

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**The Interest Rate Elasticity of Mortgage Demand: Evidence
From Bunching at the Conforming Loan Limit**

Anthony A. DeFusco and Andrew Paciorek

2014-11

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

THE INTEREST RATE ELASTICITY OF MORTGAGE DEMAND: EVIDENCE FROM BUNCHING AT THE CONFORMING LOAN LIMIT*

Anthony A. DeFusco[†] Andrew Paciorek[‡]

January 15, 2014

Abstract

The relationship between the mortgage interest rate and a household’s demand for mortgage debt has important implications for a host of public policy questions. In this paper, we use detailed data on over 2.7 million mortgages to provide novel estimates of the interest rate elasticity of mortgage demand. Our empirical strategy exploits a discrete jump in interest rates generated by the conforming loan limit—the maximum loan size eligible for securitization by Fannie Mae and Freddie Mac. This discontinuity creates a large “notch” in the intertemporal budget constraint of prospective mortgage borrowers, allowing us to identify the causal link between interest rates and mortgage demand by measuring the extent to which loan amounts bunch at the conforming limit. Under our preferred specifications, we estimate 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. We also present evidence that about one third of the response is driven by borrowers who take out second mortgages while leaving their total mortgage balance unchanged. Accounting for these borrowers suggests a reduction in total mortgage debt of between 1.5 and 2 percent per percentage point increase in the interest rate. Using these estimates, we predict the changes in mortgage demand implied by past and proposed future increases to the guarantee fees charged by Fannie and Freddie. We conclude that these increases would directly reduce the dollar volume of new mortgage originations by well under 1 percent.

*We are grateful to Manuel Adelino, Neil Bhutta, Gilles Duranton, Fernando Ferreira, Joseph Gyourko, Raven Molloy, Dan Sacks, Todd Sinai, Yiwei Zhang and seminar participants at the Board of Governors and the 2014 AREUEA-ASSA Conference for helpful comments. Corinne Land provided excellent research assistance. All remaining errors are our own. Anthony DeFusco gratefully acknowledges the financial support provided by the Wharton Risk Management and Decision Process Center through the Russel Ackoff Doctoral Fellowship. The views we express herein are not necessarily those of the Board of Governors or its staff.

[†]The Wharton School, University of Pennsylvania. E-mail: defusco@wharton.upenn.edu

[‡]Board of Governors of the Federal Reserve System. E-mail: andrew.d.paciorek@frb.gov

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.

Yet, despite these potentially important policy implications, there are relatively few empirical estimates of the extent to which individual loan sizes respond to changes in interest rates. This is due in large part to data limitations, which have led prior research in this area to focus on other aspects of mortgage choice or to rely on endogenous variation in interest rates (Follain and Dunskey, 1997; Gary-Bobo and Larriveau, 2004; Martins and Villanueva, 2006; Jappelli and Pistaferri, 2007). 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; Karlan and Zinman, 2008; Attanasio et al., 2008). 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

leveraging “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 below the limit. 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.

Recent papers in public finance have developed methods for estimating behavioral responses to nonlinear incentives in similar settings (Saez, 2010; Chetty et al., 2011; Kleven and Waseem, 2013).¹ We adapt these methods to the case of mortgage choice in the face of a notched interest rate schedule. Intuitively, the excess mass of households who bunch at exactly the conforming limit provides us with a measure of the behavioral response to the interest rate differential. We combine this estimate of bunching with estimates of the interest rate spread between jumbo and conforming loans to yield an estimate of the average interest rate (semi-)elasticity of mortgage demand.² 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 mortgage declines by between 2 and 3 percent for each 1 percentage point rise in the mortgage rate. Because both the bunching estimates and the jumbo-conforming spread estimates vary depending on the assumptions used in estimation, we 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. We also discuss heterogeneity in the responsiveness of different groups, as well as the implications of fixed adjustment costs and extensive margin responses—buyers dropping out of the market entirely—for the interpretation of our estimates and their external validity.

¹Other recent applications of these and similar methods include Sallee and Slemrod (2010); Manoli and Weber (2011); Best and Kleven (2013); Chetty et al. (2013); Gelber et al. (2013) and Kopczuk and Munroe (2013).

²More formally, our goal is to estimate the elasticity of mortgage demand with respect to the rate on the first mortgage, holding all other prices and interest rates constant.

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. We also argue that the pattern of loan-to-value ratios (LTVs) around the limit suggests that the remaining two-thirds are putting up more cash rather than buying cheaper houses.

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 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. In section 2 we provide relevant institutional details on the GSEs and the conforming loan limit. Section 3 presents our conceptual framework. In sections 4 and 5 we discuss our data and empirical research design. We then present our main results in sections 6-8. Section 9 applies these results to changes in the GSE guarantee fees and section 10 concludes by discussing avenues for future research.

