Card et al. (2008) of randomly splitting the underlying sample of housing transactions into two and using one sample to estimate the timing of the boom and the other to estimate the magnitude of price changes around that time. The idea is that if the two subsamples are independent, then estimates of $d_{i,2}$ from the second sample, which was not used to estimate the location of the break point, will have a standard distribution even under the null hypothesis of no structural break in the first sample. In practice, we randomly split our sample of unique houses into two and create separate price growth series for each sample of houses. The first price series is used to estimate the timing of the boom following the method just discussed, while the second is used to analyze the magnitude of price changes following housing booms in neighboring markets.[14]

The procedure above will generate a breakpoint estimate regardless of whether the structural break represents a positive or negative change in the price growth rate. In the cases where the estimated break point is either insignificant or implies a negative change in growth rates, we conclude that the market did not have a boom. That is the case for the 15 interior markets shown in the first panel of Figure 3.1. For those locations where we do find evidence of a statistically significant and positive break point, we also test for the existence of two breaks against the null hypothesis of

---

[14]Not accounting for this issue can result in large biases. Card et al. (2008) noted that their estimates from the full sample were somewhat larger than the estimates from the split sample approach. In our case, if we use the full sample of transactions to estimate both the timing of the beginning of the boom and the magnitude of the jump in price growth at that time, the estimated jumps are also larger, on average, than those arising out of the split sample estimation. Part of this difference, of course, could also be attributed to additional random variation that arises when reducing sample sizes by half. This may be an issue for our smaller cities that do not have as many transactions per quarter as the big cities, and are more subject to the influence of potential outliers.

only one. To do so, we closely follow Bai (1999) and Bai and Perron (1998) and we refer the reader to Appendix A.5 for the details of this procedure. About half of the MSAs were found to have experienced more than one structural break. However, for many of those cases, the secondary breaks were either small economically or not significantly different from zero. The estimation of a secondary break generally does not displace the location of the main structural break either. Moreover, comparison of histograms of timing of local booms based on one-break or two-break methods lead to similar distributions of local booms over time. Given these facts, we simplify our analysis by only using the one-break method in the empirical study below. When necessary, we also report robustness tests based on the multiple break method.

3.2.2. Contagion in Local Markets

A potential role for neighbors to influence house price growth in a focal market arises naturally within the dynamic urban model of spatial equilibrium referenced above. While users of that framework typically presume that shocks originate from own market fundamentals, they could arise from neighboring markets as well. In our case, we are interested in whether neighbors that just had housing booms influence housing market outcomes in the focal metropolitan area, all else constant. Note that any such contagion effects could arise from fundamental or behavioral factors.

An example of a fundamental factor generating spatial spillovers is a positive industry or income shock that triggers a housing boom in a local labor market. For example, if there is such a shock in the Silicon Valley, house prices in the San Jose-Sunnyvale-Santa Clara MSA will increase in the short run given supply constraints, and perhaps start a housing boom. Neighboring areas, such as the San Francisco-Oakland-Fremont MSA (or smaller, but more distant, metropolitan areas in the central valley of California), may eventually benefit from that positive income shock, as some of the Silicon Valley jobs could migrate to nearby areas. Even though such fundamental spillovers may occur with lags, house prices in the neighboring markets should immediately capitalize the expectation of future economic growth.

Housing market shocks could also be disseminated through the credit market channel. In the example above, Silicon Valley lenders may achieve extra profits since foreclosures and delinquencies tend to decline during an economic boom. If those lenders decide to reinvest profits and expand market shares in a nearby MSA—possibly because it is less costly to expand business to nearby communities—then the San Francisco metro may observe a shift in the availability of credit, which will boost its housing market both in the short and long runs.[15] A similar mechanism could exist for land owners and housing investors in the Silicon Valley. Their wealth increases after the beginning of the Silicon Valley housing boom, which could trigger an expansion of investments into neighboring MSAs.[16]

Spillovers could arise even in the absence of expected future fundamental changes in San Francisco or the central valley. Residents in those neighboring markets may be right to think that some type of positive income spillover will occur from the Silicon Valley boom, but they may incorrectly predict its magnitude. Those biased expectations can lead to short-run increases in their housing prices.[17] In addition to this type of behavior, irrational factors may lead residents in the focal market to have not only incorrect, but also non-fundamentally based expectations about future price growth in their market following a shock to a nearby neighboring area. Finally, the housing boom in the Silicon Valley may generate another type of behavioral spillover effect—namely, an increase in Silicon Valley housing prices may lead its own local residents to pay more attention to what is happening in neighboring markets, such as the city of San Francisco. Therefore, a shock that makes the focal housing market and its interactions with neighbors more salient to investors may itself lead to stronger contagion.

Regardless of its source, contagion would manifest itself in the form of abnormally large increases

---

[15]Ortalo-Magne and Rady (2006) show how increases in the availability of credit to marginal buyers can lead to overreaction in house prices.

[16]Chinco and Mayer (2012) report that out of town speculators played a significant role in the housing boom and that their presence may have exacerbated price increases in some markets. Bayer et al. (2011) similarly document the rise of speculative activity during the housing boom.

[17]A similar mechanism underlies the analysis in King and Wadhwani (1990) who document that contagion in financial markets can arise as a result of attempts by rational agents to infer information from price changes in other markets. Similarly, Clapp et al. (1995) document characteristics of house price dynamics that could be consistent with rational learning. Burnside et al. (2011) and Favara and Song (2010) show how the presence of optimistic agents in the housing market can lead to increases in house price levels and volatility.

in focal market price growth rates immediately following a boom in a neighboring market. Our empirical approach makes use of the estimates of the timing of local booms discussed above in order to test for this type of spillover and its potential mechanisms. However, before discussing the econometric specifications we estimate, we first give a brief description of the data underlying our analysis.

3.3. Data

Our house price data come from DataQuick, a private data vendor which collects the universe of housing transactions from county recorder's offices in markets across the country. The sample used is for 99 metropolitan areas, with information on over 23 million individual observations ranging from the first quarter of 1993 (1993(1)) through the third quarter of 2009 (2009(3)). We randomly split the sample into two and in each subsample we create a constant quality quarterly price index for each MSA.[18] From these indices we create the annualized growth series used in estimating the timing of the boom and in assessing how neighboring booms affect price growth in focal markets. The mean, standard deviation, and interquartile range for the price index we use to measure magnitudes are reported in the first row of Table 3.1.

We also create a number of variables to measure fundamentals that are potentially correlated with house price growth and the timing of the beginning of local housing booms. These are reported in subsequent rows of Table 3.1. We consider three types of fundamentals: (1) demand shifters, such as the average income of mortgage applicants, MSA-level unemployment rates, and net migration flows; (2) buyer characteristics and property traits, including the percentage of speculators, the percentage of minority buyers and the average square footage of transacted housing units; and (3) credit market conditions, measured by the average loan-to-value ratio of home purchases, the percentage of mortgages originated by subprime lenders and those insured by the FHA.

---

[18]We create a MSA-level constant quality house price series by quarter using hedonic regressions. Price, in logarithmic form, is modeled as a function of the square footage of the home entered in quadratic form, the number of bedrooms, the number of bathrooms, and the age of the home. We also created a version of the Case and Shiller (1987) repeat sales price index for 14 Case-Shiller markets that overlap with the DataQuick files, and found that the simple correlation of appreciation rates on the two different indexes based on DataQuick is usually higher than 0.9. We employ hedonic price indexes because their data requirements are much less onerous.

To construct many of the demographic measures of home buyers, we merge the DataQuick files with Home Mortgage Disclosure Act (HMDA) data, which provide information on the income and race of all mortgage applicants. In each time period, we calculate the average income of all local loan applicants as reported in HMDA. Similarly, the 'Percent Minority' variable reflects the fraction of African-American and Hispanic loan applicants as coded in the HMDA files. Because these measures reflect the characteristics of all mortgage applicants, and not only the set who end up purchasing a home, we take them to be an accurate description of the race and income of potential homebuyers in each market.

MSA-level unemployment rates come from the Bureau of Labor Statistics' *Local Area Unemployment Statistics* series, and net migration flows are calculated using data on county-to-county migration patterns provided on an annual basis by the Internal Revenue Service.

'Percent speculators' refers to the fraction of transactions involving a speculator on either the buyer or the seller side of the transaction. We leverage the fact that we observe the names of both the buyer and seller for each transaction in order to define speculators in a similar way as Bayer et al. (2011). Specifically we define a person as a speculator if he or she is observed to have 'flipped' at least two homes in the same metropolitan area during the entire course of the sample where a flip is defined as a purchase and sale of the same home within a two-year period. We then consider all transactions in which that person is involved as either the buyer or seller as being 'speculative.'

Credit market variables include the average loan to value ratio (LTV) among homebuyers in Data-Quick, the fraction of FHA-insured loans, and the fraction of subprime loans. We use information on the names of the underlying mortgage lenders from the DataQuick files to calculate the share of subprime loans. More specifically, we obtained lists of the top twenty subprime lenders from 1990-onward in a publication now called *Inside Mortgage Finance*.[19] 'Percent subprime lenders' is then defined as the share of mortgages issued by these top twenty lenders.

---

[19]This publication claims to capture up to 85% of all subprime originations in most years. Previously, it was named *B&C Mortgage Finance*. See Chomsisengphet and Pennington-Cross (2006) for more details on these lenders and lists. Other papers such as Mian and Sufi (2009) and Keys et al. (2010) have access to micro-level FICO scores and use that to define subprime borrowers.

3.4. Econometric Model and Estimates

3.4.1. Econometric Model

As discussed in Section 3.2, the implications of the underlying dynamic model of urban economics are readily extended to include shocks from neighboring housing markets. We estimate the following reduced form model to gauge the impact of housing market booms of close neighbors on the prices of the focal MSA m in Census Division d and quarter t:

$$
\begin{aligned}
\log\left(P_{m,d,t}\right) = & \sum_{r=-4,r\neq0}^{4} \theta_r^1 * \psi_{m,r}^1 + \sum_{r=-4,r\neq0}^{4} \theta_r^2 * \psi_{m,r}^2 + \sum_{r=-4,r\neq0}^{4} \theta_r * \psi_{m,r} \\
& + \sum_{\tau=1}^{4} \theta_\tau * \log\left(P_{m,d,t-\tau}\right) + \gamma_{d,t} + \epsilon_{m,d,t}
\end{aligned}
\tag{3.3}
$$

The first term on the right-hand side of equation (3.3) contains the primary variables of interest, the $\psi_{m,r}^1$'s, which are indicators for the years relative to the beginning of the housing boom (r) of the closest neighboring market (neighbor number 1). The coefficients, θ_r^1, on these indicator variables describe how prices in the focal market evolve over the course of its nearest neighbor's housing boom. We define Relative Year 0 to be the 12-month period prior to the beginning of the neighboring market's boom.[20] Relative year 1 then includes the quarter in which the boom starts as well as the subsequent three quarters. Relative years from -2 to $+3$ are entered individually, with all relative years greater (less) than those numbers binned together.[21] This allows us to see whether there are any important pre-trends and to track the build-up of the boom after it starts. The second term on the right-hand side of (3.3) includes an analogous set of controls, denoted $\psi_{m,r}^2$ for the second closest neighbor. In a set of robustness tests presented below, we also control for log prices of other near MSAs, to make sure that our elasticity estimates from the nearest neighbor

---

[20]We work with 12 month periods because there is noise in the quarterly data that is not due solely to error in the estimation of the break point. For example, it is common for there to be at least a one quarter difference between the time that a transactions price is agreed upon and when the actual closing occurs. In addition, we know that prices in housing markets do not follow a random walk, but move slowly and are strongly positively correlated over short horizons (Case and Shiller, 1987, 1989). Locations for which we do not estimate a statistically significant boom are still assigned a relative year according to their estimated break points.

[21]We did estimate all our models with lengthier spans of individual relative years controlled for, but they did not yield any new insights beyond those reported below.

are not confounded by other neighbors, or other neighboring shocks.

Physical proximity is the most natural measure of distance, and the specifications reported below assign the nearest neighbor based on the number of miles between the centroids of the relevant metropolitan areas.[22] We also experimented with a measure of proximity based on migration flows between pairs of markets. Those results were qualitatively and quantitatively similar, so we do not report them for space reasons.[23]

This specification also controls for the time line of the focal market's boom via the third term on the right-hand side of equation (3.3). By including the relative year effects of the focal MSA (denoted $\psi_{m,r}$, where the absence of a superscript indicates the variable refers to the focal market itself), we control for the average increase in prices of the focal market over the course of its own boom. This vector serves a similar role to the adjustments made in related financial economics research to deal with upward bias in contagion estimates arising from volatility being higher in all markets during 'crisis' periods (e.g., Forbes and Rigobon (2002)).

The remaining terms of equation (3.3) include four quarterly lags of (log) focal market prices to control for potentially unobserved time-varying characteristics of MSA m[24], as well as Census Division-by-quarter fixed effects to deal with common regional shocks that might influence close neighbors simultaneously. Finally, note that we do not control for contemporaneous focal market fundamentals in this baseline specification. This is because they could represent intermediate outcomes through which the contagion effect may be operating. In the mechanisms section be-

---

[22]Distances are calculated using the full set of MSAs according to the 2000 Census. Because we only have price data for 99 of these MSAs, some data for nearest neighbors remain empty in 24 cases. In our regressions, we create an indicator for whether we have price data for the nearest neighbor, and interact this indicator with the relevant relative years. The results are qualitatively similar when we drop MSAs with missing neighbors' price data and also when we calculate distances using only the 75 MSAs for which we have data.

[23]This economic measure of distance based on migration flows between metropolitan areas has been used in other research (Sinai and Souleles, 2009). It is strongly positively correlated with geographic distance between markets. For example, the probability of the physically closest neighbor also being the closest economic neighbor is 57%, and the probability it is one of the two closest economic neighbors is 76%. Hence, it is not surprising these two measures of distance yield similar results.

[24]We also considered specifications that dispense with the lags of the dependent variable in favor of MSA fixed effects. Results from these specifications are qualitatively similar. However, we believe that the lagged dependent variable specification is more appropriate given that omitted time-varying common factors are more likely to confound the contagion effect than unobservable but fixed MSA-specific characteristics.

low, we will directly estimate the impact of the neighbor's housing boom on those intermediate outcomes, and also test whether their inclusion in equation (3.3) mitigates the contagion effect.

3.4.2. Main Estimates

Column 1 of Table 3.2 reports baseline estimates of equation (3.3) for the metropolitan areas whose nearest neighbors had statistically significant booms. The reported coefficients show that prices were relatively stable in the three years prior to the beginning of the boom in the nearest neighboring MSA, so that there is no evidence of a pre-trend. In the year that the neighboring MSA begins its boom (Relative Year [1]), focal market prices then jump 0.87 percentage points and remain almost 1% higher for another couple of years. Column 2 reports the analogous coefficients on the relative year dummies for the set of neighboring MSAs for which we do not estimate a statistically significant or positive break point. These results confirm that a positive effect is only detected when the neighboring MSA actually had a housing boom. Hence, we find spillovers on focal market prices only if the nearest neighbor actually experienced a significantly positive shock. Columns 3 and 4 show analogous estimates for the second closest neighbors. The coefficients on the timeline of the boom for 2nd nearest neighbor generally are not statistically different from zero regardless of whether this particular neighbor had a housing boom. Hence spillovers mostly arise from the closest neighbor, and we will focus on those results in the remaining of the paper.[25]

In order to determine the elasticity of focal market housing price growth with respect to near neighbors' price growth, we need to gauge the magnitude of the housing boom for the nearest neighbors. The starting point of this exercise is to estimate a version of equation (3.3) that uses the log price of the nearest neighbor as the dependent variable. Table 3.3 reports those results. Pre-boom prices are trending down a bit in these results, and then jump 3% in the first year of the housing boom. By the third year of the boom, prices are 8% higher than the pre-boom period, and are more than 11% higher in subsequent years.[26]

---

[25]If we do not control for the time line of the focal market's boom, the point estimates of the nearest neighbor's contagion effects are about 20% higher than those reported just below for our baseline specification in Table 2. However, the results are not statistically different once the standard errors are taken into account. While this is consistent with price volatility being artificially high when the focal market itself is booming, we note that including these indicators could also be controlling for intermediate outcomes.

[26]These magnitudes are similar to those in Ferreira and Gyourko (2011) who conduct a similar exercise using the full

An upper bound on the implied elasticity can be computed by using only the estimates of price changes in the first year of the boom under the assumption that agents in these markets are myopic. Combining these figures with the estimates from Table 3.2 yields an elasticity of 0.27.[27] Smaller elasticities result if we consider cumulative price changes by the third year of the shock. In that case, the elasticity falls to about 0.10.

While differences in data and methodologies make it difficult to directly compare our estimates of contagion to others in the literature, our magnitudes appear to be smaller and we find that spillovers only arise from the closest neighbors.[28] We suspect this is partially due to our empirical strategy that minimizes specification search bias and includes various controls to deal with omitted factors related to common shocks.

A rough gauge of the relative importance of contagion in fomenting local booms can be made as follows. Ferreira and Gyourko (2011) concluded that jumps in one fundamental, local income growth, could account for one-half of the magnitude of the jump in price growth at the beginning of local booms (on average). Our estimates indicate that contagion did not play as important a role, but it still was economically meaningful, as it can account for over one-quarter of the jump at the start of the boom.[29] Using different and more conservative figures and assumptions reduce the share, but nothing reasonable can drive it below 10%.[30]

---

set of MSAs, not just those who serve as closest neighbors to some other MSA in the sample.

[27]This is the result of dividing the 0.008688 spillover estimate from column 1 of Table 3.2 by 0.03224 from column 1 of Table 3.3.

[28]For example, Cotter et al. (2011) regress housing price appreciation in eleven MSAs near San Francisco on the contemporaneous and 3 quarterly lags of San Francisco's housing appreciation rate and report coefficients ranging from 0.05-0.67 on lagged housing price appreciation of San Francisco. Fuss et al. (2011) model the volatility spillover intensity and suggest that a 1 percentage point shock in housing return in Las Vegas could result in an eventual 0.19 percentage point increase in housing return in San Diego.

[29]To see this more clearly, start with the 0.27 elasticity just discussed. Given the 3.2% average jump in Relative Year [1] reported in Table 3.3, that aforementioned elasticity implies that about 0.9 points, or about 27 percent of the jump at the start of the boom, can be explained by contagion.

[30]It is not clear what figure to use for the average level of price growth in the denominator of this ratio, which is why we focus more on the elasticity. The 3.2% number used here based on Table 3.3's results is close to the log price changes reported in Ferreira and Gyourko (2011). However, one also could use the 6.5% average jump in price growth rate (not the change in the log prices) also reported by those authors. That number arises from an estimate that does not control for any year or metropolitan area fixed effects. With that denominator, the 0.9 points amounts to about 14% of the jump at the start of the boom. As noted in the text, there are no reasonable assumptions one could make that drive the share below 10%.

Finally, the economic interpretation of the contagion estimates for the years after the beginning of the boom may be more complicated due to potential feedback effects. Feedback effects are less of a concern in the first year of the boom, as we showed that prices, roughly speaking, are in equilibrium right before that moment. It may not play a major role in subsequent years either, especially if only 10% or less of the main price effect propagates across close neighbors. Nonetheless, contagion estimates for relative years two and three, for example, are better thought of as reduced form estimates that include the impact of recent contagion, but that also embed a share of contagion from the complete path of price appreciation since the beginning of the boom.

3.4.3. *Heterogeneity in the Contagion Effect*

We test for heterogeneity in the average contagion effect along a number of dimensions. The first is distance. We already found that only the nearest neighbor matters. A natural extension is to ask whether the strength of the contagion impact associated with the closest neighbor increases with its proximity to its focal market. The first two columns of Table 3.4 show that the answer is no. Those figures are the output from a regression like that in equation (3.3) which further interacts the neighbor's relative year dummies with an indicator for whether that neighbor is more or less than the median distance of about 40 miles away from the focal market. Although the estimates tend to be imprecise, the point estimates show little difference between relatively close and farther away nearest neighbors, especially around the time when the neighbor's boom begins. Thus, contagion effects arise only from the nearest neighbor, but they do not vary materially based on how close that nearest neighbor is.[31]

It also seems natural to ask whether contagion impacts depend upon the relative sizes of the focal and neighbor markets. To investigate this, we classified focal MSAs whose population sizes are within 50% of the population of the closest neighbor as being of similar size. They are considered larger if they have at least 50% more population, and smaller if they have less than 50% of

---

[31]The interquartile range of distances between neighboring markets runs from 30-48 miles, so there is not much variation for much of the sample. The mean is larger at 74 miles, but that reflects the influence of Honolulu, whose nearest neighbor is over 2,000 miles away. The next biggest distance is 111 miles. We also experimented with alternative groupings such as dividing markets into whether their nearest neighbor was less than 30 miles away, from 36-60 miles away, and greater than 60 miles away. The results were no different from those reported here.

the population of the closest neighbor. Appendix Figure A.16 shows the distribution of relative population size, and the thresholds used to determine groups of MSAs for both geographic and economic neighbors. The estimates reported in the third and fourth columns of Table 4 indicate that the contagion effect on a focal market is larger if the nearest neighbor is substantially bigger than the focal area. Prices in the focal market are 1.3% higher immediately when the large near neighbor booms, but are little changed when the focal market is bigger (row 3 for Relative Year [1]). Prices stay higher in subsequent years for focal markets being influenced by large neighbors. In sum, size does matter and in an intuitive way in the sense that contagion effects are much larger (and consistently statistically significant) if the nearest neighbor is large relative to the focal market.

The final dimension along which we investigated whether there was any heterogeneity was by the degree of the focal market's elasticity of housing supply. For this test, we split the focal MSAs into three groups according to the supply elasticities provided by Saiz (2010).[32] Results are reported in the final two columns of Table 3.4 for the bottom third (supply inelastic) and top third (supply elastic) groups. Note that prices do not jump when the closest neighbor of an inelastically-supplied market begins to boom. However, for the most elastically-supplied metros, prices in the focal market are about 2% higher if its nearest neighbor begins to boom. The gap increases to about 3% by the third year of the boom. Performing calculations analogous to those discussed above for the economic importance of the average contagion effect show that spillovers could account for nearly two-thirds of the jump in prices at the beginning of booms in the most elastically supplied markets. Basic economics suggests that any contagion effects would be more likely to be capitalized in the inelastically supplied markets, *ceteris paribus*, so this outcome may seem counterintuitive at first glance. However, all else is not constant in this case. It turns out that a disproportionately large share of these markets are large coastal metropolitan areas with relatively small neighbors.[33] So,

---

[32]Saiz's supply elasticity estimates are available for only 76 of our metropolitan areas, so we start with a smaller sample for this particular analysis.

[33]Included in this most inelastic tercile are the metropolitan areas of Barnstable Town, MA, Boston-Cambridge-Quincy, MA-NH, Bridgeport-Stamford-Norwalk, CT, Cleveland-Elyria-Mentor, OH, Deltona-Daytona Beach-Ormond Beach, FL, Eugene-Springfield, OR, Jacksonville, FL, Los Angeles-Long Beach-Santa Ana, CA, New Haven-Milford, CT, New York-Northern New Jersey-Long Island, NY-NJ-PA, Oxnard-Thousand Oaks-Ventura, CA, Palm Bay-Melboune-Titusville, FL, Portland-Vancouver-Beaverton, OR-WA, Port St. Lucie-Fort Pierce, FL, Riverside-San Bernardino-Ontario, CA, San Diego-Carlsbad-San Marcos, CA, San Francisco-Oakland-Fremont, CA, Santa Rosa-Petaluma, CA, Sarasota-Bradenton-Venice, FL, Seattle-Tacoma-Bellevue, WA, Tampa-St. Petersburg-Clearwater, FL,

at least some of the variation we document by degree of supply elasticity could have been driven by the size results just discussed.[34]

3.4.4. Mechanisms

In this subsection, we investigate some of the potential mechanisms that could account for the influence of a near neighbor's boom on house prices in the focal market. We are particularly interested in whether our contagion effects are fundamentally-based in the sense discussed in Section 3.2. . If not, the relevance of our results for policy makers is increased, as they well may want to reevaluate their past practice of not intervening in response to asset booms in housing markets if their spread is based on some type of irrational exuberance or otherwise mistaken expectations.

We begin by asking whether there are visible economic changes in the focal market that may be driven by the neighboring market boom. We then alter our baseline specification to account for these potential fundamental drivers. Four fundamentals are investigated, with each having received prominent mention in previous academic research or by policy makers and the popular press. They are focal market income, mortgage market activity, net migration flows into the focal market and the share of house 'flips' in overall market sales. Table 3.5 reports results using these local market traits as the dependent variable in a specification similar to that in equation (3.3), with one difference being that here we include MSA fixed effects rather than own price lags.