2 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,

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.³ 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 1 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.⁴ Following the 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.⁵ Part of this savings is eventually passed on to borrowers in the form of lower interest rates on conforming loans.⁶ 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), as well as exemptions

³As 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).

⁴The 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.

⁵The implicit guarantee became explicit in 2008 when the GSEs were placed under government conservatorship.

⁶Passmore 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.

from state and local income taxes.⁷ 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.

The difference in interest rates between loans above and below the conforming limit is called the jumbo-conforming spread. Even with good mortgage data, identifying the spread is challenging because borrowers are likely to sort themselves around it, leading to differences in borrower characteristics that may or may not be observable.⁸ Although we address these issues in detail below in section 5, some insight can still be gleaned from examining the raw data. For example, figure 2 plots the interest rate for all fixed-rate mortgages in our analysis sample that were originated in 2006 as a function of the difference between the loan amount and the conforming limit.⁹ 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. There is a clear discontinuity precisely at the limit, with average interest rates on loans just above the limit being approximately 20 basis points higher than those on loans just below the limit. While 20 basis points may not reflect the “true” jumbo-conforming spread due to sorting around the limit, this figure is at least suggestive evidence of a sharp change in the cost of credit as loan size crosses the threshold.

Regardless of the precise size of the jumbo-conforming spread, its existence introduces a nonlinearity in the budget constraint of an individual deciding how much mortgage debt to incur. This nonlinearity induces borrowers who would otherwise take out loans above the conforming limit to bunch at the limit, perhaps by putting up a larger down payment or taking out a second loan. The histogram in figure 3 confirms this, showing 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. 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 precisely at the conforming limit to back out estimates of the

⁷For a full description of the direct benefits conferred on the GSEs as a result of their special legal status see [Congressional Budget Office \(2001\)](#).

⁸Many papers have attempted to overcome this challenge, using a variety of different empirical methods. See, for example, [Hendershott and Shilling \(1989\)](#), [Passmore et al. \(2002\)](#), [Passmore et al. \(2005\)](#), [Sherlund \(2008\)](#) and [Kaufman \(2012\)](#).

⁹See section 4 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 5.

interest rate elasticity of demand for mortgage debt. The next section provides a conceptual framework that we use to formalize this intuition.

3 Theoretical Framework

We begin by considering a simple two-period model of household mortgage choice.¹⁰ 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 somewhat more analogous to ours, where the budget constraint features a notch as opposed to a kink. Applications of this framework include fuel economy regulation ([Sallee and Slemrod, 2010](#)), retirement incentives ([Manoli and Weber, 2011](#)), income taxes ([Kleven and Waseem, 2013](#)), and real estate transfer taxes ([Best and Kleven, 2013](#); [Kopczuk and Munroe, 2013](#)). Ours is the first application to the mortgage market, or to a credit market of any kind.

3.1 Baseline Case: Linear Interest Rate Schedule

Households live for two periods. In our baseline model, 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 .¹¹ 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

¹⁰The underlying theory is similar to that in [Brueckner \(1994\)](#), among other papers.

¹¹Below, we relax the assumption that households cannot choose the quantity of housing services to consume.

period, denoted by c_1 and c_2 .¹² In general, the household solves:

$$\max_{c_1, c_2} \{U(c_1, c_2) = u(c_1) + \delta u(c_2)\} \quad (1)$$

$$\text{s.t. } c_1 + p = y + m \quad (2)$$

$$c_2 = p - (1 + r) m \quad (3)$$

$$0 \leq m \leq p, \quad (4)$$

where $\delta \in (0, 1)$ is the discount factor and y is first period income. Solving equation (2) for c_1 and substituting this, along with equation (3), into equation (1) allows us to rewrite the household's problem in terms of mortgage debt,

$$V = \max_m \{u(y + m - p) + \delta u(p - (1 + r) m)\}, \quad (5)$$

subject now only to the borrowing constraint (4).

To proceed, we make several simplifying assumptions. First, we assume that household preferences are given by the constant elasticity function $u(c) = \frac{1}{1-\xi} c^{1-\xi}$.¹³ Second, heterogeneity in the model is driven by the discount factor, which is assumed to be distributed smoothly in the population according to the distribution function $F(\delta)$ and density function $f(\delta)$. For illustrative purposes, we assume that y 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. Finally, we assume that households end up at an interior solution with a positive mortgage amount and a loan-to-value ratio of less than 100 percent—that is, constraint (4) does not bind.

Under these assumptions, we can solve explicitly for mortgage demand, which is given by:

$$m^* = \frac{p - (\delta(1+r))^{1/\xi} (y-p)}{(\delta(1+r))^{1/\xi} + (1+r)}. \quad (6)$$

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.¹⁴ Given the assumption of a smooth

¹²Since 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.

¹³This functional form allows us to derive a closed-form solution, but all of the basic results hold with more general utility functions.

¹⁴Technically, for this mapping to be one-to-one it must be true that $y > \frac{r}{1+r}p$. If this condition holds

distribution for δ , this mapping will induce a smooth baseline distribution of mortgage amounts, which we denote using the CDF, $G_0(m)$, density function, $g_0(m)$.

3.2 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 for reasons discussed in section 2, leading to the new schedule $r(m) = r + \Delta r \cdot \mathbb{1}(m > \bar{m})$. 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) and (3) yields the lifetime budget constraint

$$C = y - m \cdot [r + \Delta r \cdot \mathbb{1}(m > \bar{m})], \quad (7)$$

where $C = c_1 + c_2$ is lifetime consumption. This budget constraint is plotted in figure 4a 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 4a, 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 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.

This means that the density when a notch exists, $g_1(m)$, will 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 4b. 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

 then m^* is strictly decreasing in δ . This is likely to be the case for any reasonable values of r and p .

size distribution, the number of households bunching at the conforming limit is given by:

$$B = \int_{\bar{m}}^{\bar{m} + \Delta\bar{m}} g_0(m) dm \approx g_0(\bar{m}) \Delta\bar{m}, \quad (8)$$

where the approximation assumes that the counterfactual no-notch distribution is constant on the bunching interval $(\bar{m}, \bar{m} + \Delta\bar{m})$.¹⁵ This expression 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.

It is worth emphasizing that 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 of mortgage sizes be smooth. Any model for which these conditions hold would generate equation (8).

3.3 Heterogeneous Intertemporal Elasticities and Incomes

The derivation of equation (8) was under the assumption that ξ and y were constant across households. In that case, it was possible to back out the exact change in mortgage amount for the marginal bunching individual. When intertemporal elasticities and incomes 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 and income level. To see this, let the joint distribution of discount factors, intertemporal elasticities, and incomes be given by $\bar{f}(\delta, \xi, y)$, where $y \in (0, \bar{y}]$ and $\xi \in (0, \bar{\xi}]$ for some upper bounds, \bar{y} and $\bar{\xi}$. For a fixed (ξ, y) pair, the bunching interval is determined in exactly the same way as in the baseline model. Denote this interval $(\bar{m}, \bar{m} + \Delta\bar{m}_{\xi, y})$, where $\Delta\bar{m}_{\xi, y}$ is the behavioral response of the marginal bunching individual among those with intertemporal elasticity $1/\xi$ and income y . Further, let $\bar{g}_0(m, \xi, y)$ denote the joint distribution of mortgage sizes, intertemporal elasticities, and incomes in the pre-notch scenario and $g_0(m) \equiv \int_{\xi} \int_y \bar{g}_0(m, \xi, y) dy d\xi$ the unconditional mortgage size distribution. The amount

¹⁵This approximation merely simplifies the discussion. In the empirical application we allow for curvature in the counterfactual distribution.

of bunching can then be expressed as

$$B = \int_{\xi} \int_y \int_{\bar{m}}^{\bar{m} + \Delta \bar{m}_{\xi, y}} \bar{g}_0(m, \xi, y) dm dy d\xi \approx g_0(\bar{m}) E[\Delta \bar{m}_{\xi, y}]. \quad (9)$$

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

3.4 Endogenous Housing Choice

With the choice of housing fixed, as in the discussion above, 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 (5), the household's intertemporal optimization problem. Households can now choose the quantity of housing services to purchase (h), and this quantity has a direct effect on first-period utility, so that

$$V = \max_{m, h} \{u(y + m - ph, h) + \delta v(ph - (1 + r)m)\}, \quad (10)$$

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:

$$\frac{\partial V}{\partial m} = u_1 - \delta(1 + r)v_1 = 0 \quad (11)$$

$$\frac{\partial V}{\partial h} = u_2 - (pu_1 - p\delta v_1) = 0. \quad (12)$$

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,

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

against its consumption value today.

While there are no obvious functional forms that allow us to derive equivalents to equation (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^*).

4 Data

To conduct our empirical analysis, we use data on loan sizes and interest rates that come 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 municipalities across the U.S. These data serve as our primary source of information on loan size. For descriptive purposes, we have also matched the DQ data to loan application information made available through the Home Mortgage Disclosure Act (HMDA), which provides us with a limited set of borrower demographics.¹⁷ 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.

4.1 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.

¹⁷The 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 an observation in HMDA. Further details are available from the authors on request.

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 2008, when the GSEs were taken into government conservatorship, the standards for determining the conforming limit were changed in several ways, including a provision which allows it to vary across different metropolitan areas.

Another 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 led 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).¹⁸

These restrictions leave us with a primary estimation sample of approximately 2.7 million transactions representing 26 MSAs. Table 1 presents summary statistics for this sample 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 leads many of the differences between columns 1 and 3 to be 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.

¹⁸We also drop extreme outliers in appraisal value or LTV ratio.