Column 1 reports estimates for focal market income, where income is defined as the average income for all mortgage applicants in that market and quarter. If what is driving our contagion result is a real spillover such as focal market income going up because of boom in a nearby market, then we should see it changing along the timeline of the neighbor's boom. Table 3.5's results show that there is a jump from zero to 1% in Relative Year [1], but this impact is not precisely estimated

and Vallejo-Fairfield, CA. Each of these areas also has one of the 99 markets in our sample as its closest market.

[34]We also investigated whether contagion effects varied by the relative timing of booms in order to see if contagion impacts occurred primarily when the focal market itself was booming. Appendix Figure A.15 shows a histogram of the difference between the timing of the boom of the focal MSA and its closest neighbor. There is wide variation in timing of the booms of these pairs of markets. We find that contagion effect seems concentrated in focal MSAs that were already booming when the closest neighbor started to boom. This result suggests salience being a feature of the contagion effect. But as with the heterogeneity by supply elasticity, relative timing is correlated with other factors, such as market size.

(the t-statistic is 1.3). The results are similar for the second and third relative years. It is only at least four years after the nearest neighbor booms that we see focal market income higher by a statistically significant amount. One would not want to interpret these coefficients as proving that contagion does not operate via spillovers onto focal market income, but they also provide no robust evidence to the contrary.

The next two columns investigate whether the contagion effect might operate through credit markets in some fashion. We approach this question by examining two aspects of credit lender activities. First, we investigate whether the mortgage lender bases become more similar during and after the nearest neighbor booms. The intuition is that if lenders observe a boom in that neighbor, they might increase their activity in the focal market for the reasons discussed in Section 3.2. We use a proportional index to measure lender similarity.[35] Second, we investigate whether lenders speed up mortgage lending during and after the closest neighbor booms by calculating the rate at which each lender increases mortgage issuance in the focal market. The regression results in Column 2 and 3 indicate that on average, both lender similarity and lending amount largely are unaffected by the housing boom of the nearest neighbor.[36] Therefore, based on our two measures of lender activities, we did not find robust and material role of lenders in causing contagion.[37]

The fourth and fifth columns report analogous specifications using net migration between two MSAs (which uses IRS data on annual tax records) and focal market flippers (based on the fraction of transactions conducted by speculators) as the dependent variable. Both sets of results show no

---

[35]This lender similary index is calculated as $1 - \frac{1}{2}\sum_{l=1}^{L}|m_l - n_l|$, where m and n are the market shares of lender l in the focal and the nearest neighbor markets, respectively. A value of zero implies no similarity, while a value of one means that each lender has the same shares in both markets.

[36]We also found that lending amount by subprime lenders also does not respond to the neighbor's boom. If we restrict the analysis to the top 5 lenders with the highest lending amount for each market, we observe a jump in lender similarity index in Relative Year [1]. However, this result is not robust when we go from top 5 to top 10 or top 3 lenders. Including only top lenders also leads to a jump in lending growth rate in Relative Year [2] (the coefficient is around 0.5), but not in Relative Year [1]. However, lending growth also is higher before neighbor's boom, with coefficients on Relative Year [-1] and Relative Year [-2] being in the range of 0.2-0.4. Given this noticeable pre-trend, it is difficult to come to any conclusion that major credit lenders are the channel in disseminating positive housing market shocks.

[37]This average result, however, does not exclude the possibility of major credit lenders responding to neighbor's booms in other dimensions. The increase in similarity index among top lenders and the dip in their lending amount right at neighbor's boom may suggest some bank-level spillovers or substitution effects going on across markets. To fully understand the role of credit markets in this contagion context would require a closer examination of a full spectrum of lender behaviors, including those at the corporate level (e.g., mergers and acquisitions, and shocks to other business sectors of the lenders), which is beyond the scope of this paper.

discernible effect before or after the beginning of the nearest neighbor's housing boom.[38]

The fact that only focal market income shows any correlation with the timeline of the nearest neighbor's boom, especially its beginning, indicates that these fundamental factors are unlikely to be able to account for our estimated contagion elasticity. Table 3.6 provides additional support for this conclusion with alternative specifications similar to our baseline equation (3.3) that use (log) focal market price as the dependent variable. The first column reports results from a model that includes focal market fundamental controls and the log average price of near neighbors on the right-hand side, in addition to the standard controls from the baseline specification in Table 3.2. Our fundamental controls are local incomes, migration, subprime and FHA lending market shares, percentage of speculative buyers, percent minority, average LTV, average square footage, and the local unemployment rate. Note that these new estimates are very similar to the baseline results presented in Table 3.2. Controlling for these fundamentals does not change the magnitudes or time pattern of estimated contagion effects very much, indicating that the spatial spillovers are not being transmitted via the fundamentals that we consider here.

Thus far, we have abstracted from expectations of future fundamental factors, effectively treating actors as myopic. The second column of each panel in Table 3.6 begins to address this issue by adding four leads of own market income to the previous specification. Effectively, this presumes that local residents can fully predict the path of local incomes over the next four quarters. The inclusion of such stylized expectations does not change the estimated contagion effect. The third and final column of each panel reports results from adding four quarterly leads of all fundamentals, not just income. Once again, the magnitudes of the point estimates as well as the time pattern are relatively unchanged.

That we find no evidence that these spatial spillovers work through fundamental factors has po-

---

[38]We also investigated the shares of experienced and inexperienced speculators, respectively. Speculators are defined as experienced if having flipped at least four homes during the sample period, or inexperienced if otherwise. We did not find significant jumps in either experienced or inexperienced flippers when the closest neighbor begins its boom. However, the market share of experienced flippers declined by up to 0.5 percentage point (which is statistically significant) since two years after the neighbor booms, which is in line with those types of flippers being more sophisticated in timing the housing cycle and maximizing their return (Bayer et al., 2011).

tentially important implications for how policy makers should view intervention during housing booms. To the extent the booms are spread by non-fundamental forces (e.g., some type of Keynesian/Shillerian irrational exuberance), they might want to try to stop them from growing in scale and scope. Of course, fundamentals could encompass more than just income growth (or speculators, migration, etc.), and we would like to control for expectations as generally as possible. One extreme way that does eliminate the impact of contagion is to presume that the future path of price growth in the focal market is known with certainty in every period. In this context, the effect of a nearby housing boom on future prices in the focal MSA is immediately known and expectations are simply equal to the future value of price growth in the focal area. Adding four quarterly price growth leads completely wipes out the contagion impact. This indicates that the spillover could be operating primarily via expectations. Unless those expectations are based solely on fundamentals, the implications just discussed still hold. The likelihood that contagion operates via effects on expectations makes that issue an essential component of future research, but that clearly is beyond the scope of the current paper.

3.5. Additional Analysis: Instrumental Variable, Extensive Margin, and Bust

3.5.1. Alternative Model and Instrumental Variable

In this subsection, we relate price changes in the focal MSA with price changes from all neighbors, not just the nearest two. This direct estimation of the contagion elasticity has the benefit of allowing us to use an instrumental variable strategy to deal with omitted factors. Also, by interacting the price changes of the closest neighbor with a set of relative year dummies, we are able to explore how this effect varies over the course of a neighbor market's housing boom. The downside of this approach is that the specification does not allow us to fully observe the dynamic pattern of contagion, as we restrict the effect of neighboring price changes to operate through only one quarterly lag.

More specifically, we group neighboring locations into bins based on their distance from the focal

market and estimate the following equation:

$$\Delta P_{m,d,t} = \sum_{r=-4,r\neq0}^{4} \rho_r^1 * \psi_{m,r}^1 * \Delta \bar{P}_{m,t-1}^1 + \sum_{k=2}^{K} \rho^k * \Delta \bar{P}_{m,t-1}^k + \sum_{\tau=1}^{4} \rho_\tau * \Delta P_{m,d,t-\tau}$$

$$+ \sum_{r=-4,r\neq0}^{4} \rho_r * \psi_{m,r} + \gamma_{d,t} + \xi_{m,d,t} \tag{3.4}$$

where $\Delta \bar{P}_{m,t-1}^k$ is the lagged average price growth among neighboring MSAs falling into bin k for focal MSA m. In theory, we could allow each neighbor to be in its own bin based on how close it was to the focal area. However, that turns out not to be practical due to data limitations, so we bin neighbors based on distance rankings 1,2, 3–5, 6–10, 11–50, and 51+.[39] This makes the coefficients, ρ^k, the elasticity of focal area current price growth with respect to the *average* of lagged price growth among neighbors in bin k. For the closest neighbor, we further interact the lagged price growth variable with relative years to that neighbor's boom—resulting in a coefficient, ρ_r^1, for each of the neighbor's relative years, as shown in the first term on the right-hand side of (3.4). Relative year 0 again is the omitted category in all specifications. Thus, the coefficients on the lagged average price growth of the closest neighbor are interpreted relative to the effect in the 12-month period prior to that neighbor's boom.

One concern is that, even after including lags of the dependent variable (the third term on the right-hand side of (3.4)) and area-by-time fixed effects (the $\gamma_{d,t}$ vector), there still could be some common omitted factors helping drive the observed correlations. Ideally, we would like an instrumental variable that shifts the lagged average price growth of the focal MSA's closest neighbor, but does not directly affect the contemporaneous appreciation rate in the focal market itself. If taken literally, our estimating equation implies that further lags of the neighbor's price growth could potentially serve as an instrument because those variables would only affect the focal market's contemporaneous price growth through their impact on the lagged neighbor's price growth. This leads us to instrument for the lagged average price growth in each group of neighbors using one

---

[39]As in our main regressions, we include an indicator for whether we have price data for a given bin and interact this indicator with the relevant lagged average price variable. Results are qualitatively similar when we drop MSAs with missing neighbors' price data and also when we calculate distances using only the 99 MSAs for which we have data.

further lag of the average price growth among the relevant neighboring areas.[40]

Table 7 reports the results of estimating equation (3.4). Once again, there is a clear pattern that shows a shift in the importance of contagion right after the first year that close neighbors boom. Estimates for the first year of each of the neighbor's boom are 0.148, or approximately half the elasticity derived from our baseline estimates when considering the first year of the housing boom. Estimates for subsequent relative years fade relatively slowly, reaching a 0.1 elasticity that is similar to the baseline results.

3.5.2. A Hazard Model of Housing Boom Contagion

Our work above focuses entirely on the magnitude of contagion. In this section, we give empirical content to the extensive margin on the timing of booms that was suggested by Figure 3.1. We estimate simple hazard models to see if the probability of a focal market booming is influenced by the fact that neighbors boomed previously, after controlling for a host of covariates that also might account for the beginning of a boom.

Recent work on technology diffusion provides an intuitive way to generate contagion that is particularly appropriate for the specifications estimated in this subsection. In some of that research, contagion refers to a process in which people adopt a new technology when they physically meet with others who have already adopted it (Young, 2009; Comin et al., 2012). In our context of housing booms, this suggests that the probability of one metropolitan area entering a housing boom increases with its connection to other such areas that already have boomed. A standard assumption is that the intensity of the connection decays in distance between MSA pairs at a constant rate. Given that assumption, it is straightforward to generate the following two conclusions from a simple model: (a) if MSA q enters a housing boom at time t, that increases the hazard of MSA $m \neq q$ having a housing boom at time $t + 1$; and (b) this contagious effect is larger the closer is

---

[40]It turns out that the second quarterly lag of neighbor's price growth has a small and statistically insignificant impact on the focal MSA price appreciation, after controlling for all other covariates. In the specifications that interact the lagged average price growth of the neighbors with the focal MSA relative year indicators, we also interact lagged neighbors income growth and the second lag of neighbor's price growth with the focal area's relative year and include the full set of these interactions as instruments.

MSA m to MSA q.[41]

To investigate these implications, we begin by estimating whether the beginning of a boom in neighboring MSAs affects the hazard of focal MSAs entering a boom. We consider the following proportional hazard model relating the hazard of each focal MSA entering a boom in quarter t, $h_m(t)$, to a series of factors as noted in the following equation:

$$
h_m(t) = h_0(t) \exp \left\{ \sum_{k=1}^{K} \gamma_k \sum_{j \in N_m^k} Boom_{j,t-1} + \Delta X_{m,t} \beta + \sum_{\tau=1}^{4} \theta_\tau \Delta P_{m,t-\tau} \right.
$$
$$
\left. + \eta PctBoomed_{t-1} + \delta_R \right\}
$$

(3.5)

The primary coefficients of interest are γ_k's on the variable $\sum_{j \in N_m^k} Boom_{j,t-1}$, which reflects the number of MSAs among neighbors in bin k that began their boom in the previous quarter. We allow for multiple groups of neighbors (indexed by k) based on how close they are to the focal MSA. As before, we rank each focal MSA's neighbors and group them into K mutually exclusive bins based on these ranks, with N_m^k denoting the set of neighbors in bin k for focal MSA m. A positive coefficient γ_k would suggest a positive contagion effect. If the contagion effect decays with distance, we should also expect the γ_k's of closer neighbors to be larger than those of farther away neighbors. As with the estimation of magnitudes, we control for many potential correlates of a housing boom through the vector $\Delta X_{m,t}$.

We use the same bins as in Table 3.7. Table 3.8 then lists summary statistics on the number of booms started across different bins of neighbors. The very small means for the bins containing five or fewer markets document how unlikely it is that even a single boom began in any given quarter among the few markets in those bins. Thus, the lagged value of this variable ($\sum_{j \in N_m^k} Boom_{j,t-1}$) will contain many zeroes, making it difficult to estimate precise contagion effects about the timing of the beginning of local market booms. Nevertheless, we use these bins initially given the differential importance of the two closest neighbors in the contagion magnitude results and discuss findings

---

[41]See Appendix A.6 for a more formal presentation of how that approach generates specifications of the type we estimate in this section.

later that use larger bin sizes.

Our baseline results are from a common parametric specification, the exponential model, which assumes a flat baseline hazard $h_0(t) = \exp(\alpha_0)$. Implied changes in the probability of the focal area experiencing a boom this quarter (i.e., the hazard ratio) from one more neighbor having experienced a boom last quarter are reported in Table 3.9. The unconditional hazard, which is reported in column 1, indicates that the probability of the focal MSA booming this quarter more than triples if one more of its two physically closest neighbors boomed last quarter. The coefficients on the next three bins (Neighbors[3-5], Neighbors[6-10], and Neighbors[11-50]) are much smaller and none is statistically significant at conventional levels. That the lack of significance might be due to the nature of these variables having so little natural variation is indicated by the statistically significant impact of the group of furthest away markets (Neighbors[51+]). Still, it is lagged booms among the two closest neighbors that have the strongest correlation with a contemporary boom in the focal market.

Controlling for a standard set of covariates lowers the estimated effect substantially, as reported in column two. This time, the coefficients on the bins for all but the two closest neighbors are close to or below 1, indicating they have no positive impact on the probability of the focal market booming this period. And, the coefficient on the bin for the two closest neighbors falls by more than half. Even the 49% increase in probability implied by the coefficients on Neighbors[1-2] is not statistically significant, although that could be due to the nature of our data as discussed above. This conclusion is supported by the findings reported in column 3, which uses a Weibull hazard.[42] It reports a statistically significant correlation between lagged booms in very near neighbors and contemporaneous booms in the focal markets. These results show that if one of the two closest neighbors boomed last quarter, the probability of the focal market booming this quarter is about 70%–80% higher. The standardized marginal effect is smaller, of course, given the 0.18 standard deviation on Neighbors[1-2]. Using the middle of the range estimate from the Weibull hazard model, it is about 30% (i.e., $1.69 * 0.18$).

---

[42] A Weibull hazard model presumes a monotonic baseline hazard $h_0(t) = p t^{p-1} \exp(\alpha_0)$. We experimented with different functional forms to see if the pattern of results was materially affected.

These results only allow for a contagion effect from a single quarterly lag of neighbors' booms. We also estimated models that allowed for an increase in the number of booming neighbors over the past 6 to 12 months. There is a modest increase in the hazard (to 32%) if we allow for booms in any of the previous two quarters, but the result does not increase further if we allow for booms in any of the previous four quarters. Thus, the spatial effect estimated here appears to happen fairly quickly. This is also consistent with the elasticity estimates reported above.

The magnitude of the economic impact is hard to gauge on its own. We can gain some useful perspective by comparing it to the impacts on the hazard ratio of standard deviation changes in other variables. Among the underlying controls, the focal market's current income growth and previous quarter's price growth also were highly statistically significant predictors of a higher probability of booming. A one standard deviation higher own income growth rate is associated with about a 28% higher probability of the focal market booming. A one standard deviation higher rate of lagged house price appreciation is much more influential, as it is associated with an 87% higher hazard ratio. Thus, the standardized marginal effect of a boom in a very close neighbor appears to be quite influential and on a par with a standard deviation increase in its own income growth rate.[43] And, that we still find a meaningful influence of lagged near neighbors after controlling for everything else shows that the implications of the 'eyeball econometrics' from Figure 3.1 are not entirely due to many other potentially important factors.[44]

3.5.3. Was There Contagion During the Spread of the Housing Bust?

Our results indicate that contagion played a statistically and economically meaningful role in the timing and magnitude of the spread of housing booms across metropolitan areas. For completeness, we also explored the extent to which the same is true for the bust. In many respects, analysis

---

[43]The growth rate in the percentage of buyers with mortgages insured by the FHA also is a very powerful control. As expected, it is associated with a lower probability of booming. Specifically, a one standard deviation increase in that share is associated with an 86% fall in the hazard. Other statistically significant controls include the growth rate of the metropolitan area unemployment rate, as well as the second and fourth quarterly lags of focal market house price appreciation.

[44]Conditional hazard model estimates using more aggregate bins (Neighbors[1-10], Neighbors[11-50] and Neighbors[51+]) do not show any impact on the timing of the beginning of the boom in focal markets. Averaging across the ten closest neighbors masks the distinct impact of the two closest neighbors. As in the analysis of the magnitude of contagion, only the closest neighbors appear to matter.

of the bust is more challenging. One of the challenges is in deciding how it is defined and determined. We choose it to be the quarter in which nominal house prices peaked in the relevant MSA. While that may be intuitive, it also is much more *ad hoc* than our definition of the beginning of the boom, which is based on an external prediction of an economic model. Another challenge is that the bust was much more temporally correlated across markets than was the boom. This can be seen in Figure 3.3 which plots histograms of the quarters in which local market booms and busts began. This plot includes all markets that had a statistically significant boom. Note that the temporal concentration of busts is much greater, with every market experiencing a peak in prices within a 2.5-3 year period from mid-2005 through early 2008.

Figure 3.4 then provides more geographical detail with its plots of market busts over time. The first panel in this figure is identical to the first panel in Figure 3.1 and plots the MSAs for which we never estimate a statistically significant boom. The remaining panels show the timing of the bust among those MSAs that experienced booms. Unlike Figure 3.1's plots of the time line of metropolitan area booms, we see markets in all parts of the country, not just in coastal California and upper New England, with early price peaks between 2003-2005. The largest fraction of those peaks happening in last two quarters of 2005. Thus, the beginning of the bust is more national in scope than was the beginning of the boom. The subsequent plots in this figure do show a spreading out of the 'busts' to nearby markets. In the west, prices tended to peak earlier in interior markets and then spread to the coast. In Florida, the first price peaks were in markets on both coasts of that state. Peaking then occurred in a few other coastal markets before spreading to interior markets.

While this has the flavor of contagion seen in the start of the boom, more detailed analysis shows this not to be the case. For example, Table 3.10 reports hazard model estimates akin to those in Table 3.9, except here the dependent variable is the start of the bust, not the start of the boom. Unconditionally, lagged busts of neighbors are positively correlated with contemporary busts in focal markets, and near neighbors matter the most (column 1). However, column 2 of this table shows this conclusion of 'eyeball econometrics' from Figure 3.4 does not survive the inclusion of

covariates.[45]

In Table 3.11 we estimated price specifications for the bust akin to those in Table 2. First, using geographic distance, we do not see significant jumps in the magnitude of the contagion effect in relative year one. But focal prices decline by 1.2% and 3.8% in relative years two and three respectively. Results using economic distance follow a similar pattern. When compared to the magnitude of the price decline in the neighboring MSA, we find an elasticity of around 0.15. Those results are surprisingly similar to the ones observed during the spread of the housing boom.

3.6. Conclusion

We provide estimates of the role of contagion in the most recent American housing boom and bust. We find a statistically and economically meaningful role for contagion during the beginning and spread of the housing boom, and mixed evidence of contagion for the bust. Our key results are as follows. First, contagion impacts arise only from the very closest neighbors. There is no evidence of spillovers associated with more distant neighbors. The elasticity of focal market prices with respect to changes in its nearest neighbor's prices is in the range of 0.10-0.27. Back-of-the-envelope calculations suggest this is large enough to account for up to 30% of the jump in prices at the beginning of local booms, on average.

Finally, we found that local fundamentals and expectations of future fundamentals have very limited ability to account for our estimated contagion effect. That contagion transmission is not associated with local fundamentals suggests a potential role for non-rational forces. That is an issue in urgent need of new research because, if contagion does reflect some type of irrational exuberance, policy makers may want to rethink their past policy of not intervening to stop the spread of asset booms in housing.

---

[45]All markets are used in this particular estimation, including those that did not boom. The results are virtually identical if we restrict the sample to those that did boom.