4.2 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.¹⁹ 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.

Table 2 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.

5 Empirical Methodology

5.1 Estimating the Behavioral Response to the Conforming Limit

In section 3, 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 $j = -J, \dots, L, \dots, 0, \dots, U, \dots, J$,

¹⁹Although 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.

and count the number of loans in each bin, n_j . To obtain estimates of bunching and the counterfactual loan size distribution we define a 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

$$n_j = \sum_{i=0}^p \beta_i (m_j)^i + \sum_{k=L}^U \gamma_k \mathbb{1}(m_k = m_j) + \epsilon_j. \quad (13)$$

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

$$\hat{n}_j = \sum_{i=0}^p \hat{\beta}_i (m_j)^i. \quad (14)$$

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,

$$\hat{B} = \sum_{j=L}^0 (n_j - \hat{n}_j) = \sum_{j=L}^0 \hat{\gamma}_j. \quad (15)$$

This procedure is illustrated graphically in figure 5. The solid black line represents the empirical count of loans in each bin, the heavy dashed red line is the estimated counterfactual distribution, the solid shaded gray area is the bunching estimate, and the cross hatched shaded gray area is the amount of missing mass due to bunching, $\hat{M} = \sum_{j>0}^U (n_j - \hat{n}_j) = \sum_{j>0}^U \hat{\gamma}_j$.

The parameter of primary interest is $\Delta \hat{m}$, the empirical analogue of $\Delta \bar{m}$ from equation (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

$$\Delta \hat{m} = \frac{\hat{B}}{\hat{g}_0(\bar{m})}, \quad (16)$$

where $\hat{g}_0(\bar{m}) = \sum_{j=L}^0 (\hat{n}_j) / \left| \frac{m_0 - m_L}{L} \right|$ 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.

There are two key identifying assumptions necessary for equation (16) 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 would be smooth. That is, any spike in the loan size distribution at the conforming limit can be solely attributed 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 imply such a result.

In order to estimate the components of equation (16), there are several free parameters that we must choose: the bin width ($|\frac{m_0 - m_L}{L}|$), 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}). 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.

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. 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. In fact, the theory predicts that the two will likely differ.

As noted by Kleven and Waseem (2013), 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 responses among those who do not bunch. In response to the higher interest rate, some borrowers who would have located to the right of the limit may instead choose not to purchase a home at all (extensive margin responses), and jumbo borrowers who are not induced to bunch will still presumably choose to borrow slightly less than they would at conforming rates (intensive margin responses). If these types of responses have a large enough 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, most importantly, give very similar estimates of $\Delta\hat{m}$ to our preferred specification.

Finally, we calculate standard errors for all estimated parameters using a bootstrap procedure, as in [Chetty et al. \(2011\)](#). At each iteration (k) of the bootstrap loop we draw with replacement from the estimated errors, $\hat{\epsilon}_j$, in equation (13) to generate a new set of bin counts, n_j^k . We then re-estimate the amount of bunching using these new counts. Our estimate of the standard error for $\Delta\hat{m}$ is the standard deviation of the estimated $\Delta\hat{m}^k$ s. The same procedure produces standard errors for all the other bunching parameters that we report.

5.2 Estimating the Jumbo-Conforming Spread

Although our estimates of bunching provide a reliable measure of the behavioral response to the conforming loan limit, in order to convert that response 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.

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 flexibly 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 the appraised 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 semiparametrics 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

$$r_{i,t} = \alpha_{z(i),t} + \beta J_{i,t} + f^{J=0}(m_{i,t}) + f^{J=1}(m_{i,t}) + s^{LTV}(LTV_{i,t}) + s^{DTI}(DTI_{i,t}) + s^{FICO}(FICO_{i,t}) + PMI_{i,t} + PP_{i,t} + g(TERM_{i,t}) + \epsilon_{i,t}, \quad (17)$$

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 ($f^{J=0}(m_{i,t})$ and $f^{J=1}(m_{i,t})$) 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 indicating 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 ($TERM$).²⁰ 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.

If there are other unobserved determinants of interest rates which are also correlated with jumbo loan status, than estimates of β based on equation (17) 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 (17) in which we instrument for jumbo loan status using a discontinuous function of the appraised value of the home, following [Kaufman \(2012\)](#).²¹ Because mortgage contracts are frequently determined prior to the actual date of transaction, the official loan-to-value (LTV) ratio used by the bank to determine whether a borrower qualifies for a loan is often set based on an independent appraisal value, not the actual transaction price. Moreover, since many homebuyers purchase a home with an LTV of exactly 80 percent, if a home appraisal comes in just over the conforming loan limit divided by 0.8, then a buyer is substantially more likely to take out a jumbo loan. This suggests an approach in which we instrument for $J_{i,t}$ in equation (17) with whether an appraisal is above or below $\frac{\bar{m}}{0.8}$. The key to this ‘‘appraisal limit’’ being a valid instrument is that, unlike

²⁰The exact specifications are described in the results section below.