No Boom

1997-1999

2000-2001

2002

2003

2004

2005
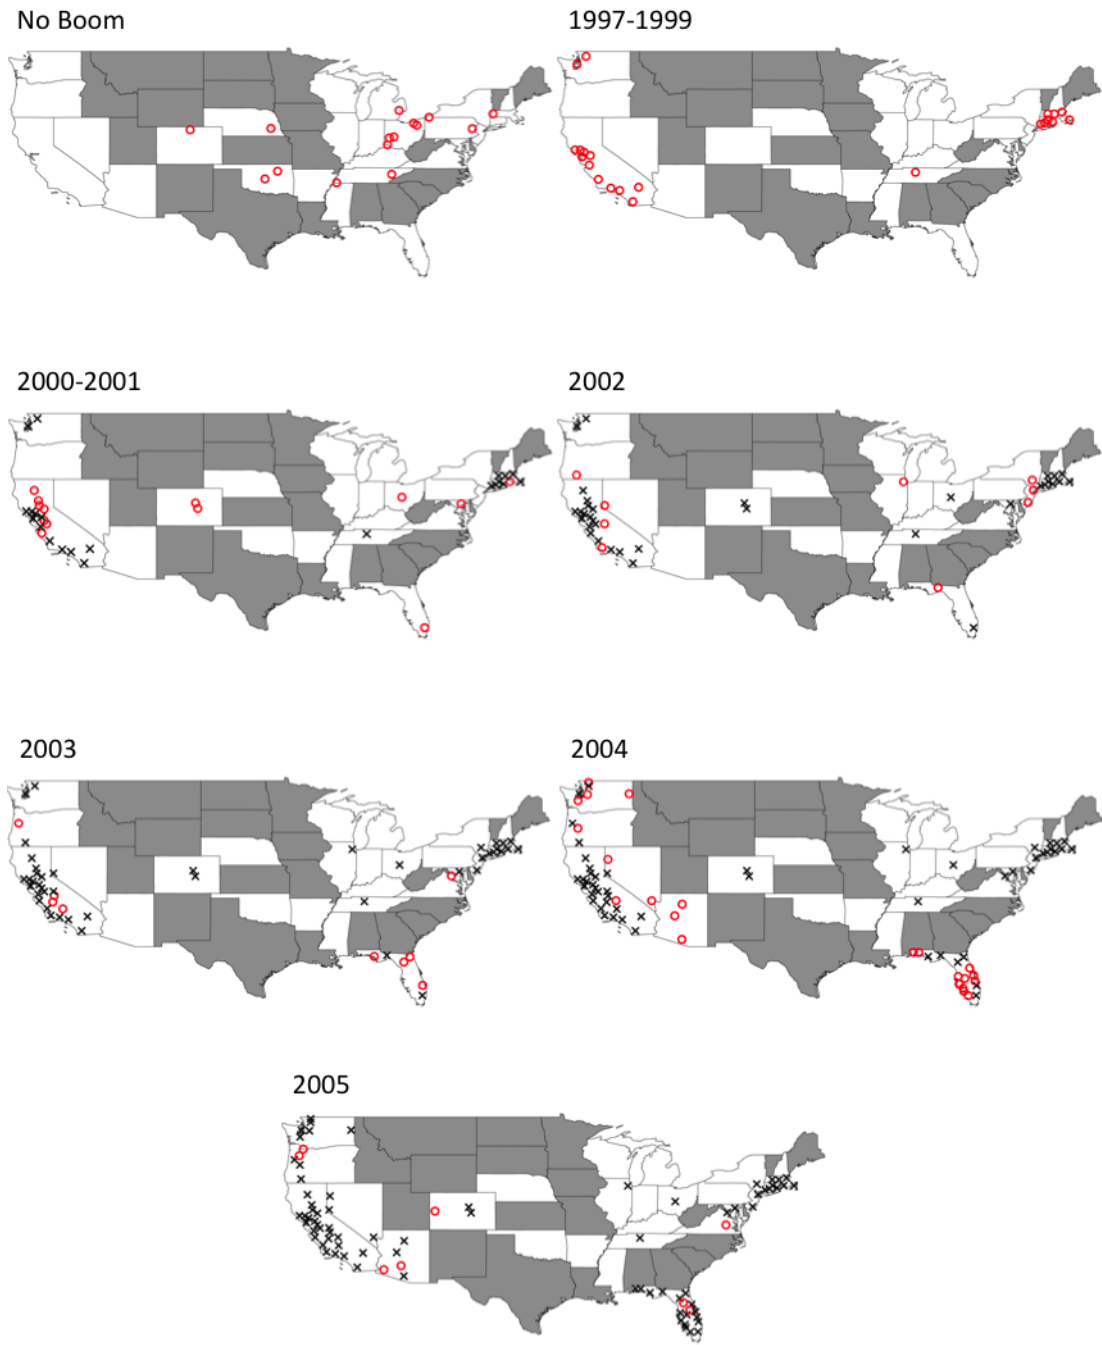
Fig. 3.1.—Timing of Housing Boom by MSA
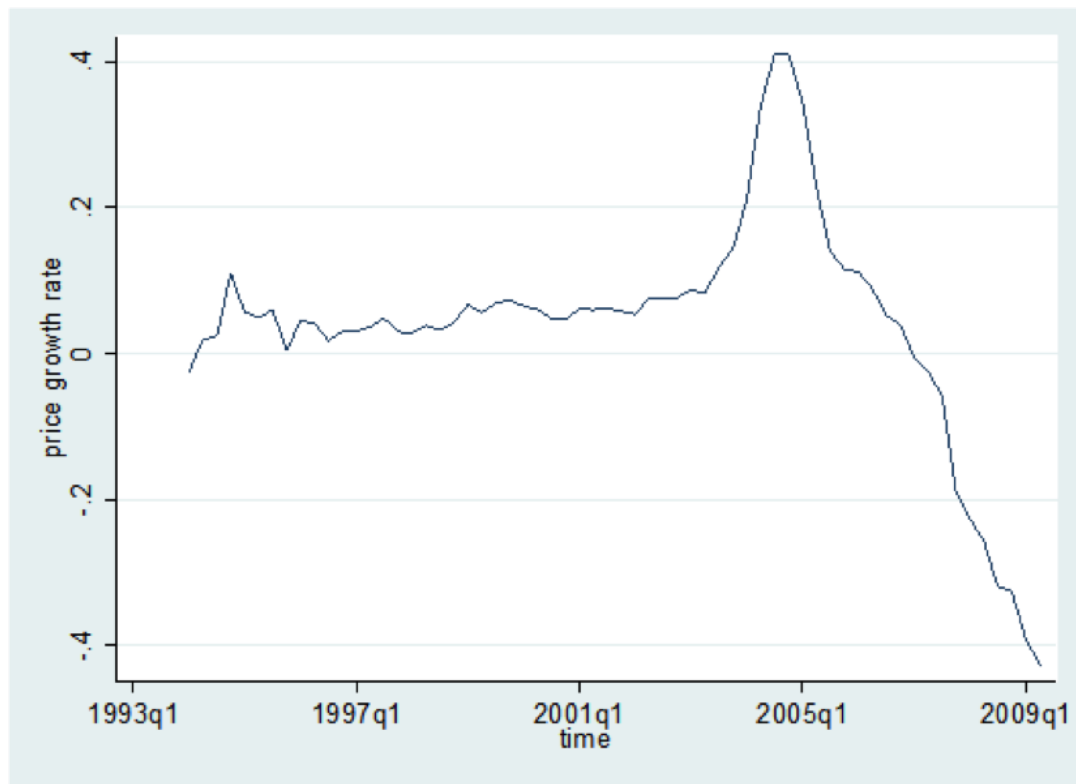
FIG. 3.2.—Las Vegas' Constant Growth Rate Before Booming. Source: Ferreira and Gyourko (2011).

FIG. 3.3.—Histograms of the Beginnings of Booms and Busts, MSAs

No Boom

2003-2005

2006(1)-2006(2)

2006(3)-2006(4)

2007(1)-2007(2)

2007(3)-2007(4)

2008-2009
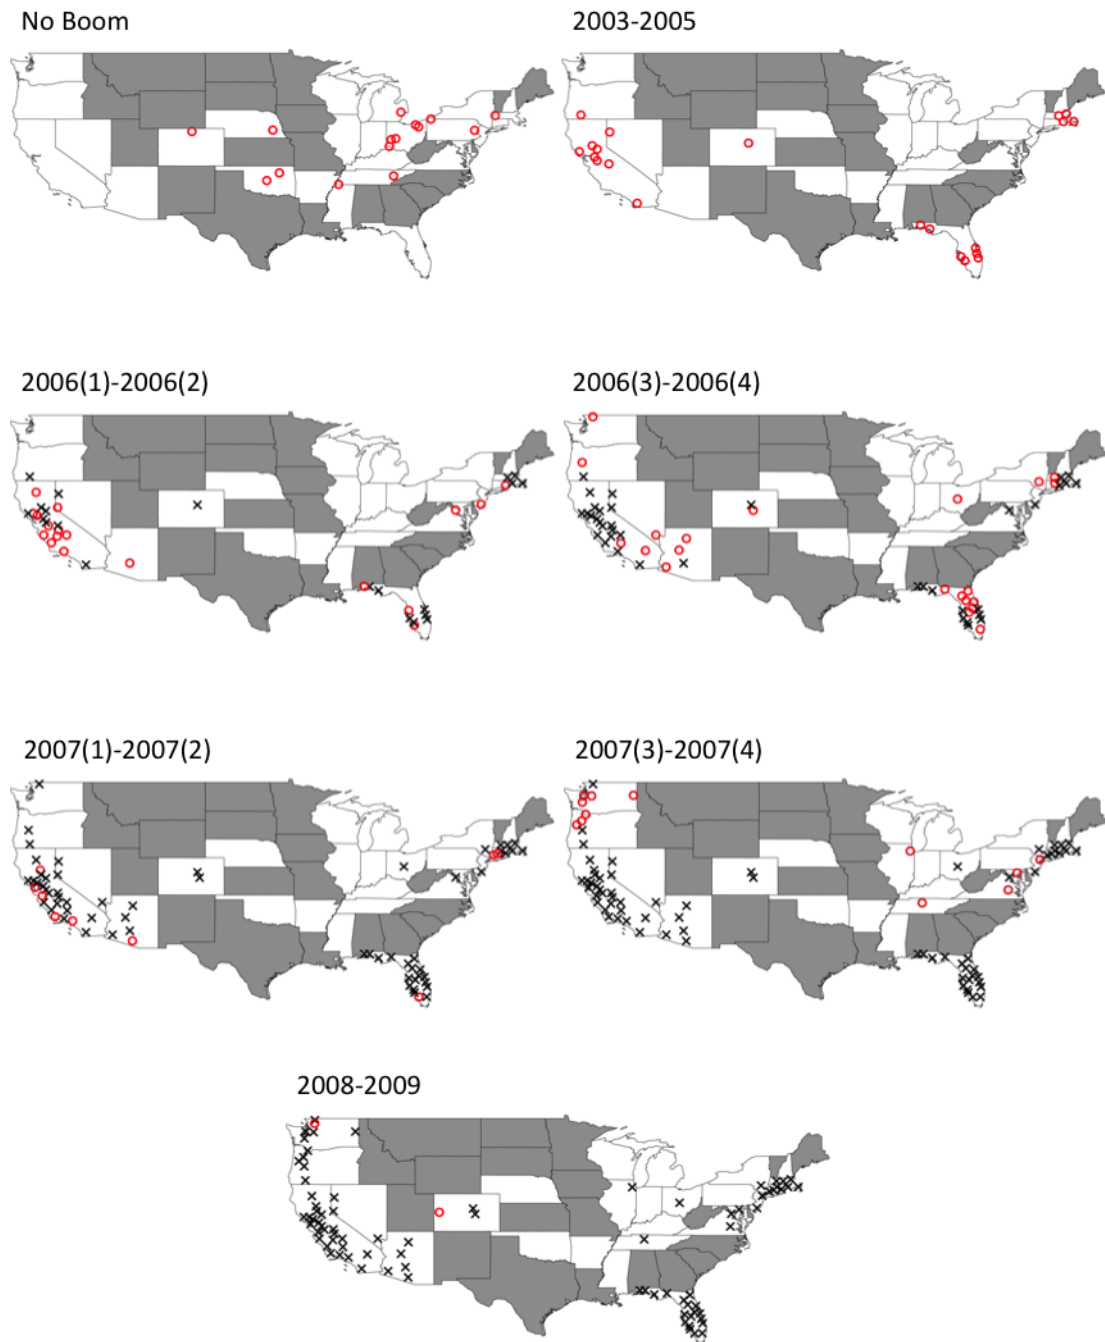
FIG. 3.4.—Timing of Housing Busts by MSA

139

TABLE 3.1
SUMMARY STATISTICS

	Mean	Std. Dev.	25th Percentile	75th Percentile
Price Index	132	52	96	156
Average Income ($1000's)	75	26	56	86
Percent Minority	.18	0.13	0.075	0.25
Percent Speculators	.054	0.032	0.031	0.071
Percent FHA Insured	.096	0.097	0.014	0.14
Percent Subprime Lenders	.14	0.075	0.087	0.19
Average LTV	.73	0.1	0.69	0.79
Average Square Footage	1672	161	1584	1764
Unemployment Rate	5.9	3.3	4	6.6
Net Migration	266	6878	-668	1408
N.	5043			

NOTE.—Columns present descriptive statistics for all MSA-quarter observations in our sample. Observation counts in regressions will vary depending on the specification and control variables used.

TABLE 3.2
THE IMPACT OF NEAREST NEIGHBOR HOUSING BOOM ON LOG FOCAL MARKET PRICE

Dep. Var: Log Focal Market Price	Neighbor 1		Neighbor 2	
Nearest Neighbors' Relative Years	Neighbor Boom Significant	Neighbor Boom Insignificant	Neighbor Boom Significant	Neighbor Boom Insignificant
Relative Year [-2]	-.003029	-0.003357	-0.00013	0.0114
	(.004662)	(0.01543)	(0.005425)	(0.008886)
Relative Year [-1]	-.000076	-.03365**	0.002323	-0.003854
	(.003184)	(0.01526)	(0.005658)	(0.006118)
Relative Year [1]	.008688**	0.00297	0.006096	0.000575
	(.004191)	(0.00563)	(0.006421)	(0.00634)
Relative Year [2]	.00977**	-.02102*	.01487***	0.002127
	(.004236)	(0.01091)	(0.005638)	(0.01006)
Relative Year [3]	.008542*	-.01698***	0.01029	0.001398
	(.004985)	(0.005502)	(0.006735)	(0.005476)
Relative Year [>=4]	.005247	-0.006412	0.004154	0.005613
	(.004774)	(0.004954)	(0.00551)	(0.005984)
Quarter-by-Census Division FE		Y		
Focal Market Relative Year FE		Y		
Four Lags of Focal Log Price		Y		
Neighbor-2 Relative Year FE		Y		
N		4584		

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest or second closest geographic neighbor. Relative year 0 indicates the 12 month period preceding the boom of the neighboring MSA and is the omitted category. Specification also includes dummy variable(s) indicating whether the closest neighbor(s) are in the sample. Standard errors are clustered at the Census division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.3
THE IMPACT OF NEAREST NEIGHBOR HOUSING BOOM
ON NEAREST NEIGHBOR LOG PRICE

Dep. Var: Log Nearest Neighbor's Price	
Nearest Neighbor Relative Years	
Relative Year [-2]	.02815**
	(.01333)
Relative Year [-1]	.01751
	(.0153)
Relative Year [1]	.03224
	(.02066)
Relative Year [2]	.06701***
	(.02199)
Relative Year [3]	.08197***
	(.02188)
Relative Year [>=4]	.1179***
	(.01976)
Quarter-by-Census Division FE	Y
Focal Market Relative Year FE	Y
Four Lags of Focal Log Price	Y
Neighbor-2 Relative Year FE	Y
N	3583

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest geographic neighbor. Relative year 0 indicates the 12 month period preceding the boom of the neighboring MSA and is the omitted category. Relative year effects are interacted with a dummy for whether the neighboring market had a statistically significant and positive break point and only the coefficients for the MSAs that actually had a boom are reported here. Specification also includes dummy variables indicating whether the closest neighbors are in the sample. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.4
HETEROGENEITY IN THE IMPACT OF NEAREST NEIGHBOR HOUSING BOOM ON LOG FOCAL MARKET PRICE

Dep. Var: Log Focal Market Price	Distance		Relative Size		Focal Market Elasticity	
	(1)	(2)	(3)	(4)	(5)	(6)
Nearest Neighbor Relative Years	<= 40 Miles	> 40 Miles	Focal Larger	Nbr. Larger	Inelastic	Elastic
Relative Year [-2]	-.004239	-.001194	-.001048	-.005334	-.001111	.001165
	(.004747)	(.005436)	(.004448)	(.005874)	(.006739)	(.007681)
Relative Year [-1]	-.001823	.001363	-.001269	.000373	-.003162	.003788
	(.003802)	(.004727)	(.003668)	(.003878)	(.004464)	(.006904)
Relative Year [1]	.008405	.008564	.00461	.01288**	.00769	.01877*
	(.005087)	(.006187)	(.00417)	(.005703)	(.007545)	(.0103)
Relative Year [2]	.008409	.009215	.01116***	.009466*	.008632	.02683**
	(.005425)	(.006344)	(.004208)	(.005423)	(.005739)	(.01069)
Relative Year [3]	.008389	.005968	.007858	.0104**	.00978	.02897***
	(.007093)	(.005006)	(.005809)	(.004913)	(.007775)	(.0102)
Relative Year [>=4]	-5.5e-06	.0118**	.005298	.008329	.01061*	.01055
	(.005004)	(.005328)	(.004551)	(.005509)	(.006347)	(.01232)
Quarter-by-Census Division FE	Y		Y		Y	
Focal Market Relative Year FE	Y		Y		Y	
Four Lags of Own Log Price	Y		Y		Y	
Neighbor-2 Relative Year FE	Y		Y		Y	
N	4647		4584		3586	

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest geographic neighbor interacted with dummies for whether the focal market is in the category indicated in the column header for various measures of heterogeneity. Relative year 0 indicates the 12 month period preceding the boom of the neighboring MSA and is the omitted category. All specifications also include dummy variables indicating whether the closest neighbors are in the sample. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.5

THE IMPACT OF NEAREST NEIGHBOR HOUSING BOOM ON FOCAL MARKET FUNDAMENTALS USING GEOGRAPHIC DISTANCE

	Dependent Variable				
Nearest Neighbor Relative Years	(1)	(2)	(3)	(4)	(5)
	Focal Market Income	Lender Similarity	Lending Amount	Net Migration	Focal Market Flippers
Relative Year [-2]	-0.013	.0127*	0.0205	84.43	.0078***
	(.009)	(.0065)	(0.0609)	(132.6)	(0.0029)
Relative Year [-1]	-0.006	0.005	0.0572	91.49	0.0039
	(.008)	(.0051)	(0.0562)	(132.9)	(0.0026)
Relative Year [1]	0.011	0.0025	0.0239	-43.23	-0.0004
	(.008)	(.0039)	(0.0505)	(85.75)	(0.0032)
Relative Year [2]	0.016	-0.0015	0.0491	-93.54	-0.0012
	(.0115)	(.0047)	(0.0416)	(78.01)	(0.003)
Relative Year [3]	0.0135	0.0004	-0.0230	-165.4	-0.0038
	(.0132)	(.0071)	(0.0448)	(114)	(0.0031)
Relative Year [>=4]	.0229*	0.0048	0.0194	-251.6**	-0.0037
	(.013)	(.0073)	(0.0618)	(101.4)	(0.0032)
Quarter-by-Census Division FE	Y	Y	Y	Y	Y
Year-by-Census Division FE	N	N	N	N	N
MSA FE	Y	Y	Y	Y	Y
Focal Market Relative Year FE	Y	Y	Y	Y	Y
Neighbor-2 Relative Year FE	Y	Y	Y	Y	Y
Focal Market Price and LTV Growth	N	N	Y	N	N
N	4877	4219	996941	1215	4829

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest geographic neighbor. Relative year 0 indicates the 12 month period preceding the boom of the neighboring MSA and is the omitted category. All specifications also include dummy variables indicating whether the closest neighbors are in the sample. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.6
THE IMPACT OF NEAREST NEIGHBOR HOUSING BOOM ON LOG FOCAL MARKET PRICE,
CONTROLLING FOR LOCAL FUNDAMENTALS AND EXPECTATIONS

Dep. Var: Log Focal Market Price			
Nearest Neighbor Relative Years			
Relative Year [-2]	-.003202	-.003899	-.002159
	(.004133)	(.003979)	(.004018)
Relative Year [-1]	-.000332	-.000577	.000934
	(.003053)	(.003016)	(.00288)
Relative Year [1]	.007821**	.008673**	.007878**
	(.003796)	(.003624)	(.003235)
Relative Year [2]	.008964**	.01033**	.008651**
	(.004278)	(.004157)	(.003704)
Relative Year [3]	.00804	.008275	.005047
	(.005421)	(.006464)	(.005994)
Relative Year [>=4]	.003512	.006228	.005626
	(.004924)	(.004178)	(.004507)
Quarter-by-Census Division FE	Y	Y	Y
Focal Market Relative Year FE	Y	Y	Y
Four Lags of Own Log Price	Y	Y	Y
Neighbor-2 Relative Year FE	Y	Y	Y
Focal Market Fundamental Controls	Y	Y	Y
Log Average Price of Neighbor Groups	Y	Y	Y
Four Leads of Own Income	N	Y	Y
Four Leads of All Fundamental Controls	N	N	Y
N	4538	4142	4142

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest geographic neighbor. Relative year 0 indicates the 12 month period preceding the boom of the neighboring MSA and is the omitted category. All specifications also include dummy variables indicating whether the closest neighbors are in the sample. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.7
THE IMPACT OF NEAREST NEIGHBOR PRICE CHANGES ON FOCAL MARKET PRICE CHANGES—IV
RESULTS

Dep. Var: Focal Market Price Changes		
Nearest Neighbor Price Changes	OLS	IV
Relative Year [-2]	0.139*	-0.045
	(0.072)	(0.228)
Relative Year [-1]	0.094**	0.041
	(0.042)	(0.131)
Relative Year [1]	0.135***	0.148***
	(0.038)	(0.037)
Relative Year [2]	0.102**	0.125***
	(0.040)	(0.035)
Relative Year [3]	0.090**	0.090**
	(0.039)	(0.037)
Relative Year [>=4]	0.089***	0.088***
	(0.027)	(0.028)
Quarter-by-Census Division FE	Y	Y
Focal Market Relative Year FE	Y	Y
Four Lags of Focal Log Price	Y	Y
Neighbor-2 Relative Year FE	Y	Y
N	4156	4131

NOTE.—Table reports the results from a regression of annualized focal market price growth rates on the lag of the nearest neighbor's annualized price growth interacted with dummies for the indicated relative years of the neighboring market. In the second column, lagged neighbor's price growth is instrumented using one further lag. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE 3.8
SUMMARY STATISTICS ON THE LAGGED NUMBER OF BOOMS FOR THE HAZARD ESTIMATION

	Lagged Number of New Booms				
	Mean	Sd	Min	Max	% Zero Boom
Neighbors [1-2]	0.0297	0.1777	0	2	97.2%
Neighbors [3-5]	0.0393	0.2058	0	3	96.2%
Neighbors [6-10]	0.0626	0.2692	0	3	94.3%
Neighbors [1-10]	0.1317	0.4246	0	5	89.0%
Neighbors [11-50]	0.3051	0.6079	0	7	76.2%
Neighbors [51+]	1.3055	1.6134	0	6	39.5%

NOTE.—Neighbors are ranked with respect to their geographic distance from the focal MSA.

TABLE 3.9

HAZARD MODEL ESTIMATES OF MSA NEIGHBORS ON THE PROBABILITY OF BOOMING BY GEOGRAPHIC DISTANCE

Independent Variables	Unconditional Hazard (1)	Baseline Results (Prop. Hazard) (2)	Weibull Hazard (3)
Lagged Number of New Booms			
Neighbors[1-2]	3.49***	1.49	1.69*
Neighbors[3-5]	1.31	0.79	0.84
Neighbors[6-10]	1.41	1.26	1.25
Neighbors[11-50]	1.14	0.86	0.8
Neighbors[51+]	1.23***	1.02	1.05
Other Controls			
Lagged % MSAs that Already Boomed	N	Y	Y
Focal Market Fundamental Controls	N	Y	Y
Four Lags of Focal Market Price Growth	N	Y	Y
Census Region FE	N	Y	Y
N	2114	2114	2114
Log-Likelihood	-29.93	29.43	34.14

NOTE.—Implied Hazard Ratios are reported along with indicators of statistical significance of the underlying regression coefficients. Significance levels 10%, 5% and 1% are denoted by *, ** and ***, respectively.

TABLE 3.10

HAZARD MODEL ESTIMATES OF MSA NEIGHBORS ON THE PROBABILITY OF BUSTING

Independent Variables	Unconditional Hazard (1)	Proportional Hazard with Controls (2)
Lagged Number of New Busts		
Neighbors[1-2]	1.66*	0.96
Neighbors[3-5]	1.14	0.74
Neighbors[6-10]	1.55**	1.08
Neighbors[11-50]	1.12*	0.96
Neighbors[51+]	1.18***	1
Other Controls		
Lagged Focal Market Fundamental Controls	N	Y
Four Lags of Focal Market Price Growth	N	Y
Census Region FE	N	Y
N	3637	3637
Log-Likelihood	2.181	100.1

NOTE.—We define the time of housing bust as the quarter when price peaks in our sample period. Implied hazard ratios are reported along with indicators of statistical significance of the underlying regression coefficients. Significance levels 10%, 5% and 1% are denoted by *, ** and ***, respectively.

TABLE 3.11

THE IMPACT OF NEAREST NEIGHBOR HOUSING BUST ON LOG FOCAL MARKET PRICE

Dep. Var: Log Focal Market Price	
Nearest Neighbor Relative Years	
Relative Year [-2]	-.01169**
	(.004509)
Relative Year [-1]	-.003846
	(.003101)
Relative Year [1]	-.006359
	(.005119)
Relative Year [2]	-.01228**
	(.006151)
Relative Year [3]	-.03804***
	(.007463)
Quarter-by-Census Division FE	Y
Focal Market Relative Year FE	Y
Four Lags of Focal Log Price	Y
Neighbor-2 Relative Year FE	Y
N	5178

NOTE.—Cells represent the coefficient on the dummy variable for indicated relative years of the closest geographic neighbor. Relative year 0 indicates the 12 month period preceding the bust of the neighboring MSA and is the omitted category. All specifications also include dummy variables indicating whether the closest neighbors are in the sample. Standard errors are clustered at the Census Division by year level and are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

APPENDIX

A.1. Additional Institutional Detail

TABLE A.1
MPDU INCOME LIMITS – 2014

Household Size	Adjustment Factor	Maximum Household Income	Minimum Household Income
1	0.70	$52,500	$35,000
2	0.80	$60,000	$35,000
3	0.90	$67,500	$35,000
4	1.00	$75,000	$35,000
5	1.08	$81,000	$35,000

NOTE.—This table shows the MPDU income limits for households of various sizes in 2014. For a four-person household, the maximum income limit is set at 70 percent of the area median income for the Washington, D.C. metropolitan area as published by the U.S. Department of Housing and Urban Development (HUD). That limit is then multiplied by the adjustment factor shown in the second column to determine the maximum income limits for households of other sizes. The minimum income limit is the same for all households and is set based on consultation with lenders in order to reflect the minimum income required to qualify for a typical mortgage on an MPDU home.

TABLE A.2
HISTORY OF MPDU CONTROL PERIOD RULES

Date MPDU Originally Offered for Sale	Control Period (Years)	Control Period Resets on Resale
Before October 1, 1981	5	No
October 1, 1981–February 28, 2002	10/15	No
March 1, 2002–March 31, 2005	10	Yes
After March 31, 2005	30	Yes

NOTE.—This table shows the history of MPDU control period rules from the inception of the program to the present. The rules governing the length of the control period and whether the control period resets upon resale prior to expiration are determined based on the date the MPDU was originally offered for sale by the developer. For MPDUs originally offered for sale between October 1, 1981 and February 28, 2002, the length of the control period was 10 years with the exception of MPDUs located in one of several DHCA designated "Annual Growth Policy Areas" for which the control period was 15 years.
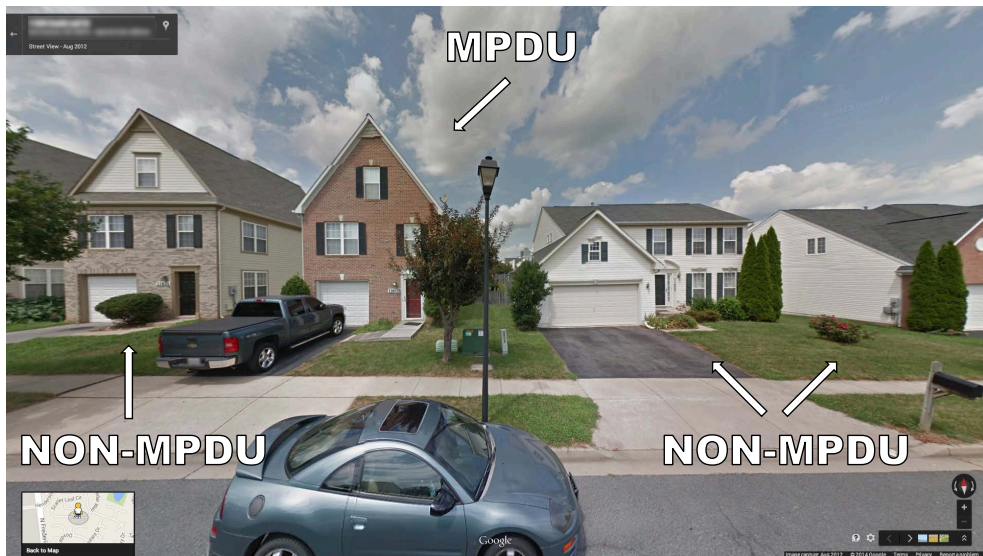
FIG. A.1.—Examples of MPDU Exterior Design. This figure presents images of the exterior of two representative MPDUs and nearby market-rate units. The MPDUs shown are located in different subdivisions. Images were accessed on June 24, 2014 and captured by Google Maps in May (top image) and August (bottom image) of 2012.

FIG. A.2.—Example Deed Restriction. This figure shows an example deed restriction for an MPDU originally sold by the developer on October 20, 1995. The deed was originally recorded in the Montgomery County Circuit Court (Land Records) MQR 13728, p. 0590, MSA_CE_63-13683 and was accessed online through MDLANDREC on June 24, 2014.

A.2. Data

A.2.1. Matching MPDU Properties to the DataQuick Assessment File

MPDU properties were matched to the DataQuick assessment file in several steps. Prior to matching, the raw addresses in both datasets were cleaned in order to correct obvious spelling errors, such as city names, and to standardize common abbreviations for street suffixes and compass directions (North, South, etc.). After being cleaned in this way, both sets of addresses were geocoded using an address locator service provided by the state of Maryland and developed in close collaboration with local jurisdictions that provides highly accurate geographic coordinates for addresses located throughout the state. Of the 8,289 MPDU properties, 91.8 percent were assigned an exact geographic location through this process. Many of the remaining 8 percent had expiration dates that occurred prior to 1980, suggesting that the reason they went unmatched was likely due to poor record keeping in the early years of the program. The DataQuick match rate was substantially higher. Over 99 percent of the properties were assigned unique geographic coordinates by the Maryland address locator. For the remaining unmatched DataQuick properties, I kept the geographic coordinates assigned by DataQuick, which uses a less accurate national address locator.

In the first step of the match, MPDU properties were assigned to DataQuick properties using exact geographic location. For each MPDU that was given a set of geographic coordinates, I first found the closest DataQuick property measured in straight-line distance. If the closest property was less than one foot away, the DataQuick property ID was assigned to that MPDU and considered a match. [1] All the remaining properties were considered unmatched and proceeded to the next step.

The remainder of the matching process used the actual address strings contained in both datasets. This part of the process proceeded in 13 iterations, re-matching unmatched addresses on increasingly lenient criteria at each step of the process. In the first step, properties were considered

---

[1] The Maryland geocoding service uses a "composite" address locator which looks to several sources to identify the geographic coordinates for a particular address. Because of this, it is possible for there to be slight differences in the geographic coordinates for the same property if it is identified using a different source. This is likely what generates distances that are less than 1 foot but still greater than zero.

matched if all components of the address string—house number, street name, unit number, zip code, and city—perfectly agreed. In the next two steps, previously unmatched properties were considered matched if all components of the address agreed except for one of either city or zip code (but not both). In the fourth step, properties were matched only if the entire address agreed, but some spelling error in the street name was accommodated by using the soundex code for the street name. [2] Steps five and six again allowed either city or zip code (but not both) to disagree while requiring an exact match on the soundex street name and the other components of the address. Steps 7–12 repeated steps 1–6, but allowing the unit number to differ at every step as long as one property had a missing unit number and the other did not. Explicit unit number disagreements were never permitted. In the final step, unit, city, and zip code were all allowed to disagree as long as the house number and street name perfectly agreed. At every step of the process, non-unique matches were randomly assigned.

Overall, the quality of the match is quite high. In total, 7,404 (90 percent) of the MPDUs were matched to a DataQuick property at some point in the process. Of these, only two percent were randomly assigned due to non-uniqueness. Eighty-two percent of matches occurred in one of the first two steps using either exact geographic location or the exact and complete address string. Ninety-five percent of the matches occurred in the first six steps, which required the unit numbers to agree.

A.2.2. Cleaning the Transaction and Loan-Level Data

The transaction and loan-level datasets were both cleaned in order to ensure that the transactions represent true ownership-changing arm's length transactions and that the loan information is accurate and consistent. The procedures for cleaning each dataset are described below.

---

[2]Soundex is a phonetic algorithm that indexes words based on their pronunciation. The goal of the algorithm to encode homophones in a similar manner. For example, "Willow Road" and "Wilow Road" have the same soundex code.

Cleaning the Transactions Data

The raw transactions dataset contains 249,264 transaction records that are coded by DataQuick as "arm's length" and involve one of the 286,484 single-family residential properties with non-missing housing characteristics contained in the assessment file. Starting with this sample, I first dropped 4,169 transactions where the year of sale preceded the year that the associated property was built (i.e. vacant land sales). A few transactions (197) recorded as having occurred in the last quarter of 2012 were dropped because the file provided by DataQuick was received in that quarter and only meant to cover up through the first three quarters of 2012. Many of the transactions recorded in 1997 were listed twice, with the first record containing the transaction price and no loan amount and the second record containing a loan amount with either no transaction price or a clearly erroneous transaction price (e.g. $1.00). In these cases, all of the other transaction characteristics were identical including the dates, buyer and seller names, and lenders. To correct this, I replaced the missing loan amount in the first record with the loan amount from the second record, and dropped the second record. This dropped 6,431 erroneous "transactions." A similar issue was present for a smaller number of transactions recorded in other years. In these cases, two transactions were recorded on the same day for the same property with transaction prices that differed by less than one percent and only one record containing a positive or realistic loan amount. In these cases, I kept the record with the non-erroneous loan amount and dropped the 2,918 duplicate records. A few transactions were exact duplicates of another transaction on property, date, price, and loan amounts, but differed along some other dimension (typically an alternate lender name). In these cases, one of the duplicates was randomly dropped. This dropped 375 records. In some cases, two transactions were recorded on the same day for the same property where the buyer for the first transaction was listed as the seller on the second transaction and the price for the first transaction was clearly not a market price. In these cases, the intermediary was typically a title company, escrow company, or some other similar entity and only the transaction containing the market price was kept (dropping 656 intermediary transactions). Finally, I randomly dropped 474 transactions that were exact duplicates of another transaction on property and date but differed along some

other dimension for unknown reasons as well as 165 transactions recorded on the same property in the same week but differing in some other way. In these cases, one of the records was kept to serve as a "placeholder" documenting the change in ownership that occurred on that day or during that week. These transactions were used in determining change of ownership, but their prices were not included in any analyses. The final sample contained 233,879 transactions involving one of the 286,484 single-family residential properties with non-missing housing characteristics contained in the assessment file.

Cleaning the Non-Purchase Loans Data

The non-purchase loans dataset was cleaned in a similar way as the transactions data. The raw dataset contains 780,927 non-purchase loans recorded on one of the 286,484 single-family residential properties with non-missing housing characteristics contained in the assessment file. Starting with this sample, I first dropped the 267 records with non-positive loan amounts. I also dropped 1,477 records where the loan amount was listed as $1.00. All of these records listed the lender name as "HUD," and were typically recorded on the same day as another loan with a more sensible loan amount. The DataQuick data contains a variable indicating whether a loan involved multiple parcels. This can happen when a real-estate investor owns multiple properties and borrows against their full portfolio using a single loan. I dropped all 4,130 loans coded in this way. As with the transactions data, I also dropped 3,074 loans recorded in years prior to the year that the associated property was built and 781 loans recorded in the last quarter of 2012. In cases where multiple loans were recorded on the same property on the same date with the same loan amount, I randomly kept one of the duplicates and dropped the remaining 743 records. Visual inspection of these records suggests that in most cases all other characteristics of the loan were also identical except for slight variations in lender name. Finally, I recoded 1,366 non-purchase loans as purchase loans if they occurred in the same week as a transaction recorded on the same property with no positive loan amounts. In these cases, the recoded loans were removed from the non-purchase dataset and recorded as the first loan on the associated transaction in the transactions dataset. The final sample of non-purchase loans contains 769,089 loans secured against one of the 286,484 single-

family residential properties with non-missing housing characteristics contained in the assessment file. These loans were used to construct the equity extraction measures used in the analysis.

A.2.3. Constructing "Debt Histories"

To accurately measure equity extraction, it is important to distinguish between three different types of non-purchase loans: (1) regular refinances, which replace an existing loan without extracting any equity; (2) cash-out refinances, which replace an existing loan with a *larger* loan, thereby extracting equity for the amount of the difference; and (3) new non-purchase originations, which directly extract equity for the amount of the new loan. In order to make this distinction, I construct a "debt history" for every property that records an estimate of the current amount of outstanding debt secured against the property at any point in time on up to two potential loans. Given this history, when a new loan is observed, I am then able to determine whether that loan represents a purchase loan, cash-out refinance, new non-purchase origination, or regular refinance by comparing the size of the new loan to the estimated outstanding balance on the relevant existing loan. This section describes the details of that procedure.

At each point in time, a given property can be thought of as having two potential "loan accounts," representing the current owner's first and second mortgage (I assume that owners carry at most two mortgages). The debt histories I construct are meant to estimate the remaining balance owed in each of these two accounts. The balances in the two loan accounts are initialized based on the first observed event in the property's history. For example, if the first observed event is a transaction with a $100,000 first loan and no second loan, then the balance in the first loan account will be initialized at $100,000 and the balance in the second loan account will be initialized at zero. If the first observed event is a non-purchase loan, then the balances are initialized based on a comparison of the loan amount with an estimate of the property's current resale value. Current resale values are estimated using quarterly constant-quality hedonic price indices constructed from the transactions data for each of 28 local planning areas designated by the Montgomery County Planning Department (see Appendix A.2.4 for details on how the price indices are constructed). These indices are used to adjust either the most recent transaction price or, for properties that

never transact, the 2011 assessed value to the relevant quarter. [3] If the loan amount is greater than 50 percent of the estimated current resale value, then the loan is used to initialize the first account balance and the second account balance is initialized at zero. If the loan amount is less than 50 percent of the current resale value, then the loan is used to initialize the second account balance and the first account balance is initialized at zero.

When a new transaction occurs, the balances in each account are replaced with the loan amounts associated with that transaction and a new ownership-spell is initiated. For non-purchase loans that occur between transactions, the balances in each loan account are updated based on a comparison of the new loan amount with the amortized balances remaining in the two accounts as of the date of the new loan. Since the DataQuick data does not contain information on loan terms or interest rates, all loans are amortized using the average offered interest rate on a 30-year fixed rate mortgage in the month that the loan was originated. Monthly average offered interest rates are taken from the Freddie Mac Primary Mortgage Market Survey (PMMS). Similarly, since the data does not distinguish between closed-end liens and HELOCs, all loans are treated as fully amortizing with an initial principal balance equal to the origination amount, which, for HELOCs, represents the maximum draw-down amount.

Several rules are used to determine whether the new loan updates the first loan account or the second loan account. When both accounts have a positive remaining balance, the new loan updates the account with the remaining balance that is closest to the new loan amount. If the new loan amount is at least five percent larger than the old loan, then the new loan is considered a cash-out refinance and replaces the old balance. [4] In this case, the difference between the two loans is counted as equity extraction. If the new loan is less than five percent larger than the old loan, then the new loan is considered a regular refinance and replaces the old balance, with no equity extraction recorded. If the second loan account has a zero remaining balance while the first loan

---

[3]Since prices for MPDU properties are not permitted to appreciate faster than the rate of inflation during the control period, I use the Consumer Price Index for All Urban Consumers (CPI-U) to adjust prices for these properties before the end of the control period and the hedonic price indices afterwards. Only transactions on the relevant side of the expiration date are used to derive the current resale price for MPDUs.

[4]The five percent threshold is chosen to reflect the cutoff used by Freddie Mac in its definition of cash-out refinancing.

account has a positive balance, then one of two things will happen. First, if the new loan is larger than 50 percent of the current first loan balance, then the new loan will replace the first loan and equity extraction will be determined using the same rules as above. Second, if the new loan is less 50 percent of the current first loan balance, then the new loan will replace the zero balance in the second loan account and be counted as a new loan origination. For new originations, the entire loan amount is counted as equity extraction. In the rare case in which there is a positive second loan balance and zero first loan balance the same rules are followed; in this case, however, the comparison is made with respect to the estimated current resale value rather than the first loan balance. Finally, if both loan accounts have a zero current balance, the new loan always replaces the first loan and is counted as a new origination.

Validating the Equity Extraction Measure

While the deeds data provide exhaustive coverage of all loans secured against a property, the assumptions needed to determine whether a new loan adds to or replaces existing debt introduce measurement error in the equity extraction variable. [5] To gauge the magnitude of this error, Figure A.4 presents aggregate time series evidence comparing my equity extraction measure against two external measures that were calculated using data from which it is possible to directly determine whether a new loan adds to a borrower's existing debt. Panel A. plots the yearly average probability of equity extraction. The dashed grey line plots a measure of equity extraction that was calculated by Bhutta and Keys (2014) using nationally representative data from the Equifax Consumer Credit Panel (CCP). The CCP data tracks individual debt obligations at a quarterly frequency and provides a near complete picture of the liability side of household balance sheets. Using this data, Bhutta and Keys (2014) define equity extraction as any instance in which an existing homeowner's total mortgage debt increases by more than five percent. The two solid lines were calculated using the DataQuick data and the measure of equity extraction discussed above. They plot the fraction of properties from which equity was extracted in each year for all proper-

---

[5]Under standard assumptions, such measurement error should only affect the precision of my estimates and not their accuracy. In particular, measurement error in the dependent variable does not introduce bias or inconsistency as long as the measurement error is uncorrelated with the explanatory variables. See, for example, Bound et al. (2001).

ties in Montgomery County (orange circles) and for the restricted set of properties in my analysis sample (blue squares). [6] The two methods of measuring equity extraction generate remarkably similar time series. Both the DataQuick series and the Equifax series increase rapidly during the period 1999–2003 before reaching a peak of roughly 20 to 25 percent and eventually declining and leveling off at around 5 percent by 2010. [7] The correlations between the Equifax measure and the two DataQuick series are also reported in the figure and are greater than or equal to 0.95 in both cases. Panel B. provides an alternative way of validating my measure—in this case, plotting the fraction of all refinance loans that in the DataQuick data that I code as cash-out. The dashed grey line plots a similar series taken from Freddie Mac's Quarterly Cash-Out Refinance Report. This series reports the share of all refinance mortgages in Freddie Mac's portfolio that were at least five percent larger than the loan they replaced. Again, the two methods for measuring the cash-out share generate very similar time series. Both the DataQuick series and the Freddie Mac series show clear cyclicality in the early and late 2000s, mimicking the cyclicality of interest rates during that period. The correlations in this case are slightly lower but nonetheless still quite high (0.87 for the full sample and 0.82 for the analysis sample). Taken together, the evidence presented in Figure A.4 suggests that any measurement error in my equity extraction variable is not substantial enough to affect the ability of that variable to accurately measure changes in equity extraction over time.

A.2.4. *Estimating Local House Price Indices*

The current resale values used to generate the "debt histories" for each property are estimated using quarterly constant-quality hedonic price indices constructed from the transactions data for each of 28 local planning areas designated by the Montgomery County Planning Department. I use planning areas to construct the house price indices because their boundaries are drawn in order to specifically take into account the homogeneity of interests, land use types, and local economic conditions of each respective area. [8] I use hedonic indices instead of repeat sales indices in order

---

[6]The criteria used to select the analysis sample are described in detail in Section 1.4.3.

[7]The series in Panel A. are only shown for 1999–2010 because Bhutta and Keys (2014) only report their measure for that period.

[8]Planning areas also have the added advantage that they cover the entire county and are large enough to provide enough data to reliably estimate a local house price index. The median planning area contains approximately 8,000 properties, which is about six times more than the median census tract.

to maximize the number of local indices available since repeat sales indices generally require much more data and therefore need to be estimated over larger geographies.

To construct the price indices, I begin by estimating the following hedonic regression:

$$\log(P_{ijmt}) = \alpha + X_i'\beta + \gamma_m + \psi_j \times \eta_t + \epsilon_{ijmt}, \tag{A.1}$$

where P_{ijmt} denotes the transaction price of property i, in planning area j, that transacts in calendar month m and quarter t, X_i is a set of property characteristics, γ_m is a set of calendar month fixed effects, $\psi_j \times \eta_t$ is a set of fully interacted planning area by quarter fixed effects, and ϵ_{ijmt} is the error term. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, dummies for the year the property was built, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with all the other characteristics. The property characteristics are included to control for changes in the composition of the transacted housing stock, while the calendar month fixed effects are included to control for the well-known seasonality of the housing market.

Having estimated this regression, I then obtain the (exponentiated) predicted values for each property, leaving out the contribution of the property characteristics and calendar month dummies. These predicted values, which are constant within planning area and quarter, are then used to construct the price index. Specifically, let $\widehat{P_{jt}}$ denote the predicted value for planning area j in quarter t. Then the price index for that planning area and quarter is given by

$$HPI_{jt} = 100 \times \frac{\widehat{P_{jt}}}{\widehat{P_{p0}}}, \tag{A.2}$$

where quarter zero is the base period used to normalize the index. Figure A.5 plots the price indices for all 28 planning areas normalized to 100 in the first quarter of 2000. In general, prices in Montgomery County evolved similarly to the national housing market over this period; however, there is substantial heterogeneity even within the county. Some rural areas barely saw any changes

162

in prices over this period, while prices nearly tripled in some of the more volatile areas of the county.

A.2.5. Matching the DataQuick Transactions Data to HMDA

To gauge the economic and demographic representativeness of my sample, I match a subset of the DataQuick transactions data to data on mortgage applications reported under the Home Mortgage Disclosure Act (HMDA) of 1975. HMDA requires lenders to report loan-level information on all loan applications received in a given year for home purchases, home purchase pre-approvals, home improvements, and refinances involving 1 to 4 unit and multifamily dwellings. This data is made publicly available by the Federal Financial Institutions Examination Council (FFIEC). I match the HMDA data to the transaction-level data from DataQuick using information on the primary loan amount, lender name, loan type (Conventional, FHA, VA), census tract, and year in which the transaction occurred. This section contains an overview of the matching process.

Prior to matching the data, I first selected a subsample of eligible transactions and loan applications based on two criteria. First, I restricted each dataset to include only transactions or loan applications pertaining to single-family homes with a positive first lien amount. [9] Second, I restricted the HMDA data to include only home purchase loan applications for which a loan was actually originated and presumably resulted in a completed transaction.

I then matched the data using a straightforward iterative process that proceeded in 6 steps, re-matching unmatched transactions and loans on increasingly lenient criteria at each step. In the first step, each transaction was matched to a loan using the year in which the transaction occurred, the census tract number, the loan type (Conventional, FHA, VA), the exact lender name, and the exact loan amount. [10] In cases where there were multiple matches, one of them was randomly

---

[9]HMDA only reports property type and whether a loan was a first or subordinate lien starting in 2004. However, this information is reported for all years in DataQuick. I restrict the DataQuick sample in all years and the HMDA sample in the years in which the information is available. There is no discernible difference in match rates or match quality between the years preceding and following 2004.

[10]HMDA rounds loan amounts to the nearest $1,000. Accordingly, I used a rounded version of the DataQuick loan amount to conduct the merge, but required that the rounded amounts agree exactly in both datasets. Census tracts were all converted to reflect boundaries as of the 2000 census using a crosswalk file provided by the U.S. Census Bureau.

assigned as being the true match while the rest were considered unmatched.

Remaining unmatched observations were then re-matched again based only on year, census tract, exact lender name, and exact loan amount, with multiple matches being randomly assigned as in the first step. The next two steps repeated the first two but used a truncated version of the lender name containing only the first 5 letters. The last two steps omitted the lender name entirely. Any observations remaining after this process are considered unmatched.

In total, 70.8 percent of the 233,870 transactions in the cleaned DataQuick file were matched at some point in the procedure. Of those, approximately 76 percent were matched in the first step requiring an exact match on the full lender name and including the loan type. Of the matched observations, 12.6 percent were randomly assigned due to having multiple matches.

To validate the quality of the match among matched observations, I further merged the surnames provided in the DataQuick transactions file with a Census generated list of the 1,000 most common surnames in the U.S. that also tabulates the percent of people with each name by race. For each race—Black, White, Hispanic, and Asian—I then grouped the Census percentages into 5 percent bins and used the race reported in HMDA to calculate the fraction of matched transactions in that bin for which the buyer had the same race. Since race was not used as a matching criteria, a strong positive correlation between the Census shares and matched HMDA shares would imply that the HMDA/DataQuick match does a good job of identifying the demographic characteristics of a particular home purchaser. Figure A.6 plots these correlations separately for each race. In each panel the blue circles plot the fraction of matched transactions in each 5 percent bin who report having the indicated race on their loan application. The black line in each panel plots the fitted values from a regression estimated in the underlying unbinned microdata, while the slope coefficient from that regression and its standard error are reported in the top left of each panel. The dashed orange line is the 45-degree line. For each race, the correlations are close to one and precisely estimated, suggesting that the match successfully identifies the demographic characteristics of individual home buyers.