²¹[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.

their actual loan amount, borrowers likely have little control over the exact outcome of their appraisal.

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 in response to a high appraisal. 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 high appraisal. 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.

6 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 7](#) to calculate elasticities. In [section 8](#), we conclude with a discussion on the ways in which borrowers appear to be adjusting their loan sizes.

6.1 Bunching at the Conforming Limit

6.1.1 Results for all Borrowers

As a starting point for our empirical analysis, [figure 6](#) 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 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 0 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 precisely 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 5.1.

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 3 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 3 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 3, 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.

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 7 plots the empirical loan size distribution separately for the years 2000, 2002, and 2004.²² 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.

6.1.2 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. 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 replicate this stylized fact in figure 8 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.²³ 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 8 might suggest an extensive margin response. That is, in response to the higher jumbo spread for FRMs some borrowers may choose to substitute toward

²²We avoid plotting consecutive years to preserve the clarity of the picture, but the results hold at any horizon.

²³We replicate this well-documented difference in spreads using our own sample of loans from LPS in section 6.2. 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.

adjustable-rate mortgages.

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, and not because individual borrowers are choosing ARMs over FRMs. 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 substituting 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.

Figures 9 and 10 show the results from this exercise for fixed and adjustable-rate mortgages, respectively. As with the combined results, the standard errors and additional bunching parameters are also reported separately for each loan type in columns 2 and 3 of table 3. While figure 9 shows a substantial downward shift in the FRM distribution to the right of the limit, figure 10 shows no corresponding upward shift in the ARM distribution. In fact, much like in the figures for the FRM and combined samples, the ARM 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. These results suggest, at least in aggregate, that the drop in the FRM share at the conforming limit is driven primarily by differential bunching *within* loan type rather than borrowers substituting from one type of mortgage to another.²⁴ This is further reflected in the estimate of excess mass reported in table 3, which is roughly 4.4 times larger for FRMs than for ARMs.

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.

6.1.3 Heterogeneity by Borrower Type

In addition to investigating bunching behavior by loan type, 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 using the subset of DQ transactions that

²⁴Of 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.

we matched to HMDA loan applications. For this exercise we restrict attention to fixed-rate mortgages, where the sample sizes are largest. Using the race and ethnicity information in 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 median income reported across all loans and “high-income” if the reported income is above the median.

Figure 11 shows results from estimating bunching separately in each of these four subsamples. 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.²⁵ Table 4 confirms the visual impressions given by figure 11, 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.

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

²⁵Part 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.

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 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 below 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.

6.2 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 5](#) presents estimates of the jumbo-conforming spread, following the strategies discussed in [section 5.2](#). 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. Analogous estimates of the spread, estimated in logs, are provided in [table 6](#). These estimates are used in calculating the elasticities below, but we discuss the results in levels here because it is more intuitive to think of changes in interest rates in terms of basis or percentage points.

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

For fixed-rate mortgages, 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.²⁶

However, as discussed above, the OLS specifications do not control for borrower selection around the conforming limit. As we have just shown, a large fraction of borrowers who would otherwise get loans just above the limit instead choose to bunch at the limit to get conforming loans and reduce their rates. If these borrowers differ in any unobserved ways from those who do not reduce their rates, the estimates of the spread could be biased.

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 appraised value of the house exceeds the “appraisal limit” (the conforming limit divided by 0.8). As would be expected, these estimates are somewhat less precise than the OLS results but are nonetheless still significant 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.²⁷ 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.²⁸

Our estimates of the jumbo-conforming spread for initial ARM rates, presented in the lower half of table 5, 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, we saw in the previous section that a significant fraction of borrowers with ARM balances that would be just above the limit under the counterfactual instead choose to bunch at the limit, although the effect is smaller than in

²⁶Specifically, [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.

²⁷Regressing the jumbo limit indicator on the appraisal limit indicator, 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 are tiny and the F-statistics very large.

²⁸[Kaufman \(2012\)](#) uses LPS data that cover the entire U.S. from 2003 to 2007.

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 “appraisal limit” 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.²⁹

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.

7 Elasticities

7.1 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 the change in the marginal price of the loan.

Figure 12 illustrates this point using numbers from 2006 as an example. Panel (a) shows the overall interest rate on a 30-year FRM given the size of the loan.³⁰ The rate jumps at the conforming loan limit of \$417,000—the dashed blue line—from 6.38 percent to 6.55

²⁹Kaufman (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.