164

A.2.6. Residential Building Permits Data

The data used to measure residential investment activity was obtained from the Montgomery County Department of Permitting Services and includes address-level information on all residential building, home improvement, and mechanical permits issued by the department since 2000. Each record includes information on the date the permit was applied for, the street address for the associated property, the type of work to be performed, and the type of structure on which the work will be performed. The permits data covers all areas of the county except for the cities of Gaithersburg and Rockville, which have their own permitting departments. The data covers both new construction and any major renovations, alterations, improvements, or additions to a home as well as any work performed on the heating, ventilation, and air-conditioning (HVAC) system. I drop any records where the listed structure type is clearly non-residential (e.g. "Restaurant," "Commercial," "Industrial") as well as any records for which the listed work type is not related to the actual improvement or alteration to the property (e.g. "Inspect and Approve", "Information"). The remaining dataset contains 148,647 unique permit applications, which I match to the DataQuick assessment file using the same approach used to match the list of MPDU addresses (see Appendix A.2.1) but allowing for multiple permits to match to the same property. Of the original 148,647 permits, a total of 137,510 (92.5 percent) were matched to a DataQuick property at some stage of the matching procedure. These permits were then used to construct the annual property-level panel used in the analysis as described in Section 1.7.
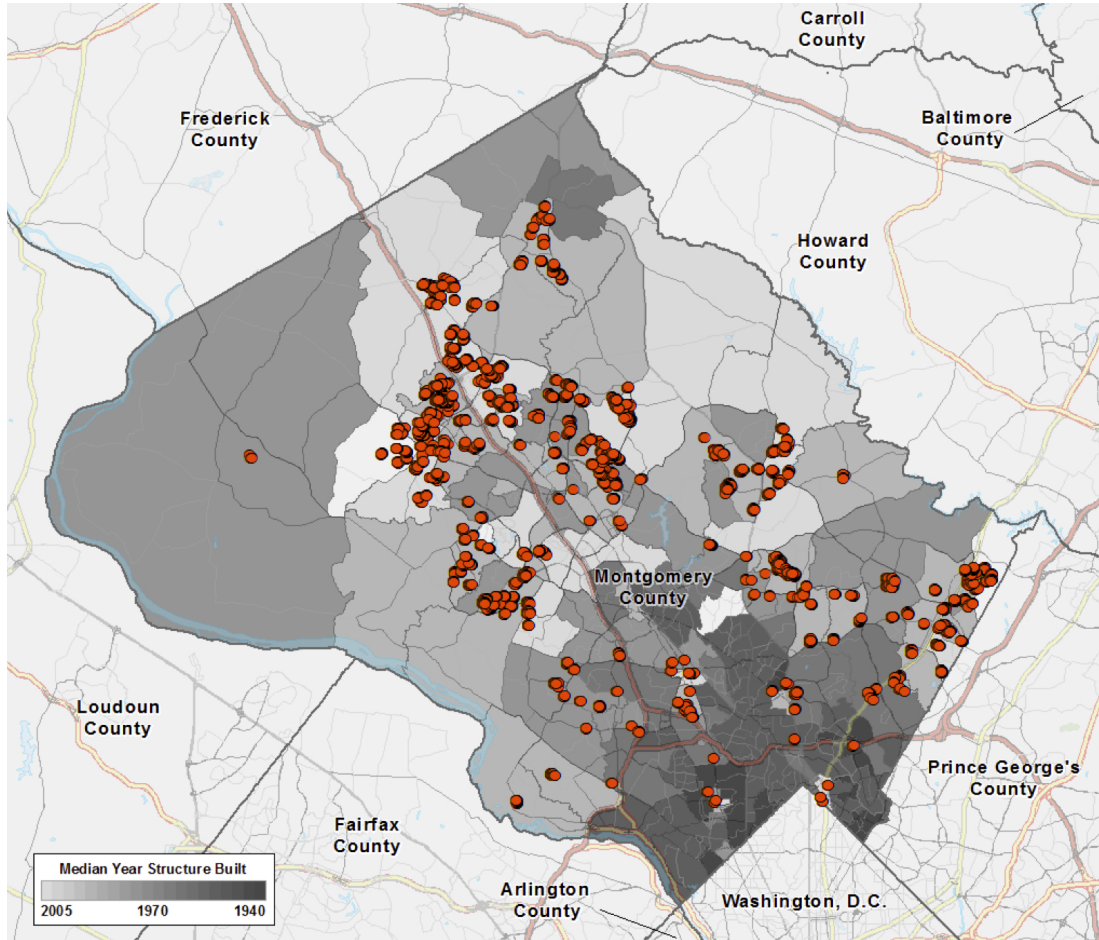
FIG. A.3.—Geographic Distribution of MPDU Properties within Montgomery County, Maryland. This figure shows the location of all MPDU properties that were successfully matched to a property in the DataQuick assesment file (N=7,404). MPDU properties are marked with an orange circle. Census tracts within Montgomery County are shaded according to the median year built for all housing units in the census tract as reported in the 2010 American Community Survey.
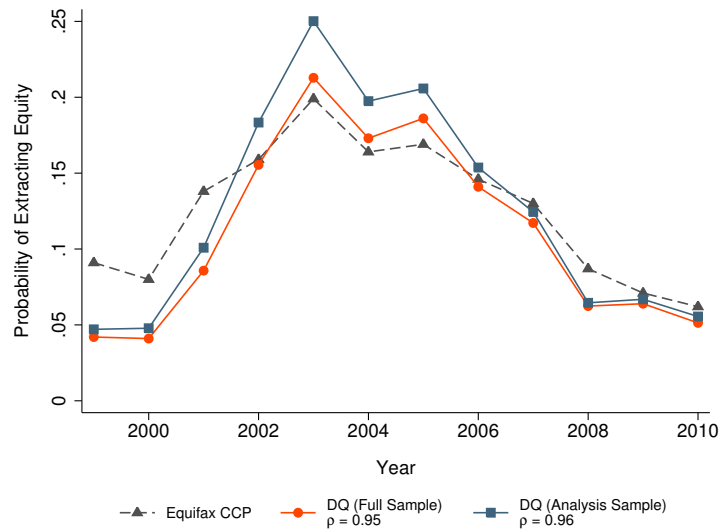
Panel A. Probability of Extracting Equity   *Panel B. Cash-Out Refinance Share*

FIG. A.4.—Validating the Home Equity Extraction Measure. This figure provides evidence validating the accuracy of the home equity extraction measure derived from the DataQuick deeds records against nationally representative aggregate series derived from other sources. Panel A. plots the yearly aggregate probability of extracting equity. Panel B. plots the yearly fraction of refinance mortgages that were cash-out. In both panels, the solid lines were generated using the equity extraction measure for Montgomery County derived from DataQuick for the full sample (orange circles) and the analysis sample (blue squares). In Panel A., the dashed grey line was taken from Bhutta and Keys (2014) and constructed using nationally representative borrower-level data from the Equifax Consumer Credit Panel. In Panel B., the dashed grey line was constructed using data from Freddie Mac's Quarterly Cash-Out Refinance Report. The correlations between the DataQuick measures and the corresponding national aggregate measures in each panel are reported in the legend. The time span in Panel A. is shorter than that of Panel B. because Bhutta and Keys (2014) only report their measure for the period 1999–2010.
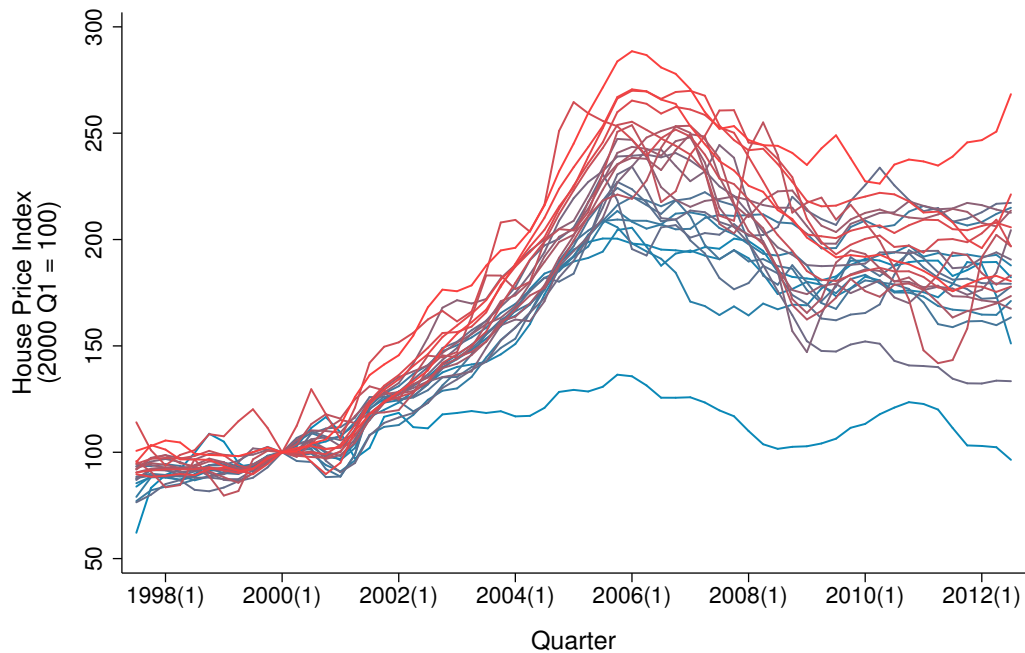
FIG. A.5.—Quarterly Hedonic House Price Indices for Montgomery County Planning Areas. This figure plots quarterly constant-quality hedonic price indices over the period 1997–2012 for each of the 28 local planning areas designated by the Montgomery County Planning Department. Each index was generated from the predicted values of a regression of (log) transaction price on a series of property characteristics, seasonal dummies, and planning area by quarter fixed effects as described in ??. All series are normalized to 100 in the first quarter of 2000 and are shaded according to their maximum value obtained over the entire period.
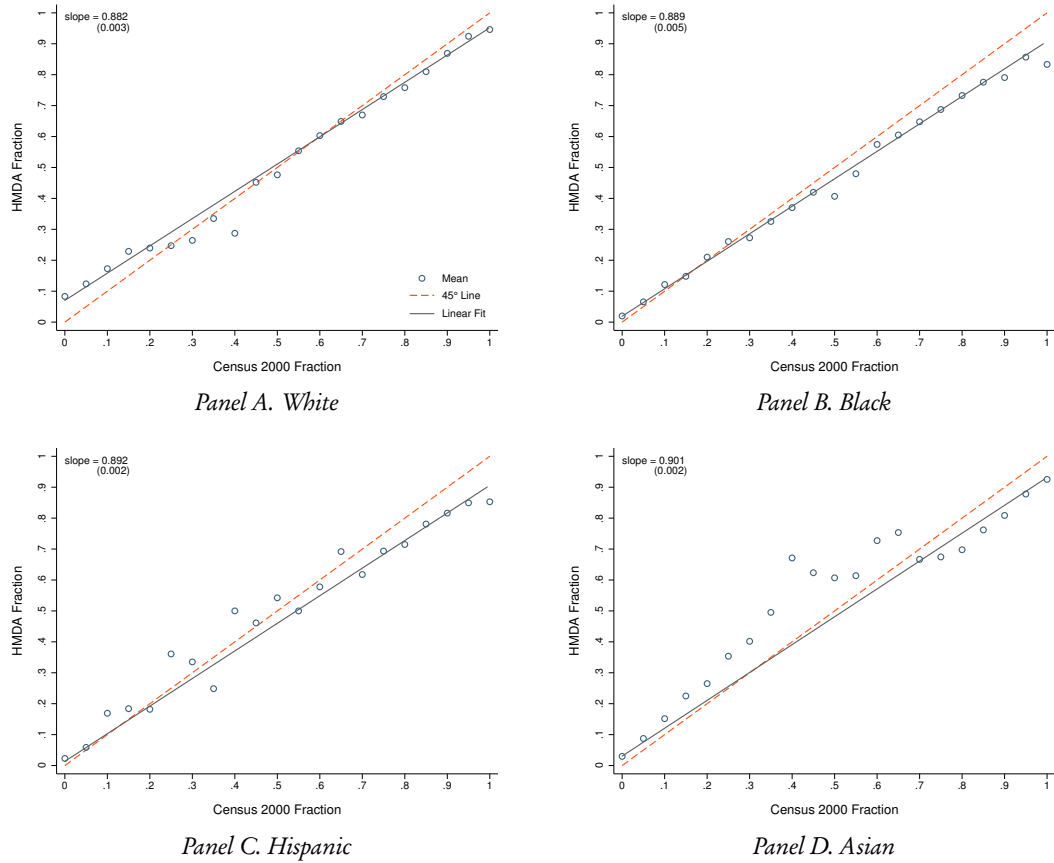
FIG. A.6.—Validating the DataQuick HMDA Match. This figure presents evidence validating the quality of the match between DataQuick housing transactions and HMDA loan applications. In each panel, the blue circles plot the fraction of matched transactions belonging to the indicated race as reported in HMDA on the y-axis against the fraction of households with the same surname as the home buyer who belong to that race as implied by the list of the 1,000 most popular surnames provided by the U.S. Census on the x-axis. The solid black line in each panel is the fit from a linear regression fit in the underlying microdata. The slope coefficient from that regression and its standard error are also reported in each panel. The dashed orange line is the 45-degree line.

Summary Stats in $1,000s
Mean: 107.06
Med: 79.98
SD: 101.14
IQR: [53.16, 122.99]

Panel A. National Sample

Summary Stats in $1,000s
Mean: 130.81
Med: 105.06
SD: 97.22
IQR: [75.39, 152.75]

Panel B. Full Analysis Sample

Summary Stats in $1,000s
Mean: 84.28
Med: 73.96
SD: 47.78
IQR: [55.29, 99.83]

Panel C. MPDU Only Sample

FIG. A.7.—Distribution of Homebuyer Income. This figure plots the distribution of (log) homebuyer income (in real 2012 $1,000s) as reported on loan applications contained in the HMDA data for three separate samples. Panel A. shows the distribution for the set of all approved purchase-mortgage applications filed in the United States between 1997 and 2012 for which the borrower reported a non-missing income ($N = 62,267,113$). Panel B. restricts the sample to include only mortgage applications that were successfully matched to a transaction contained in the primary analysis sample ($N = 21,696$). Panel C. further restricts the sample to include only applications that matched to a transaction in the analysis sample that involved an MPDU property ($N = 1,523$). Each panel also reports the mean, median, standard deviation, and interquartile range in levels for the plotted log income distribution. A bin width of 10 log points is used in all three panels.
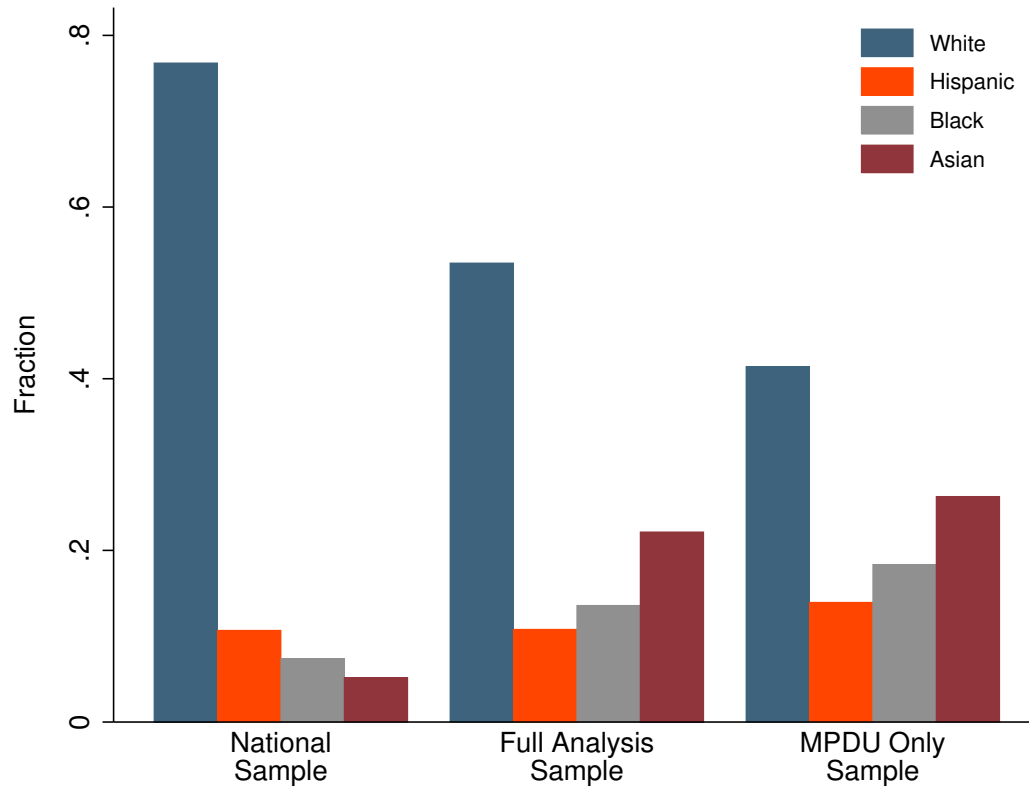
FIG. A.8.—Distribution of Homebuyer Race. This figure plots the distribution of homebuyer race as reported on loan applications contained in the HMDA data for three separate samples. The first set of bars shows the racial breakdown for the set of all approved purchase-mortgage applications filed in the United States between 1997 and 2012 for which the borrower reported a non-missing race or ethnicity of either White, Black, Hispanic, or Asian ($N = 57,710,449$). The second set of bars restricts the sample to include only mortgage applications that were successfully matched to a transaction contained in the primary analysis sample ($N = 19,780$). The third set of bars further restricts the sample to include only applications that matched to a transaction in the analysis sample that involved an MPDU property ($N = 1,427$). Racial shares are calculated only within the sample of applications with one of the indicated races so that the height of the bars adds to one within each sample. The four categories are mutually exclusive, meaning that a borrower is categorized as Hispanic if she reports an ethnicity of Hispanic regardless of which race she reports.

A.3. Additional Results and Robustness Checks

A.3.1. The Effect of Expiring Price Controls on MPDU Turnover

One potential concern with implementing a difference-in-differences research design at the property level is that in addition to the increase in collateralized debt capacity, the expiration of the price control also creates incentives for MPDU owners to sell their homes, which could lead to a differential increase in turnover at MPDU properties. While this concern is explicitly addressed in the main analysis through the inclusion of property and ownership-spell fixed effects, it is nonetheless interesting to empirically gauge the magnitude of any changes in turnover induced by expiring price controls.

To do so, I construct an annual property-level panel which records for each property in the main analysis sample whether that property was sold in a given year. For properties built prior to 1997, the panel covers the full sample period from 1997–2012; for properties built afterwards, the construction year is used as the first year of observation. Using this panel, I then estimate regressions of the following form:

$$Sold_{ist} = \alpha_s + \delta_t + X'_{it}\gamma + \beta_1 \cdot MPDU_i + \beta_2 \cdot MPDU_i \times Post_{st} + \epsilon_{ist}, \qquad \text{(A.3)}$$

where $Sold_{ist}$ is an indicator for whether property i in subdivision s was sold in year t, and all other variables are as described in Section 1.5 in reference to equation (1.10). The coefficient of interest is β_2, which measures the differential change in the turnover rate for MPDUs relative to non-MPDUs following the expiration of the price control, holding constant individual housing characteristics and aggregate differences in turnover rates across subdivisions and over time.

Table A.3 presents results from estimating this regression using various specifications. In the first column, I include only time-invariant property characteristics and fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and

the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. In the second column I also include fixed effects for the subdivision the property is located in. In column 3, I further interact the subdivision fixed effects with a linear time trend to allow for differential aggregate trends in turnover across subdivisions. In the fourth column, I include property fixed effects, causing the time-invariant property characteristics and the MPDU main effect to drop out. In this specification, the effect of expiring price controls is identified by comparing within-property changes in turnover probabilities for properties that are and are not MPDUs. Finally, in columns 5 and 6, I dispense with the linear probability model and report probit and logit marginal effects using the same specification as in column 3.

The estimated effects are relatively stable across specifications and imply that expiring price controls lead to an increase in the annual turnover rate at MPDU properties of roughly three to five percentage points. These effects are large relative to the pre-period average annual turnover rate of 4.8 percent among MPDUs reported in the bottom panel of the table. Comparing the MPDU main effect with the interaction term shows that expiring price controls close between 50 to 80 percent of the gap in turnover rates between MPDUs and non-MPDUs that existed during the period of price control. However, the estimates are small in absolute terms, reflecting the fact that turnover is a relatively rare event. For example, adding the estimated three percentage point increase in column six to the pre-period mean turnover rate among MPDUs implies an annual post-expiration turnover rate of 7.8 percent. At that rate, it would take almost 13 years for the entire MPDU housing stock to turnover.

To give a sense of the dynamics of the turnover effect, Figure A.9 plots estimates from a version of equation (A.3) that allows the effect of the price control to differ separately for MPDUs and non-MPDUs by year relative to the first control period expiration (as described in the discussion of equation (1.11) in Section 1.5). Specifically, the series in orange squares plots the coefficient estimates on a set of dummies indicating whether the year of observation falls in a given *relative* year as measured from the year the first MPDU in the relevant subdivision expired (relative year

173

zero). This series measures the trend in the turnover rate for non-MPDU properties around the time the price control expired. Relative year -1 is the omitted category so that all estimates should be interpreted as relative to the year prior to when the first price control in the subdivision expired. Similarly, the series in blue circles shows the trend for MPDU properties. This line plots the sum of the relative year main effects (the series in orange squares) and the interaction of those effects with an indicator for whether the property is an MPDU. The figure also reports the 95 percent confidence interval for that sum. All controls included in column 3 of Table A.3 were also included in the regression. As the figure makes clear, the turnover rate at MPDU properties exhibits a sharp departure from it's pre-period trend precisely in the year the first price control expires while there is no corresponding change for non-MPDU properties. Moreover, the trends for MPDUs and non-MPDUs are statistically indistinguishable in the period prior to the expiration of the price control and only diverge beginning in the year of expiration.

A.3.2. *Propensity Score Matching Estimates of the Borrowing Response*

Another potential concern with the main difference-in-differences estimates provided in Section 1.6 is that they rely on standard OLS estimation, which can be sensitive to differences in the distribution of covariates across treatment and control groups and relies heavily on extrapolation in areas where the covariates do not overlap (Imbens, 2004). In this section, I explore the sensitivity of the main results to an alternative estimation approach which restricts attention to the set of properties with overlapping characteristics and constructs the counterfactual outcome for each MPDU property using a locally weighted average of the outcomes among the non-MPDU properties whose characteristics are most similar. Specifically, I provide estimates based on the local linear propensity score matching difference-in-differences estimator developed in Heckman et al. (1997, 1998).

Adopting the notation in Smith and Todd (2005), let t' and t denote the time periods before and after the expiration of the first price control within a property's subdivision. Let Y_{1ti} denote the observed outcome for property i if it receives the "treatment" in period t, where here the treatment is defined as being an MPDU in the post-period. Similarly, let Y_{0ti} denote the out-

come for property i without treatment. Further, let the dummy variable $D_i = 1$ if a property is an MPDU and $D_i = 0$ if a property is not an MPDU. Given a vector of fixed property characteristics, X_i, the propensity score is defined as the conditional probability that a property is an MPDU, $P(X_i) = \Pr(D_i = 1 | X_i)$. A property is said to be in the region of common support, S_P, if its propensity score has positive density in both the MPDU and non-MPDU distributions of propensity scores: $S_P = \{P : f(P | D = 1) > 0 \text{ and } f(P | D = 0) > 0\}$.

Having established this notation, the propensity score matching difference-in-differences estimator can be expressed as:

$$\hat{\Delta}_{D=1}^{DID} = \frac{1}{n_{1t}} \sum_{i \in I_{1t} \cap S_P}^{n_{1t}} \left\{ Y_{1ti}(X_i) - \hat{E}\left(Y_{0ti} | P(X_i), D_i = 0\right)\right\} - $$
$$\frac{1}{n_{1t'}} \sum_{j \in I_{1t'} \cap S_P}^{n_{1t'}} \left\{ Y_{0t'j}(X_j) - \hat{E}\left(Y_{0t'j} | P(X_j), D_j = 0\right)\right\}, \qquad (A.4)$$

where $I_{1t'}$ and I_{1t} denote the set of MPDU properties with outcomes observed in the pre- and post-periods, respectively, and n_{1t} and $n_{1t'}$ are the number of observations for MPDU properties in those two sets that are also in the region of common support. Implementing this estimator requires determining the region of common support and estimating the two expectations $\hat{E}(Y_{0ti} | P(X_i), D_i = 0)$ and $\hat{E}\left(Y_{0t'j} | P(X_j), D_j = 0\right)$, which serve as the counterfactual outcomes for MPDU properties in the two periods.

Both the determination of the region of common support and the estimation of the counterfactual outcomes depend on the propensity score, which I estimate using a simple probit model. Specifically, I take the estimated propensity score for property i to be the fitted values from a probit regression of the MPDU dummy on a set of property characteristics which includes the interior square footage of the home, the year it was built, the number of bathrooms and stories in the home and an indicator for whether the property is a condo or townhome. In order to get an accurate prediction of the propensity score, these covariates are entered into the model in a highly flexible fashion. I include cubic splines in both the square footage and year built as well as the linear interaction of both variables with a fully interacted set of dummies for the number of bathrooms and

the number of stories. All of these terms are then further interacted with the condo dummy.

Having estimated the propensity score for each property, I then define the region of common support as the set of all propensity scores that are larger than the maximum of the first percentile in the distribution of propensity scores in both sets of properties and smaller than the minimum of the 99th percentile of propensity scores in both distributions. Figure A.10 plots the distribution of propensity scores for properties that fall within the region of common support separately for MPDUs and non-MPDUs. Dropping properties outside the region of common support leaves a total of 1,836 MPDU properties and 8,443 non-MPDU properties. Table A.4 provides an assessment of how well the estimated propensity score does in balancing covariates across these properties. The table shows the means of the covariates used to estimate the propensity score separately for MPDUs and non-MPDUs within terciles of the combined propensity score distribution. For each set of means, I also report the t-statistic from the test of the null hypothesis of no difference in means. While the propensity score does not do a perfect job of balancing the covariates, as is evidenced by the statistically significant differences for several of the variables, in nearly all cases, the differences in means are not economically meaningful and are far less stark than the differences in the full sample shown in columns 5 and 7 of Table 1.1.

I use a local linear regression estimator to construct the matched counterfactual outcomes for each observed MPDU outcome. Implementing this estimator is relatively straightforward, and the full details can be found in Todd (1999). I focus the discussion here on how I construct the counterfactual outcome for observed MPDU outcomes in the post-period, Y_{1ti}. The process for constructing counterfactual outcomes for the pre-period outcomes, Y_{0ti}, is completely analogous. To construct the matched outcome for a particular MPDU property, i, observed in the post-period, I first match the post-period outcome for that property to all observed post-period outcomes among non-MPDU properties. I then calculate the difference in propensity scores between the MPDU property and each of the matched non-MPDU properties. For a particular non-MPDU property, j, denote this difference as $P(X_i) - P(X_j)$. I then run a weighted least squares regression of the outcomes for the matched non-MPDU properties on a constant and a linear term in this differ-

ence. I weight each observation according to the difference in propensity scores using a quartic kernel function and an bandwidth of 0.1. This means that outcomes among non-MPDUs whose propensity scores are more than 0.1 away from the propensity score for MPDU i receive no weight in the regression while those with identical propensity scores receive a weight that is close to 1. The estimated counterfactual outcome, $\hat{E}(Y_{0ti}|P(X_i), D_i = 0)$, is given by the constant from this this regression. This process is then repeated for all observed MPDU outcomes in both periods in order to obtain all of the matched outcomes.

With the matched outcomes in hand, I can then directly calculate the propensity score matching difference-in-differences estimate given by equation (A.4). To calculate the standard error for this estimate, I bootstrap the entire process using a stratified resampling procedure that randomly samples properties with replacement in a way that ensures that the number of properties sampled from each subdivision stays the same. Table A.5 reports the matching estimates for each of the three main outcomes—log transaction prices, the annual probability of extracting equity, and the total amount of equity extracted per year. The estimated effect for all three outcomes is positive and precisely estimated. The equity extraction estimates are almost identical to the main OLS difference-in-differences estimates reported in Table 1.3 and Table 1.4. Similarly, the price effect is slightly larger but qualitatively similar to the main estimates reported in Table 1.2. Together, these results suggest that the main estimates are not being greatly affected by the fact that OLS relies on extrapolation in regions of the covariate space with poor overlap.

FIG. A.9.—Dynamic Effects of Expiring Price Controls on Turnover at MPDU Properties. This figure reports estimates of the effect of expiring price controls on the annual turnover rate at MPDU properties derived from a flexible difference-in-differences regression that allows the effect to vary by year relative to the expiration of the price control. Estimates were constructed by regressing an indicator for whether a given property sold in a particular year on an indicator for whether that property is an MPDU and the interaction of the MPDU indicator with a series of dummy variables indicating whether the year of observation falls in a given relative year as measured from the year the first MPDU in the relevant subdivision expired. Relative year zero denotes the year the first price control in the subdivision expired. Relative year −1 is the omitted category so that all estimates should be interpreted as relative to the year prior to expiration. Results are shown for five years preceding and following the expiration of the price control, with all years outside that window grouped into the effects for relative years −5 and 5. The series in orange squares plots the coefficient estimates on the relative year main effects, which represent the trend in turnover rates among non-MPDU properties. The series in blue circles plots the estimate and 95 percent confidence interval for the sum of the relative year main effects and the interaction of those effects with the MPDU indicator, representing the trend among MPDU properties. The 95 percent confidence intervals are based on standard errors which were clustered at the subdivision level. The regression also included year fixed effects, subdivision fixed effects and their interaction with a linear time trend and a set of property characteristics. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms, stories, and property age, as well as an indicator for whether the property is a condo or townhome and the interaction of that indicator with the year fixed effects and all of the other property characteristics.

FIG. A.10.—Propsensity Score Overlap. This figure shows the overlap in the distribution of estimated propensity scores for MPDU and non-MPDU properties in the region of common support. The propensity score was estimated using a simple probit regression of the MPDU dummy on a set of property characteristics that included the interior square footage of the home, the year it was built, the number of bathrooms, the number of stories and an indicator for whether the property is a condo or townhouse. These covariates were entered in a flexible fashion that included cubic splines in square footage and year built as well as the linear interaction of both of those variables with a fully interacted set of dummies for the number of bathrooms and the number of stories. All of these terms were then further interacted with the condo dummy. The region of common support is defined as the set of all propensity scores that are larger than the maximum of the first percentile in the distribution of propensity scores in both sets of properties and smaller than the minimum of the 99th percentile of propensity scores in both distributions.

TABLE A.3
The Effect of Expiring Price Controls on Turnover at MPDU Properties

	OLS				Probit	Logit
	(1)	(2)	(3)	(4)	(5)	(6)
MPDU	-0.072***	-0.064***	-0.063***		-0.057***	-0.055***
	(0.007)	(0.007)	(0.008)		(0.007)	(0.008)
MPDU × Post	0.050***	0.049***	0.049***	0.032***	0.034***	0.030***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)
Property Characteristics	X	X	X		X	X
Year and Age FEs	X	X	X	X	X	X
Subdivision FEs		X	X		X	X
Subdivision Trend			X	X	X	X
Property FEs				X		
Pre-Expiration MPDU Mean	0.048	0.048	0.048	0.048	0.048	0.048
Number of Observations	483,805	483,805	483,805	483,805	483,805	483,805

NOTE.—This table reports difference-in-differences estimates of the effect of expiring MPDU price controls on the annual turnover rate at MPDU properties. Each column reports a separate regression estimated at the property-year level where the dependent variable is an indicator for whether the property sold in a particular year. Coefficients are reported for the "treatment" dummy, denoting whether the property is an MPDU, and the interaction of that dummy with an indicator for whether the year of observation falls on or after the year the first price control within the relevant subdivision expired. All specifications include fixed effects for both the year of observation and the age of the property in that year. The property characteristics include a quadratic in the interior square footage of the home, dummies for the number of bathrooms and the number of stories, and an indicator for whether the property is a condo or townhome as well as the interaction of that indicator with the year fixed effects and all of the other property characteristics including property age. Subdivision trends are estimated by interacting the subdivision fixed effects with a linear time trend. Columns 1–4 report coefficient estimates from linear probability models, while columns 5–6 report marginal effects from probit and logit specifications. The mean of the dependent variable among MPDU properties in the period prior to the first price control expiration is reported in the second to last row. Standard errors are reported in parentheses and are clustered at the subdivision level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE A.4

COVARIATE BALANCE WITHIN TERCILES OF THE PROPENSITY SCORE DISTRIBUTION

	First P-Score Tercile			Second P-Score Tercile			Third P-Score Tercile		
	Non-MPDU	MPDU	t-stat	Non-MPDU	MPDU	t-stat	Non-MPDU	MPDU	t-stat
Fraction Condo	0.824	0.871	-1.164	0.774	0.843	-3.884***	0.916	0.924	-0.797
Square Footage (1000's)	1.333	1.249	1.977*	1.179	1.134	3.687***	1.159	1.141	2.715***
Number of Bathrooms	2.391	1.971	2.921***	2.124	2.187	-1.455	2.084	2.056	1.004
Number of Stories	1.690	1.557	1.691*	1.635	1.623	0.470	1.852	1.893	-2.369**
Age (Years)	24.207	23.186	1.315	24.033	23.832	0.802	23.562	23.400	1.065
Number of Observations			3,392			3,395			3,492

NOTE.—This table presents means of the covariates used to estimate the propensity score. Properties are grouped based on whether their propensity score falls in the bottom, middle, or top third of the combined distribution of propensity scores. Means are then calculated separately for MPDUs and non-MPDUs within these three terciles. For each set of means, the table also reports the t-statistic from a test of the null hypothesis of no difference in means between MPDUs and non-MPDUs. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE A.5
PROPENSITY SCORE MATCHING DIFFERENCE-IN-DIFFERENCES ESTIMATES

	Log Transaction Price	Probability of Extracting Equity	Amount Extracted ($1,000s)
	(1)	(2)	(3)
DID Matching Estimate	0.610***	0.035***	2.783***
	(0.060)	(0.013)	(0.686)
Number of Matched MPDUs	1,846	1,846	1,846
Number of Matched Non-MPDUs	8,443	8,443	8,443
Number of Bootstrap Replicates	100	100	100
Bandwidth	0.1	0.1	0.1

NOTE.—This table presents propensity score matching difference-in-differences estimates of the effect of expiring price controls on log transaction prices, the annual probability of equity extraction, and the total amount of equity extracted per year among MPDU properties and their owners. Estimates were constructed as described in **??**. Bootstrap standard errors are reported in parentheses. Significance levels 10%, 5%, and 1% are denoted by *, **, and *** and were determined based on the assumption that the bootstrap distribution is normally distributed.

A.4. Supplementary Material

A.4.1. The GSEs and the Conforming Loan Limit

The two large government sponsored enterprises—the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac)—were created to encourage mortgage lending. The GSEs purchase mortgages from lenders and either hold them in portfolio or package them into mortgage-backed securities (MBS), which are guaranteed by the GSEs and sold to investors in the secondary market. By purchasing mortgages, the GSEs free up lender capital, allowing the lenders to make additional loans and thus expanding the general availability of mortgage credit.

The GSEs play a large role and exert a substantial amount of influence in the mortgage market. [11] However, they are only allowed to purchase loans which satisfy a specific set of criteria as outlined by their regulator. These criteria include requirements for loan documentation, debt-to-income ratios, leverage, and a nominal cap on the dollar amount of any purchased loan. Loans which meet these criteria and are therefore eligible to be purchased by the GSEs are referred to as "conforming loans." In this paper we are primarily interested in the cap on loan size, known as the "conforming limit." Mortgages exceeding this limit are not eligible for GSE purchase and are referred to as "jumbo loans."

Figure A.11 plots the conforming limit in nominal terms (the solid black line) and in real 2007 dollars (the dashed red line) for each year during our sample period. During this period, the GSEs were regulated by the Office of Federal Housing Enterprise Oversight (OFHEO), which set the limit each year based on changes in the national median house price. The limit was the same for all mortgages in a given year irrespective of local housing market conditions. [12] Following the

---

[11]As of 2010, the GSEs were responsible for nearly 50 percent of the approximately $10.5 trillion in outstanding mortgage debt, either directly or through outstanding MBS (Jaffee and Quigley, 2012). More than 75 percent of all mortgages originated in 2011 passed through the hands of one of the GSEs (Kaufman, 2012).

[12]The only exceptions to this rule were Alaska, Hawaii, the Virgin Islands, and Guam, which were deemed to be high cost areas and had a 50 percent higher conforming limit prior to 2008. Since the housing crisis, the national conforming loan limit has been replaced by a more complicated series of limits set at the metropolitan level. All of the analysis in this paper pertains to the pre-2008 period.

trend in national house prices, the nominal limit increased from around $215,000 in 1997 to its peak in 2006 and 2007 at approximately $420,000. In real terms, the limit also rose sharply over this period, especially during the house price boom of the mid-2000s.

Interest rates on loans above the conforming limit are typically higher than those on comparable loans below the limit for two reasons. First, because the debt underlying the MBS issued by the GSEs is backed by an implicit government guarantee, investors are willing to accept lower yields in exchange for that guarantee. [13] Part of this savings is eventually passed on to borrowers in the form of lower interest rates on conforming loans. [14] Second, the GSEs are also granted several special privileges that private securitizers are not. These include access to a line of credit at the U.S. Treasury, exemption from disclosure and registration requirements with the Securities and Exchange Commission (SEC), and exemptions from state and local income taxes. [15] These advantages lower the cost of securitizing mortgages for the GSEs relative to private market securitizers, with some of the savings passed on to borrowers in the form of lower interest rates on loans below the conforming limit.

A.4.2. Endogenous Housing Choice

With the choice of housing fixed, as in the discussion in the main text, borrowers can only respond to the presence of a notch by adjusting their mortgage balance. In other words, all households buy the same house at the same price as in the absence of a notch, but some households respond to the notch by making a larger down payment or taking out a second mortgage. In reality, some households may instead choose to buy a lower quality home, leading to a lower level of h.

Our model extends to cover endogenous housing choice, albeit at the cost of a closed-form solution. Consider again equation (2.5), the household's intertemporal optimization problem. Households can now choose the quantity of housing services to purchase (h), and this quantity has a

---

[13]The implicit guarantee became explicit in 2008 when the GSEs were placed under government conservatorship.
[14]Passmore et al. (2002) and Passmore et al. (2005) provide several theoretical explanations for how the savings from the guarantee are eventually passed down to mortgage borrowers.
[15]For a full description of the direct benefits conferred on the GSEs as a result of their special legal status see Congressional Budget Office (2001).

direct effect on first-period utility, so that

$$V = \max_{m,h}\{u(y+m-ph,h)+\delta v(ph-(1+r)m)\}, \tag{A.5}$$

with $v(c_2)$ now denoting second-period utility, as distinct from $u(c_1,h)$ in which housing enters directly.

The optimal h and m must now satisfy two first-order conditions:

$$\frac{\partial V}{\partial m} = u_1 - \delta(1+r)v_1 = 0 \tag{A.6}$$

$$\frac{\partial V}{\partial h} = u_2 - (pu_1 - p\delta v_1) = 0. \tag{A.7}$$

Intuitively, the first condition captures the trade-off, using mortgage debt, between consumption today and consumption tomorrow. The second condition says that households trade off the cost of purchasing housing today, less the amount recovered tomorrow when it is sold, against its consumption value today.

While there are no obvious functional forms that allow us to derive equivalents to equation (2.6), the intuition remains the same. Under standard conditions, there are optimal m^* and h^*, both of which can shift in response to the notch in the interest rate schedule. Our bunching estimation will capture the shifts in m^*, which could result in part from changes in housing consumption (h^*).

A.4.3. Summary Statistics

Table A.6 presents summary statistics for our primary estimation sample from DataQuick as well as the sub-sample of transactions with first loan amounts within $50,000 of the conforming limit that was in place in the year of the transaction. All dollar amounts here and throughout the analysis are converted to real 2007 dollars.

In the full sample, shown in column 1, the mean first loan size is approximately $350,000 and the mean transaction price is $465,000. Column 3 shows the means from the restricted sample. Although the large sample size means that many of the differences between columns 1 and 3 are statistically significant, they are qualitatively similar along all dimensions. Interestingly, because the restricted sample drops both high priced houses and low priced houses, the average transaction price and loan amount near the conforming limit are actually a bit lower than the averages for the entire sample. In many states with lower average house prices, there are relatively few loans made substantially above the limit, but in California such transactions are much more common.

Table A.7 presents summary statistics from the LPS data for fixed-rate (FRM) and adjustable-rate (ARM) loans separately. Columns 1 and 3 report statistics for the full analysis sample while columns 2 and 4 restrict the sample to loans within $50,000 of the conforming loan limit. In general, the restricted samples for each loan type are quite similar to the full sample, suggesting that loans near the limit are reasonably representative of the entire sample, at least along these dimensions.

A.4.4. Heterogeneity by Borrower Type

In addition to investigating bunching behavior by loan type, as in subsubsection 2.5.1, it is also interesting to examine whether bunching varies with the observable characteristics of borrowers. While the available information on borrower demographics is somewhat limited, we are able to provide several rough cuts of the data based on race and income by matching a subset of DQ transactions to loan application information made available through the Home Mortgage Disclosure Act (HMDA). [16] For this exercise we restrict attention to fixed-rate mortgages, where the sample sizes are largest. Using the race and ethnicity information in the HMDA data, we define a loan as belonging to a "minority" borrower if the primary loan applicant reports his race as black or his ethnicity as Hispanic and as belonging to a "non-minority" borrower otherwise. Similarly, we define a borrower as "low-income" if the income reported on the loan application was below the

---

[16]The matching procedure uses information on the primary loan amount, lender name, Census tract, property type, and year. We successfully match about 60 percent of the larger DQ sample to observations in HMDA. Further details are available from the authors on request.

median income reported across all loans and "high-income" if the reported income is above the median.

Figure A.12 shows results from estimating bunching separately in each of these four sub-samples. Panels (A) and (B) show results for non-minority and minority borrowers while panels (C) and (D) present the results for high-income and low-income borrowers, respectively. In each case there is substantial evidence of heterogeneous responses, with far less bunching among minority and low-income borrowers than among non-minority and high-income borrowers. For non-minority and high-income borrowers the estimated percentage reduction in loan size is roughly 7 to 8 percent while for minority and low-income borrowers it is closer to 4 to 5 percent. [17]

These differences could arise from at least three sources: heterogeneous preferences, heterogeneous costs of adjusting first mortgage balances, or borrower-level differences in the magnitude of the jumbo-conforming spread. While we cannot test for differences in the magnitude of the spread along these dimensions because the LPS data do not contain information on race or income, we do not find robust evidence of differences along other dimensions that are likely correlated with these characteristics, such as borrowers' credit scores. This finding suggests that one of the first two sources of heterogeneity are likely operative.

In the context of income taxes, Kleven and Waseem (2013) are able to distinguish between the role of preferences and adjustment costs using the fact that a notch can sometimes create a dominated region in which no wage earner, regardless of tax elasticity, would choose to locate in the absence of adjustment costs. By counting the number of wage earners observed in the dominated region they are able to back out an estimate of adjustment costs. Unfortunately, we cannot perform a similar exercise here because there is no such dominated region in our setting. In the terminology of

---

[17]Table A.8 confirms the visual impressions given by Figure A.12, reporting the point estimates and standard errors for the bunching parameters we estimate. For each of the reported parameters, the standard errors are small enough to reject the null that bunching behavior is the same across the high-income and low-income samples, as well as across the non-minority and minority samples. Part of the smaller response among low-income borrowers could be driven by the fact that there are virtually no borrowers in the low-income sample who take out jumbo loans. That is, almost all of the low-income borrowers who would locate anywhere to the right of the limit have chosen to bunch, which limits the possible magnitude of the estimated response. Consequently, the bunching borrowers may all be infra-marginal, and the bunching estimate cannot be interpreted as an average marginal response, as in the theory.

Kleven and Waseem (2013), the jumbo-conforming spread creates a "downward notch" where, for any finite loan amount above the limit, there exists a first mortgage demand elasticity sufficiently close to zero such that some borrower would be willing to take out that loan. It is therefore not possible to estimate the magnitude of any differences in adjustment costs across these four groups.

While it seems more likely that these differences are driven by adjustment costs than by differences in underlying preferences for first mortgage debt, we leave a full analysis of this issue for future research. However, it is important to note that because we are not able to determine the magnitude of such adjustment costs, the elasticities we estimate are necessarily "reduced form," in the sense that they incorporate the effect of adjustment costs and are not driven entirely by the intertemporal elasticity of substitution alone.

A.4.5. *More Cash or Cheaper Houses?*

In Section 2.6.3 we provide results suggesting that roughly 65 percent of borrowers who bunch at the conforming limit do so without the use of a second mortgage. In order to lower their first mortgage amount, these borrowers must either be putting up more cash or spending less on housing than they otherwise would. If they are spending less on housing while holding their leverage roughly constant, then both their first and combined LTVs should be little changed relative to the counterfactual in which the conforming limit does not exist. [18] Therefore, if all of the bunching borrowers were either taking out second mortgages or buying cheaper houses, then we would expect the average combined LTV at the conforming limit to be about the same as in nearby bins. Figure A.13, which plots the combined LTV against the first mortgage amount, makes clear that this (admittedly extreme) scenario is far from true. The combined LTV at the limit is about 75 percent, well below the 80 to 85 percent that would be predicted based on the red line, which is a polynomial fit using the data outside of the same excluded region that was used to estimate bunching in Figure 2.7A. [19] Thus, a significant portion of the 65 percent of borrowers who bunch

---

[18]Some borrowers may buy a cheaper house but target an ideal monthly payment, rather than an ideal LTV. Such borrowers would slightly reduce their LTV when they bunch, but by much less than buying the same house at the same price.

[19]The 85 to 90 percent combined LTVs to the right of the limit, which are higher than any of the other points on the plot, also stand out. One possibility is that the borrowers in these bins who do not bunch have different characteristics

without a second mortgage must be doing so by putting up more cash as opposed to spending less on housing.

In the most extreme case, all of these borrowers are putting up more cash. Since both putting up more cash and taking out a second mortgage reduce a borrower's first-mortgage LTV while spending less on housing does not, we can gauge the plausibility of this extreme case by examining the relationship between first-mortgage LTV ratios and loan size near the limit. To do so, Figure A.14 plots the first-mortgage LTV against the first mortgage amount. This figure is analogous to Figure A.13 except for the blue "X", which is a first-mortgage LTV calculated under the assumption that none of the bunching borrowers adjust their housing expenditure.

To calculate this LTV, we first generate a "counterfactual" mean house price for each loan size bin by fitting a 5th degree polynomial to the observed mean price in each bin omitting the bins in the excluded region used to estimate bunching. We then take the weighted average of these mean counterfactual house prices in the bin containing the limit and in each bin to the right of the limit in the excluded region. In calculating this average, the weight assigned to the mean price at the limit is the estimated counterfactual bin count from the bunching procedure, and the weights assigned to each of the mean prices to the right of the limit are equal to the difference between the counterfactual and observed bin counts from Figure 2.7A. To calculate the LTV plotted in the figure, we then compute the average conforming limit for all loans observed at the limit and divide by the weighted average counterfactual house price.

Somewhat surprisingly, the observed LTV at the limit is even lower than under the extreme scenario used to calculate the "X", in which no borrower adjusts her house price. Even allowing for some noise in the estimates, this comparison suggests that very few borrowers who bunch at the limit do so by buying cheaper houses. This LTV result is consistent with the results in Adelino et al. (2012), who also document that borrowers who purchase homes using mortgages at the conforming limit have substantially lower first loan LTVs than those just above the limit.

than those who do and are thus a selected sample. Indeed, these borrowers are clearly somewhat "abnormal," in that they do not bunch despite the seemingly large gains from doing so. However, as we know from Figure 2.7A, there are relatively few borrowers remaining in these bins.

Although these calculations suggest that there is no direct impact of the interest rate differential on house prices and the demand for housing itself, we do not necessarily want to draw that inference, for three reasons. First, the LTV calculations above are "back of the envelope" and there are several untested assumptions involved. Second, the interest rate differential at the limit is relatively small and may not be as informative about larger changes in rates over time. Finally, it is likely that other methods, in particular those used by Adelino et al. (2012), are better suited to studying the effects of rates on house prices. They estimate the elasticity of house prices to interest rates by following similar houses over time as they become more or less difficult to finance with a conforming loan at a constant LTV of 80 percent, essentially using a difference-in-difference approach rather than the bunching approach we follow. This approach has advantages over ours for this particular question, especially because we cannot observe borrowers who drop out of the market entirely (the extensive margin).

FIG. A.11.—Conforming Loan Limit Over Time. This figure plots the annual conforming loan limit for single family homes in nominal and real 2007 dollars for each year between 1997 and 2007. Historical conforming limits were taken from the Office of Federal Housing Enterprise Oversight's (OFHEO) 2007 annual report to congress. Nominal dollars are inflated using the Consumer Price Index for all Urban Consumers (CPI-U).