³⁰For the purposes of this example, we assume that the relationship between loan size and interest rates is constant, except at the limit where it jumps discretely. In practice we only need to assume a *locally* constant relationship near the limit.

percent.³¹ Panel (b) shows the equivalent monthly mortgage payment, which jumps by about \$45 at the limit.³²

Although the marginal increase in payment per dollar of the loan—the slope of the line—is only slightly higher above the limit than below it, the discontinuity implies a much more substantial effect on incentives. To see this, let $P(m, r)$ denote the monthly payment on a 30-year mortgage of amount m at fixed interest rate r . Then the total increase in payment for a loan of size $m > \bar{m}$ relative to a loan at the conforming limit, \bar{m} , is equal to

$$P(m, \hat{r} + \Delta\hat{r}) - P(\bar{m}, \hat{r}), \quad (18)$$

where \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

$$P^*(m) = \frac{P(m, \hat{r} + \Delta\hat{r}) - P(\bar{m}, \hat{r})}{m - \bar{m}} \quad (19)$$

Panel (d) of figure 12 plots $P^*(m)$ for a range of different mortgage sizes. A loan that is \$1,000 over the limit faces an increase in monthly payment of about \$53, or a P^* of 5.3 cents per dollar of the loan above the limit. The marginal price is quite high just above the limit, so it gives a strong incentive for borrowers to reduce their loan balances. This incentive declines as the loan amount rises. Our estimate of $\Delta\hat{m} = .0627$ for FRMs from column 2 of table 3 implies a counterfactual loan size about \$26,000 over the limit, denoted in the figure with a solid blue line. At this point, the increase in payment relative to the limit totals about \$210, or a P^* of 0.8 cents per dollar of the loan. This amount is still about 0.15 cents higher than the P^* on loans far above the limit, for which the notch is essentially irrelevant.

While the monthly payment is an intuitive metric, the correct theoretical price per dollar of the loan is the underlying interest rate. Define $r^*(m)$ such that

$$(m - \bar{m}) \cdot r^*(m) = m \cdot (\hat{r} + \Delta\hat{r}) - \bar{m} \cdot \hat{r} \quad (20)$$

This $r^*(m)$ is the implicit interest rate on the loan amount in excess of the conforming limit

³¹We estimate that 6.38 percent was the average rate paid on loans just under the conforming limit in 2006, while the 0.17 percentage point differential comes from the OLS estimate in column 3 of table 5. Since this is just an example, we use the estimate of the spread for the entire sample rather than estimating it for just 2006.

³²We use a standard formula to compute the monthly payment given these parameters: $P(m, r) = m \frac{\frac{r}{12}(\frac{r}{12} + 1)^{360}}{(\frac{r}{12} + 1)^{360} - 1}$, where m is the mortgage amount and r is the annual interest rate.

$(m - \bar{m})$, taking into account the jump in the overall rate.³³ Solving explicitly for $r^*(m)$ yields

$$r^*(m) = \hat{r} + \Delta\hat{r} + \Delta\hat{r} \cdot \frac{\bar{m}}{m - \bar{m}}. \quad (21)$$

Equation (21) makes clear that $r^*(m)$ is equal to the jumbo rate ($\hat{r} + \Delta\hat{r}$) plus a term that is increasing 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.

Panel (c) of figure 12 plots $r^*(m)$ for our 2006 example. We calculate that $r^*(m)$ for a loan \$1,000 over the limit to be about 77 percent. This implies that, all else equal, borrowers in 2006 who took out a jumbo loan for \$418,000 at $\hat{r} + \Delta\hat{r} = 6.55\%$ must have weakly preferred that loan to the combination of a conforming loan for \$417,000 at $\hat{r} = 6.38\%$ and a second mortgage at $r^*(m) = 77\%$. Of course, second mortgages in 2006 were available at far lower interest rates, so the average borrower who wanted a \$418,000 mortgage would have been much better off with the conforming/second combination, ignoring any additional costs associated with the second mortgage.³⁵

As the borrower's desired mortgage balance increases, however, $r^*(m)$ declines. At our estimate of $\Delta\hat{m} = .0627$, which implies a loan size of \$443,000, we calculate $r^*(m)$ to be 9.3 percent, which is about 2.9 percentage points above the conforming rate of 6.38 percent. By comparison, in a sample of matched first and second 30-year FRMs from 2006 in LPS, we find that the second mortgage has a rate about 2.4 percentage points higher than the associated first mortgage.³⁶

Since $\Delta\hat{m}$ is, on average, the loan amount in excess of the conforming limit at which borrowers are indifferent between taking out a jumbo loan and reducing their first mortgage balance to the limit, the similarity of the implied $r^*(m)$ and the estimated conforming/second

³³This calculation is analogous to the “reduced-form approximation” to the tax elasticity proposed by [Kleven and Waseem \(2013\)](#).

³⁴For a loan that is $\Delta\hat{m}$ above the limit in approximate percentage terms, equation (21) simplifies to $r^*(\bar{m} + \Delta\hat{m}) = \hat{r} + \Delta\hat{r} + \frac{\Delta\hat{r}}{\Delta\hat{m}}$.