Panel A. Non-Minority

Panel B. Minority
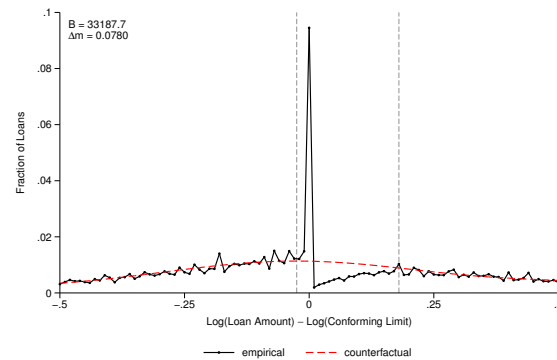
Panel C. High-Income

Panel D. Low-Income

FIG. A.12.—Bunching at the Conforming Loan Limit by Borrower Type, Fixed Rate Mortgages Only. This figure plots the empirical and counterfactual density of (log) loan size relative to the conforming limit estimated separately for: (a) non-minority borrowers, (b) minority borrowers, (c) high-income borrowers and (d) low-income borrowers. The connected black line is the empirical density. Each dot represents the fraction of loans in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed red line is the estimated counterfactual density obtained by fitting a 13th degree polynomial to the bin counts, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The figure also reports the estimated number of loans bunching at the limit (B) and the average behavioral response among marginal bunching individuals (Δm), calculated as described in Section 2.4.1.
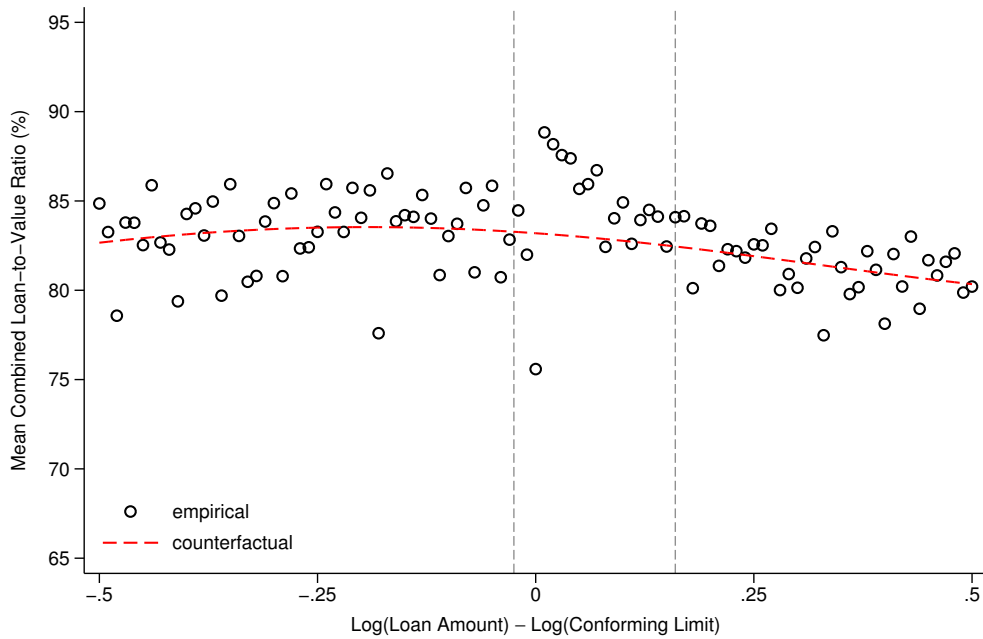
FIG. A.13.—Combined Loan-to-Value Ratio by First Mortgage Amount. This figure plots the average combined loan-to-value ratio (CLTV) as a function of the first loan amount relative to the conforming limit. Each dot represents the average CLTV in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed red line is the counterfactual mean CLTV obtained by fitting a 5th degree polynomial to the bin averages, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The excluded region is the same region used to estimate bunching for the sample of fixed-rate mortgages. CLTV is calculated as the ratio of the sum of up to three mortgages used to finance a transaction to the recorded purchase price. Sample includes only transactions with a fixed-rate first mortgage.
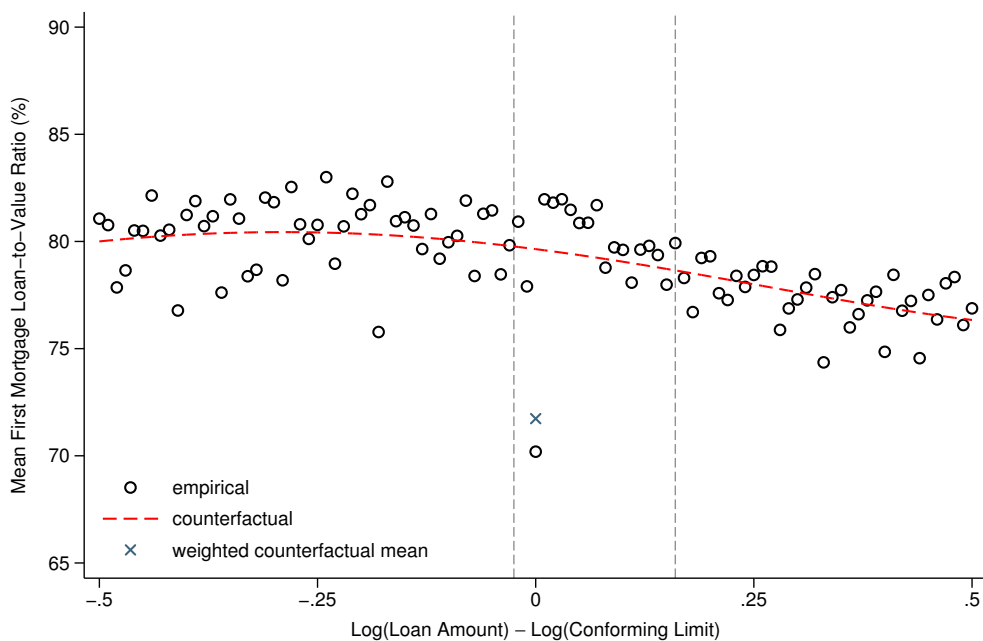
FIG. A.14.—First Mortgage Loan-to-Value Ratio by First Mortgage Amount. This figure plots the average first mortgage loan-to-value ratio (LTV) as a function of the first loan amount relative to the conforming limit. Each dot represents the average LTV in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed red line is the counterfactual mean LTV obtained by fitting a 5th degree polynomial to the bin averages, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The excluded region is the same region used to estimate bunching for the sample of fixed-rate mortgages. The blue "X" is an LTV calculated assuming that borrowers who bunch at the limit do so without adjusting their housing price. See the text for the details of this calculation. LTV is calculated as the ratio of the first loan amount to the recorded purchase price. Sample includes only transactions with a fixed-rate first mortgage.

SUMMARY STATISTICS, DATAQUICK SAMPLE

	Full Sample		Within $50k of CLL	
	(1)	(2)	(3)	(4)
	Mean	SD	Mean	SD
Transaction and Loan Characteristics				
First Loan Amount ($1,000)	349	(229)	342	(55)
Transaction Price ($1,000)	465	(340)	449	(123)
Has Second Loan	0.37	(0.48)	0.41	(0.49)
First Loan ARM	0.49	(0.50)	0.48	(0.50)
Housing Characteristics				
Square Footage	1,764	(4,614)	1,787	(2,979)
Property Age (Years)	29	(25)	29	(24)
Number of Bedrooms	3.2	(1.3)	3.3	(1.4)
Number of Bathrooms	2.1	(1.0)	2.2	(0.9)
Borrower Characteristics				
Applicant Income ($1,000)	142	(181)	133	(127)
White	0.50	(0.50)	0.50	(0.50)
Black	0.03	(0.18)	0.03	(0.17)
Hispanic	0.21	(0.41)	0.19	(0.39)
Observations	2,739,775		637,369	

NOTE.—Means and standard deviations for select variables from DataQuick data set. Columns (1) and (2) are based on the full sample of all DataQuick transactions recorded in California between 1997 and 2007. Columns (3) and (4) restrict the sample to only transactions with first mortgage amounts within $50,000 of the conforming limit in effect at the time of origination. All dollar amounts are in real 2007 dollars. Statistics for transaction and housing characteristics are calculated using all available transactions. Statistics for borrower characteristics are calculated using only the subset of transactions that match to a HMDA loan application. See text for details on sample construction.

TABLE A.7
SUMMARY STATISTICS, LPS SAMPLE

	FRMs		ARMs	
	(1) Full Sample	(2) Within $50k of CLL	(3) Full Sample	(4) Within $50k of CLL
Interest Rate (Initial for ARMs)	6.69	6.70	5.05	5.04
	(0.91)	(0.90)	(2.13)	(2.09)
Jumbo	0.17	0.14	0.50	0.39
	(0.38)	(0.35)	(0.50)	(0.49)
First Loan Amount ($1,000)	278.03	322.55	452.13	376.95
	(174.40)	(56.81)	(283.42)	(47.29)
Appraisal Amount ($1,000)	396.50	447.16	611.28	497.01
	(281.31)	(156.74)	(438.51)	(105.31)
Loan-to-Value Ratio	74.09	76.09	76.31	77.33
	(16.35)	(13.63)	(10.07)	(9.17)
Debt-to-Income Ratio	35.62	36.79	35.51	36.50
	(12.47)	(11.96)	(11.96)	(11.46)
Missing DTI Ratio	0.70	0.75	0.45	0.45
	(0.46)	(0.44)	(0.50)	(0.50)
FICO Score	731.18	731.70	719.70	717.97
	(51.81)	(49.98)	(52.82)	(52.52)
Missing FICO Score	0.32	0.35	0.20	0.22
	(0.47)	(0.48)	(0.40)	(0.41)
Term (Months)	345.91	350.14	365.78	365.63
	(52.58)	(46.44)	(30.73)	(28.83)
30-Year	0.90	0.93	0.93	0.94
	(0.30)	(0.26)	(0.25)	(0.23)
Observations	1,062,164	264,654	947,565	224,475

NOTE.—Means and standard deviations (in parentheses) for select variables from the LPS data set. Columns (1) and (3) are based on the full sample of fixed-rate and adjustable-rate purchase mortgages originated in California between 1997 and 2007. Columns (2) and (4) restrict these samples to only loans that fall within $50,000 of the conforming limit in effect at the time of origination. All dollar amounts are in real 2007 dollars. See text for details on sample construction.

TABLE A.8
BUNCHING ESTIMATES BY BORROWER TYPE, FRMs ONLY

	(1) High-Income	(2) Low-Income	(3) Non-Minority	(4) Minority
Bunched Loans ($\hat{B}$)	33187.7 (548.4)	8899.3 (464.6)	34013.6 (736.5)	3715.2 (150.3)
Behavioral Response ($\Delta\hat{m}$)	0.0780 (0.0028)	0.0487 (0.0036)	0.0741 (0.0033)	0.0378 (0.0019)
Excess Mass ($\hat{B}/\sum_{j=L}^{0}\hat{n}_j$)	7.801 (0.276)	4.868 (0.362)	7.406 (0.328)	3.777 (0.196)
Upper Limit (m_H)	0.180 (0.021)	0.120 (0.019)	0.180 (0.023)	0.110 (0.015)

NOTE.—Each column reports the estimated number of loans bunching at the conforming limit ($\hat{B}$), the average (log) shift in mortgage balance in response to the conforming limit among marginal bunching individuals ($\Delta\hat{m}$), the excess mass at the conforming limit ($\hat{B}/\sum_{j=L}^{0}\hat{n}_j$), and the upper limit of the excluded region used in estimation (m_L). Estimates are reported separately for high- and low-income borrowers and for minority and non-minority borrowers. High-income borrowers are those who report an income on their loan application that is higher than the median in the pooled sample. Low income borrowers are those below the median. Minority borrowers are those who identify as either black or Hispanic on their loan applications. Sample includes only transactions with a fixed-rate first mortgage which could be successfully matched to a mortgage application in the HMDA data and for which the borrower reported their income as well as both a race and an ethnicity. Standard errors (in parentheses) were calculated using the bootstrap procedure described in section 2.4.1.

TABLE A.9
JUMBO-CONFORMING SPREAD ESTIMATES, LOG POINTS

	(1) Baseline	(2) Splines	(3) Within $50k of CLL	(4) Within $10k of CLL
Fixed-Rate Mortgages				
OLS	0.026	0.027	0.023	0.025
	(0.000)	(0.000)	(0.002)	(0.006)
IV	0.017	0.018	0.017	0.017
	(0.002)	(0.001)	(0.005)	(0.010)
Observations	1,061,738	1,061,738	263,641	87,617
Adjustable-Rate Mortgages				
OLS	-0.051	-0.053	-0.168	-0.210
	(0.003)	(0.003)	(0.015)	(0.034)
IV	0.005	0.009	0.014	0.002
	(0.008)	(0.008)	(0.022)	(0.045)
Observations	686,970	686,970	156,403	39,198

NOTE.—Standard errors in parentheses. Estimates of jumbo-conforming spread, in log points, using OLS and IV with an indicator for home value being greater than $1.25\bar{m}$ used as an instrument for jumbo-status, as described in the text. Controls include distance to CLL (cubic), LTV ratio, DTI ratio, missing LTV and DTI ratios, FICO score, missing FICO score, PMI, prepayment penalty, and mortgage term, as well as month by zip-code fixed effects. Column 1 includes linear effects of LTV and DTI ratios. Column 2 includes cubic B-splines in LTV and DTI ratios, as well as FICO score. Columns 3 and 4 limit the sample to loans near the CLL.

A.5. Estimating Multiple Breakpoints

In estimating the break points, we allow for the possibility that a given market might experience more than one housing boom during the course of our sample period. Our method is recursive in that we first test for the existence of one break point against the null hypothesis of zero. Given the existence of at least one break point, we can then test the hypothesis of $m + 1$ break points against the null of m using the results from Bai (1999). Bai and Perron (1998) show that the test for one break is consistent in the presence of multiple breaks, which is what allows for this sequential estimation procedure.

More specifically, let $0 < \psi_{i,1} < \cdots < \psi_{i,m} < 1$ mark the proportions of the sample generated by the m break points estimated under the null hypothesis for MSA i. For technical reasons, we require that $\psi_{i,j} - \psi_{i,j-1} > \pi_{i,0}$ for some small $\pi_{i,0}$ where we define $\psi_{i,0} = 0, \psi_{i,m+1} = 1$. Further, let $\eta_{i,j} = \frac{\pi_{i,0}}{\psi_{i,j} - \psi_{i,j-1}}, j = 1,, m + 1$. The likelihood ratio test compares the maximum of the likelihood ratio obtained when allowing for $m + 1$ breaks to that from only allowing for m. The distribution of this likelihood ratio statistic is given by

$$P(LR > c) = 1 - \prod_{i=1}^{m+1} \left(1 - P\left(sup_{\pi_i \in [\eta_{i,j}, 1 - \eta_{i,j}]} Q_1(\pi_i)\right)\right), \tag{A.8}$$

which we calculate by recursive application of the method provided in Estrella (2003).

We apply this procedure to test for the existence of two break points against the null of one as well as three against the null of only two among those MSAs for which we find at least two statistically significant break points. There are some noteworthy practical issues involved with carrying out this procedure. We have not until this point said where the sample proportions $\pi_{i,0}, \pi_{i,1}, \pi_{i,2}$ come from. In practice, we restrict the full sample period for each MSA to lie between the first quarter in the data and the peak of price growth. We then do not allow any break points to lie in either the first or last two quarters of this sample for each MSA. This determines the fractions $\pi_{i,1}$ and $\pi_{i,2}$ which, because different MSAs have a different number of quarters, will vary across areas.

When estimating multiple break points, we further require that any two break points be at least four quarters apart. This determines the fraction $\pi_{i,0}$ which, again, will vary across areas due to differing sample sizes. Because of these restrictions, we are not able calculate p-values for many MSAs in the case of multiple breaks. The reason for this can be seen from the expression in (A.8). Because this expression requires that $\eta_j < 0.5$, we must require that $\frac{\pi_{i,0}}{\psi_{i,j}-\psi_{i,j-1}} > 0.5$ for all j. This implies that we will not be able to calculate p-values for the two-break case in MSAs (neighborhoods) where the first break is less than $\pi_{i,0}/0.5$ from the beginning of the sample period. Naturally, this restriction is more burdensome when trying to calculate p-values in the three break case.

A.6. A Hazard Model of Housing Boom Contagion

Consider an economy with N metro areas. The probability of MSA m entering a housing boom increases if it 'meets' other MSAs which already have entered a housing boom. Assume α is the frequency of such a meeting and that the connection between MSA pairs decays in distance at rate δ. Following Comin et al. (2012), we can write MSA m's probability of not entering a housing boom at time $t + h$ conditional on not having a housing boom in time t as

$$P(0,m,t+h) = P(0,m,t) \left[\frac{\sum_{q \neq m} P(0,q,t)e^{-\delta r_{mq}}}{\sum_{q \neq m} e^{-\delta r_{mq}}} \right]^{\alpha h} \tag{A.9}$$

where r_{mq} denotes the distance between MSA m and MSA q. Taking $h \to 0$, we have

$$\frac{\partial ln P(0,m,t)}{\partial t} = \alpha ln \left(\sum_{q \neq m} P(0,q,t)e^{-\delta r_{mq}} \right) - \alpha ln \left(\sum_{q \neq m} e^{-\delta r_{mq}} \right) \tag{A.10}$$

By assumption, $P(0,q,t) = 0, \forall t \tau$ if MSA q enters a housing boom at time τ. As long as some MSA enters a boom, equation (A.9) implies that $\frac{\partial ln P(0,m,t)}{\partial t} < 0$, so that the hazard of entering a housing boom increases over time. To consider the contagious effect of housing booms, suppose

MSA q booms at time t. That increases the hazard of MSA m having a boom because

$$\frac{\partial lnP(0,m,t)}{\partial t\,\partial P(0,q,t)} = \frac{\alpha e^{-\delta r_{mq}}}{\sum_{q\neq m} P(0,q,t)e^{-\delta r_{mq}}} > 0 \tag{A.11}$$

In addition, the contagious effect of a housing boom in MSA q decreases over geographical distance, because

$$\left.\frac{\partial^2 lnP(0,m,t)}{\partial t\,\partial r_{mq}}\right|_{P(0,q,t)=0} = \frac{\alpha\delta e^{-\delta r_{mq}}}{\sum_{q\neq m} e^{-\delta r_{mq}}} > 0 \tag{A.12}$$

In sum, this framework provides two relevant implications for our purposes:

Implication 1: The fact that MSA q enters a housing boom at time t increases the hazard of MSA m (mq) having a housing boom at time $t+1$.

Implication 2: The contagion effect in **Implication 1** becomes larger when MSA q is closer to MSA m.
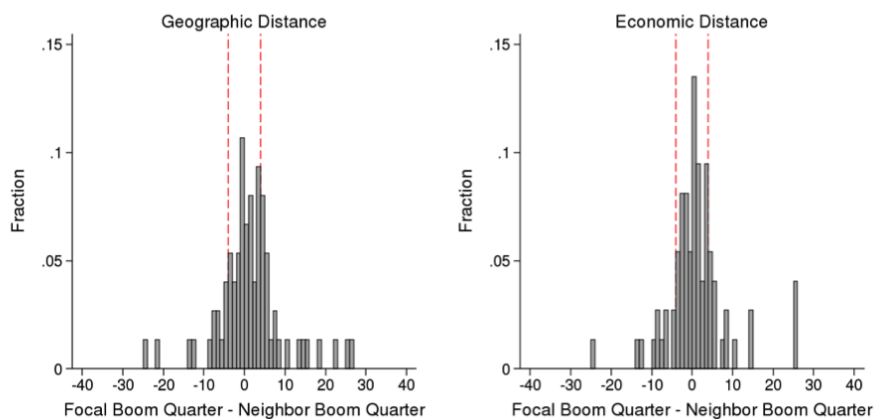
A.7. Additional Figures

Fig. A.15.—Histogram of number of quarters between timing of the booms of focal markets and nearest neighbors. Note: The level of observation is the MSA-quarter.
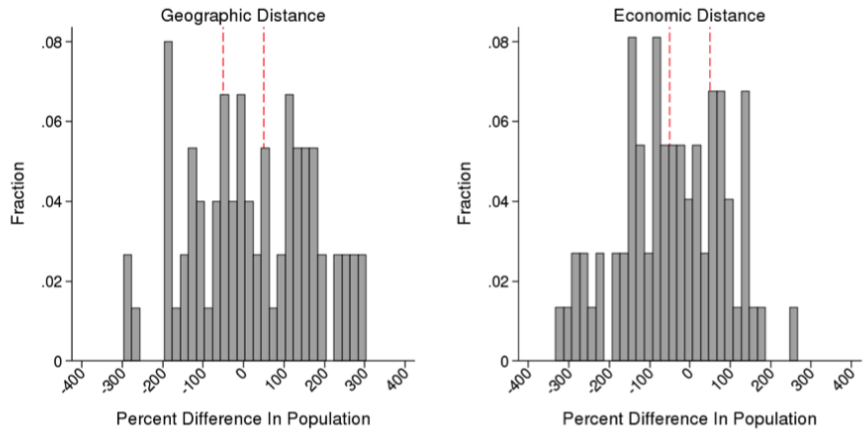
FIG. A.16.—Histogram of percentage difference in population of focal markets and nearest neighbors. Note: The level of observation is the MSA-quarter.

BIBLIOGRAPHY

Aaronson, Daniel, Sumit Agarwal, and Eric French, "The Spending and Debt Response to Minimum Wage Hikes," *American Economic Review*, 2012, *102* (7), 3111–3139.

Abdallah, Chadi S. and William D. Lastrapes, "Home Equity Lending and Retail Spending: Evidence from a Natural Experiment in Texas," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 94–125.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Credit Supply and House Prices: Evidence From Mortgage Market Segmentation," February 2012. NBER Working Paper No. 17832.

_ , _ , and _ , "House Prices, Collateral and Self-Employment," 2014. Forthcoming, *Journal of Financial Economics*.

_ , **Kristopher Gerardi, and Paul Willen**, "Why Don't Lenders Renegotiate More Home Mortgages? Redefaults, Self-Cures and Securitization," 2009. FRB Atlanta Working Paper 2009-17.

Agarwal, Sumit and Wenlan Qian, "Access to Home Equity and Consumption: Evidence from a Policy Experiment," 2014. Working Paper.

_ , **Chunlin Liu, and Nicholas S. Souleles**, "The Reaction of Consumer Spending and Debt to Tax Rebates–Evidence from Consumer Credit Data," *Journal of Political Economy*, 2007, *115* (6), 986–1019.

_ , **John C. Driscoll, and David I. Laibson**, "Optimal Mortgage Refinancing: A Closed-Form Solution," *Journal of Money, Credit and Banking*, 2013, *45* (4), 591–622.

Aladangady, Aditya, "Household Balance Sheets and Monetary Policy," 2013. Working Paper.

Alessie, Rob, Stefan Hochguertel, and Guglielmo Weber, "Consumer Credit: Evidence from Italian Micro Data," *Journal of the European Economic Association*, 2005, *3* (1), 144–178.

Almeida, Heitor, Murillo Campello, and Crocker Liu, "The Financial Accelerator: Evidence from International Housing Markets," *Review of Finance*, 2006, *10* (3), 321–352.

Andrews, Donald, "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 1993, *61* (1), 821–856.

Arce, Oscar and David Lopez-Salido, "Housing Bubbles," *American Economic Journal: Macroeconomics*, 2011, *3*, 212–241.

Attanasio, Orazio P., Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou, "Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans," *International Economic Review*, 2008, *49* (2), 401–436.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Avery, Robert and Kenneth Brevoort, "The Subprime Crisis: How Much Did Lender Regulation Matter," 2010. Division of Research and Statistics. Board of Governors of the Federal Reserve System.

Bai, Jushan, "Likelihood Ratio Tests for Multiple Structural Changes," *Journal of Econometrics*, 1999, *91* (2), 299–323.

_ **and Pierre Perron**, "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 1998, *66* (1), 47–78.

Bajari, Patrick, Chenghuan Sean Chu, and Minjung Park, "An Empirical Model of Subprime Mortgage Default from 2000-2007," 2008. NBER Working Paper No. 14625.

Baker, Scott R., "Debt and the Consumption Response to Household Income Shocks," 2014. Working Paper.

Bayer, Patrick, Christopher Geissler, and James Roberts, "Speculators and Middlemen: The Role of Flippers in the Housing Market," 2011. NBER Working Paper No. 16784.

Benjamin, John D., Peter Chinloy, and G. Donald Jud, "Real Estate Versus Financial Wealth in Consumption," *Journal of Real Estate Finance and Economics*, 2004, *29* (3), 341–354.

Benmelech, Efraim, Mark J. Garmaise, and Tobias J. Moskowitz, "Do Liquidation Values Affect Financial Contracts? Evidence from Commercial Loan Contracts and Zoning Regulation," *Quarterly Journal of Economics*, 2005, *120* (3), 1121–1154.

Bernanke, Ben and Mark Gertler, "Agency Costs, Net Worth, and Business Fluctuations," *American Economic Review*, 1989, *79* (1), 14–31.

Best, Michael and Henrik J. Kleven, "Property Transaction Taxes and the Housing Market: Evidence from Notches and Housing Stimulus in the UK," February 2015. Working Paper.

Bhutta, Neil, "Regression Discontinuity Estimates of the Effects of the GSE Act of 1992," 2009. Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C.

_ **and Benjamin J. Keys**, "Interest Rates and Equity Extraction During the Housing Boom," 2014. University of Chicago Kreisman Working Papers Series in Housing Law and Policy No. 3.

Bostic, Raphael, Stuart Gabriel, and Gary Painter, "Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data," *Regional Science and Urban Economics*, 2009, *39* (1), 79–89.

Bound, John, Charles Brown, and Nancy Mathiowetz, "Measurement Error in Survey Data," in James J. Heckmand and Edward Leamer, eds., *Handbook of Econometrics*, Vol. 5, Elsevier Science, North-holland, 2001, chapter 59, pp. 3705–3843.

Browning, Martin and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 1996, *34* (4), 1797–1855.

⎯ **and M. Dolores Collado**, "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates," *American Economic Review*, 2001, *91* (3), 681–692.

⎯ **and Thomas F. Crossley**, "The Life-Cycle Model of Consumption and Saving," *Journal of Economic Perspectives*, 2001, *15* (3), 3–22.

⎯ **, Mette Gørtz, and Søren Leth-Petersen**, "Housing Wealth and Consumption: A Micro Panel Study," *The Economic Journal*, 2013, *123* (568), 401–428.

Brueckner, Jan K., "The Demand for Mortgage Debt: Some Basic Results," *Journal of Housing Economics*, 1994, *3*, 251–262.

Bubb, Ryan and Alex Kaufman, "Securitization and Moral Hazard: Evidence from a Lender Cut-Off Rule," 2009. Working Paper.

Buiter, Willem H., "Housing Wealth Isn't Wealth," 2008. NBER Working Paper No. 14204.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, "Understanding Booms and Busts in Housing Markets," 2011. NBER Working Paper No. 16734.

Calormiris, Charles W., Stanley D. Longhofer, and William Miles, "The Housing Wealth Effect: The Crucial Roles of Demographics, Weath Distribution and Weath Shares," *Critical Finance Review*, 2013, *2* (1), 49–99.

Calvo, Sara and Carmen Reinhart, "Capital Flows to Latin America: Is There Evidence of Contagion Effects?," in Guilermo A. Calvo, Morris Goldstein, and Eduard Hochreiter, eds., *Private Capital Flows to Emerging Markets After the Mexican Crisis*, Institute for International Economics, 1996.

Campbell, John Y. and João F. Cocco, "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics*, 2007, *54* (3), 591–621.

⎯ **, Stefano Giglio, and Parag Pathak**, "Forced Sales and House Prices," *American Economic Review*, 2011, *101* (5), 2108–2131.

Caplin, Andrew, Charles Freeman, and Joseph Tracy, "Collateral Damage: Refinancing Constraints and Regional Recessions," *Journal of Money, Credit and Banking*, 1997, *29* (4), 496–516.

Card, David, Alexandre Mas, and Jesse Rothstein, "Tipping and the Dynamics of Segregation," *Quarterly Journal of Economics*, 2008, *123* (1), 177–218.

Carroll, Christopher D. and Xia Zhou, "Dynamics of Wealth and Consumption: New and Improved Measures for U.S. States," *B.E. Journal of Macroeconomics*, 2012, *12* (2).

_ , **Misuzu Otsuka, and Jiri Slacalek**, "How Large Are Housing and Financial Wealth Effects? A New Approach," *Journal of Money, Credit and Banking*, 2011, *43* (1), 55–79.

Case, Karl and Robert Shiller, "Prices of Single Family Homes Since 1970: New Indexes for Four Cities," *New England Economic Review*, 1987, pp. 46–56.

_ **and** _ , "The Efficiency of the Market for Single Family Homes," *American Economic Review*, 1989, *70* (1), 125–137.

Case, Karl E., John M. Quigley, and Robert J. Shiller, "Comparing Wealth Effects: The Stock Market versus the Housing Market," *Advances in Macroeconomics*, 2005, *5* (1), 1–32.

_ , _ , **and** _ , "Wealth Effects Revisited 1978-2009," *Critical Finance Review*, 2013, *2* (1), 101–128.

Chaney, Thomas, David Sraer, and David Thesmar, "The Collateral Channel: How Real Estate Shocks Affect Corporate Investment," *American Economic Review*, 2012, *102* (6), 2381–2409.

Chen, Hui, Michael Michaux, and Nikolai Roussanov, "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty," 2013. NBER Working Paper No. 19421.

Chetty, Raj, John N. Friedman, and Emmanuel Saez, "Using Differences in Knowledge Across Neighborhoods to Uncover the Impact of the EITC on Earnings," 2013. Forthcoming, American Economic Review.

_ , _ , **Tore Olsen, and Luigi Pistaferri**, "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126* (749-804).

Chinco, Alex and Chris Mayer, "Distant Speculators and Asset Bubbles in the Housing Market," 2012. Working Paper.

Chomsisengphet, Souphala and Anthony Pennington-Cross, "The Evolution of the Subprime Mortgage Market," *Federal Reserve Bank of St. Louis Review*, 2006, *88* (1), 31–56.

Clapp, John M., Walter Dolde, and Dogan Tirtiroglu, "Imperfect Information and Investor Inferences From Housing Price Dynamics," *Real Estate Economics*, 1995, *23* (3), 239–263.

Comin, Diego, Mikhail Dmitriev, and Esteban Rossi-Hansberg, "The Spatial Diffusion of Technology," 2012. Working Paper.

Congressional Budget Office, "Federal Subsidies and the Housing GSEs," 2001.

Cooper, Daniel, "Did Easy Credit Lead to Overspending? Home Equity Borrowing and Household Behavior in the Early 2000s," 2010. Federal Reserve Bank of Boston Public Policy Discussion Papers No. 09-7.

___ , "House Price Fluctuations: The Role of Housing Wealth as Borrowing Collateral," *Review of Economics and Statistics*, 2013, *95* (4), 1183–1197.

Corsetti, Giancarlo, Marcello Pericoli, and Massimo Sbracia, "Some Contagion, Some Interdependence: More Pitfalls in Tests of Financial Contagion," *Journal of International Money and Finance*, 2005, *24*, 1177–1199.

Cotter, John, Stuart Gabriel, and Richard Roll, "Integration and Contagion in US Housing Markets," 2011. Working Paper.

Cvijanovič, Dragana, "Real Estate Prices and Firm Capital Structure," 2014. Forthcoming, *Journal of Financial Economics*.

Danis, Michelle A and Anthony Pennington-Cross, "The Delinquency of Subprime Mortgages," *Journal of Economics and Business*, 2008, *60*, 67–90.

Davidoff, Thomas, "Supply Elasticity and the Housing Cycle of the 2000s," 2011. Forthcoming, *Real Estate Economics*.

___ , "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors," 2014. Working Paper.

Demyanyk, Yuliya and Ottto van Hemert, "Understanding the Subprime Mortgage Crisis," *Review of Financial Studies*, 2011, *24* (6), 1848–1880.

Deng, Yongheng, Joseph Gyourko, and Jing Wu, "Is There Evidence of a Real Estate Collateral Channel Effect on Listed Firm Investment in China?," 2013. NBER Working Paper No. 18762.

Disney, Richard and John Gathergood, "House Price Growth, Collateral Constraints and the Accumulation of Homeowner Debt in the United States," *B.E. Journal of Macroeconomics*, 2011, *11* (1).

Dolde, Walter and Dogan Tirtiroglu, "Temporal and Spatial Information Diffusion in Real Estate Price Changes and Variances," *Real Estate Economics*, 1997, *25* (4), 539–565.

Dungey, Mardi, Renee Fry, Brenda Gonzalez-Hermosillo, and Vance Martin, "Empirical Modeling of Contagion: A Review of Methodologies," *Quantitative Finance*, 2005, *5* (1), 1–16.

Dunsky, Robert M. and James R. Follain, "Tax-Induced Portfolio Reshuffling: The Case of the Mortgage Interest Deduction," *Real Estate Economics*, 2000, *28* (4), 683–718.

Dynan, Karen E. and Donald L. Kohn, "The Rise in U.S. Household Indebtedness: Causes and Consequences," 2007. FEDS Working Paper 2007-37. Board of Governors of the Federal Reserve System.

Engelhardt, Gary V., "House Prices and Home Owner Saving Behavior," *Regional Science and Urban Economics*, 1996, *26* (3-4), 313–336.

Estrella, Arturo, "Critical Values and *P* Values of Bessel Process Distributions: Computation and Application to Structural Break Tests," *Econometric Theory*, 2003, *19* (6), 1128–1143.

Favara, Giovanni and Zheng Song, "House Price Dynamics with Dispersed Information," 2010. Working Paper.

Favilukis, Jack, Sydney Ludvigson, and Stijn Van Nieuwerburgh, "Macroeconomic Implications of Housing Wealth, Housing Finance and Limited Risk Sharing in Equilibrium," 2011. NYU Stern Working Paper.

Federal Housing Finance Agency, "A Strategic Plan for Enterprise Conservatorships: The Next Chapter in a Story that Needs an Ending," February 2012.

— , "FHFA Takes Further Steps to Advance Conservatorship Strategic Plan by Announcing an Increase in Guarantee Fees," December 2013.

Federal Housing Finance Agency Office of Inspector General, "FHFA's Initiative to Reduce the Enterprises' Dominant Position in the Housing Finance System by Raising Gradually Their Guarantee Fees," July 2013.

Ferreira, Fernando and Joseph Gyourko, "Anatomy of the Beginning of the Housing Boom: U.S. Neighborhoods and Metropolitan Areas," 2011. NBER Working Paper No. 17374.

Follain, James R. and Robert M. Dunsky, "The Demand for Mortgage Debt and the Income Tax," *Journal of Housing Research*, 1997, *8* (2), 155–199.

Foote, Chris, Kristopher Gerardi, and Paul Willen, "Negative Equity and the Foreclosure: Theory and Evidence," *Journal of Urban Economics*, 2008, *14* (2), 234–245.

Forbes, Kristin, "The "Big C": Identifying Contagion," 2012. NBER Working Paper No. 18465.

— **and Roberto Rigobon**, "No Contagion, Only Interdependence: Measuring Stock Market Co-movements," *Journal of Finance*, 2002, *57* (5), 2223–2261.

Fuss, Nico B., Roland Rottke, and Bing Zhu, "Correlated Risk and Volatility Spillovers Across U.S. Regional Housing Markets," 2011. Working Paper.

Fuster, Andreas and Basit Zafar, "The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey," 2014. Federal Reserve Bank of New York Staff Report No. 702.

— **and James Vickery**, "Securitization and the Fixed-Rate Mortgage," January 2013. Federal Reserve Bank of New York Staff Report No. 594.

Gan, Jie, "Collateral, Debt Capacity, and Corporate Investment: Evidence from a Natural Experiment," *Journal of Financial Economics*, 2007, *85* (3), 709–734.

— , "Housing Wealth and Consumption Growth: Evidence from a Large Panel of Households," *Review of Financial Studies*, 2010, *23* (6), 2229–2267.

Gelber, Alexander M., Damon Jones, and Daniel W. Sacks, "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test," September 2013. Working Paper.

Gerardi, Kristopher, Adam Shapiro, and Paul Willen, "Subprime Outcomes, Risky Mortgages, Homeownership Experiences and Foreclosures," 2007. Federal Reserve Bank of Boston Working Paper.

Gerardii, Kristopher, Andreas Lehnert, Shane Sherlund, and Paul Willen, "Making Sense of the Subprime Crisis," *Brookings Papers on Economic Activity*, 2008, *Fall* (69-145).

Glaeser, Edward, Joseph Gyourko, Eduardo Morales, and Charles Nathanson, "Housing Dynamics," 2012. Working Paper (previously circulated as NBER Working Paper No. 13704).

Goetzmann, William, Liang Peng, and Jacqueline Yen, "The Subprime Crisis and House Price Appreciation," 2009. NBER Working Paper No. 15334.

Greenspan, Alan and James Kennedy, "Estimates of Home Mortgage Originations, Repayments, and Debt On One-to-Four-Family Residences," 2005. FEDS Working Paper 2005-41. Board of Governors of the Federal Reserve System.

_ **and** _ , "Sources and Uses of Equity Extracted from Homes," *Oxford Review of Economic Policy*, 2008, *24* (1), 120–144.

Gross, David B. and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics*, 2002, *117* (1), 149–185.

Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 1988, *96* (2), 339–357.

Hansen, Bruce E., "Approximate Asymptotic P Values for Structural-Change Tests," *Journal of Business and Economic Statistics*, 1997, *15*, 60–67.

_ , "Testing for Structural Change in Conditional Models," *Journal of Econometrics*, 2000, *97*, 93–115.

Haughwout, Andrew, Ebiere Okah, and Joseph Tracy, "Second Chance: Subprime Mortgage Modification and Re-Default," 2010. Federal Reserve Bank of New York Staff Reports, report no. 417.

Haurin, Donald R. and Stuart S. Rosenthal, "House Price Appreciation, Savings, and Consumer Expenditures," 2006. Working Paper.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd, "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *Review of Economic Studies*, 1997, *64* (4), 605–654.

— , — , **Jeffrey A. Smith, and Petra E. Todd**, "Characterizing Selection Bias Using Experimental Data," *Econometrica*, 1998, *66* (5), 1017–1098.

Hendershott, Patric H. and James D. Shilling, "The Impact of the Agencies on Conventional Fixed-Rate Mortgage Yields," *Journal of Real Estate Finance and Economics*, 1989, *2* (101-115).

Hickey, Robert, Lisa Sturtevant, and Emily Thaden, "Achieving Lasting Affordability through Inclusionary Housing," 2014. Lincoln Institute of Land Policy Working Paper WP14RH1.

Ho, Giang and Anthony Pennington-Cross, "Predatory Lending Laws and the Cost of Credit," *Real Estate Economics*, 2008, *36* (2), 175–211.

Holly, Sean M., Hashem Pesaran, and Takashi Yamagata, "A Spatio-Temporal Model of House Prices in the USA," *Journal of Econometrics*, 2010, *158*, 160–173.

— , — , **and** — , "The Spatial and Temporal Diffusion of House Prices in the UK," *Journal of Urban Economics*, 2011, *69*, 2–23.

Hurst, Erik and Annamaria Lusardi, "Liquidity Constraints, Household Wealth, and Entrepreneurship," *Journal of Political Economy*, 2004, *112* (2), 319–347.

— **and Frank Stafford**, "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption," *Journal of Money, Credit and Banking*, 2004, *36* (6), 985–1014.

Iacoviello, Matteo, "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 2005, *95* (3), 739–764.

Imbens, Guido W., "Nonparametric Estimation of Average Treatment Effects Under Exogeneity: A Review," *Review of Economics and Statistics*, 2004, *86* (1), 4–29.

Jaffee, Dwight and John M. Quigley, "The Future of the Government Sponsored Enterprises: The Role for Government in the U.S. Mortgage Market," in Edward Glaeser and Todd Sinai, eds., *Housing and the Financial Crisis*, University of Chicago Press, 2012.

Jappelli, Tullio, "Who is Credit-Constrained in the US Economy?," *Quarterly Journal of Economics*, 1990, *105* (1), 219–234.

— **and Luigi Pistaferri**, "Do People Respond to Tax Incentives? An Analysis of the Italian Reform to the Deductibility of Home Mortgage Interest," *European Economic Review*, 2007, *51*, 247–271.

— **and** — , "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2* (1), 479–506.

— , **Jörn-Steffen Pischke, and Nicholas S. Souleles**, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," *Review of Economics and Statistics*, 1998, *80* (2), 251–262.

Jensen, Thais Lærkholm, Søren Leth-Petersen, and Ramana Nanda, "Housing Collateral, Credit Constraints and Entrepreneurship - Evidence from a Mortgage Reform," 2014. NBER Working Paper No. 20583.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Karlan, Dean and Jonathan Zinman, "Long-Run Price Elasticities of Demand for Credit: Evidence from a Countrywide Field Experiment in Mexico," 2014. Working Paper.

Karlan, Dean S. and Jonathan Zinman, "Credit Elasticities in Less-Developed Economies: Implications for Microfinance," *American Economic Review*, 2008, *98* (3), 1040–1068.

Kaufman, Alex, "The Influence of Fannie and Freddie on Mortgage Loan Terms," 2012. Federal Reserve Board Finance and Economics Discussion Series.

Kennedy, Peter E., "Estimation with Correctly Interpreted Dummy Variables in Semilogarithmic Equations," *American Economic Review*, 1981, *71* (4), 801.

Keys, Benjamin J., Devin Pope, and Jaren Pope, "Failure to Refinance," 2014. NBER Working Paper No. 20401.

Keys, Benjamin, Tanmoy Mukherjee, Amit Seru, and Vikrankt Vig, "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," *Quarterly Journal of Economics*, 2010, *125* (1), 307–362.

King, Mervyn and Sushil Wadhwani, "Transmission of Volatility between Stock Markets," *The Review of Financial Studies*, 1990, *3* (1), 5–33.

Kiyotaki, Nobuhiro and John Moore, "Credit Cycles," *Journal of Political Economy*, 1997, *105* (211-248).

Kleven, Henrik J. and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics*, 2013, *128* (2), 669–723.

Kopczuk, Wojciech and David J. Munroe, "Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market," May 2014. Working Paper.

Lai, Rose and Robert van Order, "Momentum and House Price Growth in the United States: Anatomy of a Bubble," *Real Estate EconomicsE*, 2010, *38* (4), 753–773.

Laufer, Steven, "Equity Extraction and Mortgage Default," 2014. Working Paper (previously circulated as FEDS Working Paper 2013-30. Board of Governors of the Federal Reserve System.).

Leamer, Edward E., *Specification Searches: Ad Hoc Interence with Nonexperimental Data*, Wiley, 1978.

Lee, Sang B. and Kwang Jung Kim, "Does the October 1987 Crash Strengthen the Comovements Among National Stock Markets?," *Review of Financial Economics*, 1993, *3*, 89–102.

Lehnert, Andreas, "Housing, Consumption, and Credit Constraints," 2004. FEDS Working Paper 2004-63. Board of Governors of the Federal Reserve System.

Leth-Petersen, Søren, "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?," *American Economic Review*, 2010, *100* (3), 1080–1103.

Ling, David C. and Gary A. McGill, "Evidence on the Demand for Mortgage Debt by Owner-Occupants," *Journal of Urban Economics*, 1998, *44* (3), 391–414.

Lustig, Hanno and Stijn Van Nieuwerburgh, "How Much Does Household Collateral Constrain Regional Risk Sharing?," *Review of Economic Dynamics*, 2010, *13* (2), 265–294.

Manoli, Dayanand S. and Andrea Weber, "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions," August 2011. NBER Working Paper No. 17320.

Martins, Nuno C. and Ernesto Villanueva, "The Impact of Mortgage Interest-Rate Subsidies on Household Borrowing," *Journal of Public Economics*, 2006, *90* (8-9), 1601–1623.

Mayer, Christopher and Karen Pence, "Subprime Mortgages: What, Where and to Whom?," 2008. Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C., No. 2008-29.

_ , _ , **and Shane Sherlund**, "The Rise of Mortgage Defaults," *Journal of Economic Perspectives*, 2009, *23* (1), 27–50.

Mian, Atif, Amir Sufi, and Francesco Trebbi, "The Political Economy of the U.S. Mortgage Default Crisis," *American Economic Review*, 2010, *100* (6), 1967–1988.

_ **and** _ , "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *Quarterly Journal of Economics*, 2009, *124* (4), 1449–1496.

_ **and** _ , "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 2011, *101* (5), 2132–2156.

_ **and** _ , "House Price Gains and U.S. Household Spending from 2002 to 2006," 2014. NBER Working Paper No. 20152.

_ , **Kamalesh Rao, and Amir Sufi**, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

Mishkin, Frederic S., "Symposium on the Monetary Transmission Mechanism," *Journal of Economic Perspectives*, 1995, *9* (4), 3–10.

Ortalo-Magne, Francois and Sven Rady, "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints," *The Review of Financial Studies*, 2006, *73* (2), 459–85.

Parker, Jonathan A., "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 1999, *89* (4), 959–973.

— , **Nicholas S. Souleles, David S. Johnson, and Robert McClelland**, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–2553.

Passmore, Wayne, Roger Sparks, and Jamie Ingpen, "GSEs, Mortgage Rates, and the Long-Run Effects of Mortgage Securitization," *Journal of Real Estate Finance and Economics*, 2002, *25* (2), 215–242.

— , **Shane M. Sherlund, and Gillian Burgess**, "The Effect of Housing Government Sponsored Enterprises on Mortgage Rates," *Real Estate Economics*, 2005, *33* (3), 427–463.

Piskorski, Tomacz, Amit Seru, and Vikrankt Vig, "Securitization and Distressed Loan Renegotiation: Evidence from the Subprime Mortgage Crisis," 2009. Chicago Booth School of Business Research Paper No. 09-02.

Poterba, James, "Tax Subsides to Owner-Occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics*, 1984, *99* (4), 729–752.

— **and Todd Sinai**, "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income," *American Economic Review Papers and Proceedings*, 2008, *98* (2), 84–89.

— **and** — , "Revenue Costs and Incentive Effects of the Mortgage Interest Deduction for Owner-Occupied Housing," *National Tax Journal*, 2011, *64*, 531–564.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economyol*, 1982, *90* (4), 1257–78.

Rosen, Sherwin, "Wage-Based Indexes of Urban Quality of Life," in Peter Mieszkowski and Mahlon Straszheim, eds., *Current Issues in Urban Economics*, Baltimore: Johns Hopkins University Press, 1979.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points," *American Economic Journal: Economic Policy*, 2010, *2*, 180–212.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Sallee, James M. and Joel Slemrod, "Car Notches: Strategic Automaker Responses to Fuel Economy Policy," December 2010. NBER Working Paper No. 16604.

Schmalz, Martin C., David A. Sraer, and David Thesmar, "Housing Collateral and Entrepreneurship," 2013. NBER Working Paper No. 19680.

Shapiro, Matthew D. and Joel Slemrod, "Consumer Responses to the Timing of Income: Evidence from a Change in Tax Withholding," *American Economic Review*, 1995, *85* (1), 274–283.

Sherlund, Shane M., "The Jumbo-Conforming Spread: A Semiparametric Approach," 2008. Federal Reserve Board Finance and Economics Discussion Series.

Shiller, Robert, *Irrational Exuberance*, 2 ed., Princeton, NJ: Princeton University Press, 2005.

Sinai, Todd, "House Price Moments in Boom-Bust Cycles," 2012. NBER Working Paper No. 18059.

__ **and Nicholas S. Souleles**, "Owner-Occupied Housing as a Hedge Against Rent Risk," *Quarterly Journal of Economics*, 2005, *120* (2), 763–789.

__ **and Nicholas Souleles**, "Can Owning a Home Hedge the Risk of Moving?," 2009. NBER Working Paper No. 15462.

Skinner, Jonathan S., "Housing Wealth and Aggregate Saving," *Regional Science and Urban Economics*, 1989, *19* (2), 305–324.

__ , "Is Housing Wealth a Sideshow?," in David A. Wise, ed., *Advances in the Economics of Aging*, University of Chicago Press, 1996, pp. 241–272.

Smith, Jeffrey A. and Petra E. Todd, "Does Matching Overcome LaLonde's Critique of Non-experimental Estimators?," *Journal of Econometrics*, 2005, *125* (1-2), 305–353.

Souleles, Nicholas S., "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 1999, *89* (4), 947–958.

Stephens, Melvin Jr., "The Consumption Response to Predictable Changes in Discretionary Income: Evidence from the Repayment of Vehicle Loans," *Review of Economics and Statistics*, 2008, *90* (2), 241–252.

Todd, Petra E., "A Practical Guide to Implementing Matching Estimators," 1999. Working Paper.

Wheaton, William and Gleb Nechayev, "The 1998-2005 Housing 'Bubble' and the Current 'Correction': What's Different This Time?," *Journal of Real Estate Research*, 2008, *30* (1), 1–26.

Yamashita, Takashi, "House Price Appreciation, Liquidity Constraints, and Second Mortgages," *Journal of Urban Economics*, 2007, *62* (3), 424–440.

Young, Peyton H., "Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning," *American Economic Review*, 2009, *99* (5), 1899–1924.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

Zhang, C. Yiwei, "Consumption Responses to Pay Frequency: Evidence from 'Extra' Paychecks," 2014. Working Paper.

Zinman, Jonathan, "Household Debt," *Annual Review of Economics*, 2015, *7* (1).