³⁵As we discuss further in the next section, taking out a second mortgage is only one way to reduce a first mortgage balance. Here we consider the jumbo and conforming/second options in isolation for illustrative purposes.

³⁶Since LPS is a loan-level data set, there is no direct link between first and second mortgages on the same home purchase. We construct a sample of matched loans using the property type, zip code, closing date, appraisal amount, and FICO score of the borrower, which together almost uniquely identify loans. Because we prefer to be very conservative with this calculation, the match is quite strict: We match only about 13 percent of first loans from California in 2006 to an associated second loan. HMDA data indicate that the national figure for 2006 was about 30 percent, and the DQ data suggest that the figure for California was more than half.

spread lends substantial credence to our bunching estimate. However, they need not be exactly equal since taking out a second mortgage may entail additional unobserved costs that make the jumbo loan slightly more attractive. Indeed, there are some jumbo loans made only slightly above the limit, which can only be justified by having extremely low elasticities or facing substantial adjustment costs.

Given estimates of $P^*(m)$ and $r^*(m)$, we can calculate the elasticities of (first) mortgage demand implied by our estimate of $\Delta\hat{m}$. Let $P'(\bar{m}) = \frac{P(\bar{m}, \hat{r})}{\bar{m}}$ be the (constant) marginal payment per dollar of the loan for loans made at the conforming interest rate \hat{r} . Then we can define

$$\epsilon_s^r = \frac{\Delta\hat{m}}{r^*(\bar{m} + \Delta\hat{m}) - \hat{r}} \quad (22)$$

and

$$\epsilon^P = \frac{\Delta\hat{m}}{\log(P^*(\bar{m} + \Delta\hat{m})) - \log(P'(\bar{m}))} \quad (23)$$

as the semi-elasticity of mortgage demand with respect to the interest rate and the elasticity of mortgage demand with respect to the increase in payment, respectively. As before, $\Delta\hat{m}$ is estimated in logs, so it represents the approximate percentage change in mortgage demand induced by the conforming limit. The denominators in both expressions represent the change in the effective marginal costs to borrowers induced by the limit.³⁷ 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 different loans or to compare owning to the monthly cost of renting (Attanasio et al., 2008).³⁸

7.2 Estimated Elasticities of Mortgage Demand

The first two columns of table 7 report the semi-elasticities we calculate for a range of estimates of $\Delta\hat{m}$ 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 $\Delta\hat{m}$ at the beginning of that row.

Our preferred estimate of bunching for FRMs from table 3 (0.063) is shown in the middle row. The other two estimates (0.052 and 0.083) are the smallest and largest estimates of $\Delta\hat{m}$

³⁷We 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.

³⁸The 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

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.³⁹ 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 5 and correspond to the OLS (column 1) and IV (column 2) estimates, respectively.

The estimated semi-elasticities range from about -0.015 to about -0.053, with our preferred estimates in the middle row at -0.022 and -0.31. 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, which strikes us as a reasonably small but plausible estimate.

Columns 3 and 4 report the analogous elasticities with respect to the marginal monthly payment for the same set of estimates of $\Delta\hat{m}$ and equivalent estimates of the jumbo-conforming spread in logs, taken from table 6. The preferred estimates in the middle row are -0.27 using the OLS estimate of 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.

8 How Do Borrowers Adjust?

Although the elasticities reported in table 7 are valid estimates of the amount by which borrowers reduce their first mortgage balance in response to the jumbo-conforming spread, they do not provide any information regarding the margins along which this adjustment occurs. In this section we present suggestive evidence that helps to distinguish between several methods that borrowers may be using to adjust their mortgage balance.

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.⁴⁰ Second, she could take out an additional mortgage

³⁹We 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.

⁴⁰“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

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.⁴¹ While borrowers likely use various combinations of these three strategies to lower their mortgage balances, if we consider the extreme cases in which only one of the three is used, we can provide a rough sense of the extent to which each may be contributing to the bunching behavior we observe.

To measure the extent to which borrowers are using second mortgages to lower their first mortgage balance, figure 13 plots the number of transactions financed using second loans as a function of the associated first mortgage value relative to the conforming limit for all fixed-rate first mortgages in the DQ sample. 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. Returning to table 3, we estimated that about 70,000 FRM borrowers bunched at the limit, suggesting 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. 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.⁴² 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 14, 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 9.⁴³

allow us to distinguish between these cases.

⁴¹This could happen either through substitution to a lower quality home or through direct price negotiation with the seller.

⁴²Some 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.

⁴³The 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

Thus, a significant portion of the 65 percent of borrowers who bunch 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 15 plots the first-mortgage LTV against the first mortgage amount. This figure is analogous to figure 14 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.⁴⁴ 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.

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, we think that another paper, Adelino et al. (2012), is likely to be 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 a difference-in-difference approach rather than the bunching approach

have different characteristics 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 9, there are relatively few borrowers remaining in these bins.

⁴⁴To 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 9. 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.

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).

Regardless of how they do so, 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. Unfortunately, we only know the average reduction in first loan size ($\Delta\hat{m}$) among marginal bunching borrowers, so we cannot solve for the response of total mortgage debt. If, however, the fraction of marginal bunching individuals who take out a second loan is the same as the fraction who do so in the population of all bunchers, then scaling down our elasticity estimates by a factor of one-third should 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.

9 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 provide some basic calculations using our estimated mortgage demand elasticities to illustrate how these 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. 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). 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.⁴⁵ 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).⁴⁶ 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).⁴⁷

Our estimates of the interest rate elasticity of mortgage demand can be used to gauge the potential 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. While our elasticity is well suited to examining the direct effects of these relatively small changes in fees, we emphasize that we cannot draw any inference on

⁴⁵The 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.

⁴⁶FHFA reports this 11 basis point number as the combined overall average 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.

⁴⁷The future of this plan remains uncertain, particularly after it was delayed by incoming director Mel Watt in January 2014.

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.

10 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 rates.

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 in response to a 1 percentage point increase in the first mortgage interest rate. The remaining two-thirds of bunchers appear to be bringing more cash to the table rather than spending less on housing. Applying these elasticity estimates to recently proposed increases in GSE guarantee fees implies a reduction in fixed-rate 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 two 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, the differential responses of minority versus non-minority and high-income versus low-income borrowers suggests that there could be heterogeneity in elasticities or adjustment costs along a wide range of characteristics. Understanding this heterogeneity may be as important as pinning down the overall average elasticity of demand.

References

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino**, “Credit Supply and House Prices: Evidence From Mortgage Market Segmentation,” February 2012. NBER Working Paper No. 17832.
- 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.
- 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.
- Best, Michael and Henrik J. Kleven**, “Property Transaction Taxes and the Housing Market: Evidence from Notches and Housing Stimulus in the UK,” September 2013. Working Paper.
- Browning, Martin and Annamaria Lusardi**, “Household Saving: Micro Theories and Micro Facts,” *Journal of Economic Literature*, 1996, 34 (4), 1797–1855.
- Brueckner, Jan K.**, “The Demand for Mortgage Debt: Some Basic Results,” *Journal of Housing Economics*, 1994, 3, 251–262.
- 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).
- Congressional Budget Office**, “Federal Subsidies and the Housing GSEs,” 2001.
- Federal Housing Finance Agency**, “A Strategic Plan for Enterprise Conservatorships: The Next Chapter in a Story that Needs and Ending,” February 2012.
- , “FHFA Takes Further Steps to Advance Conservatorship Strategic Plan by Announcing and 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 Gaurantee Fees,” July 2013.
- 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.
- Fuster, Andreas and James Vickery**, “Securitization and the Fixed-Rate Mortgage,” January 2013. Federal Reserve Bank of New York Staff Report No. 594.
- Gary-Bobo, Robert J. and Sophie Larribeau**, “A Structural Econometric Model of Price Discrimination in the French Lending Industry,” *International Journal of Industrial Organization*, 2004, 22, 101–134.
- Gelber, Alexander M., Damon Jones, and Daniel W. Sacks**, “Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test,” September 2013. Working Paper.
- 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.
- 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).
- 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 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.
- 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.
- 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 Residential Real Estate Market,” August 2013. Working Paper.
- 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.
- Mishkin, Frederic S.**, “Symposium on the Monetary Transmission Mechanism,” *Journal of Economic Perspectives*, 1995, 9 (4), 3–10.
- 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.
- Poterba, James**, “Tax Subsidies 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.
- Saez, Emmanuel**, “Do Taxpayers Bunch at Kink Points,” *American Economic Journal: Economic Policy*, 2010, 2, 180–212.

Sallee, James M. and Joel Slemrod, “Car Notches: Strategic Automaker Responses to Fuel Economy Policy,” December 2010. NBER Working Paper No. 16604.

Sherlund, Shane M., “The Jumbo-Conforming Spread: A Semiparametric Approach,” 2008. Federal Reserve Board Finance and Economics Discussion Series.




FIG. 1.—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).




FIG. 2.—Mean Interest Rate Relative to the Conforming Limit, Fixed-Rate Mortgages Only (2006). This figure plots the mean interest rate for fixed rate mortgages originated in 2006 as a function of the loan amount relative to the conforming limit. 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. See text for details on sample construction.




FIG. 3.—Loan Size Distribution Relative to the Conforming Limit. This figure 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 each loan is centered at the conforming limit in effect at the date of origination, so that a value of 0 represents a loan at exactly the conforming limit. Sample includes all transactions in the primary DataQuick sample that fall within \$400,000 of the conforming limit. See text for details on sample construction.

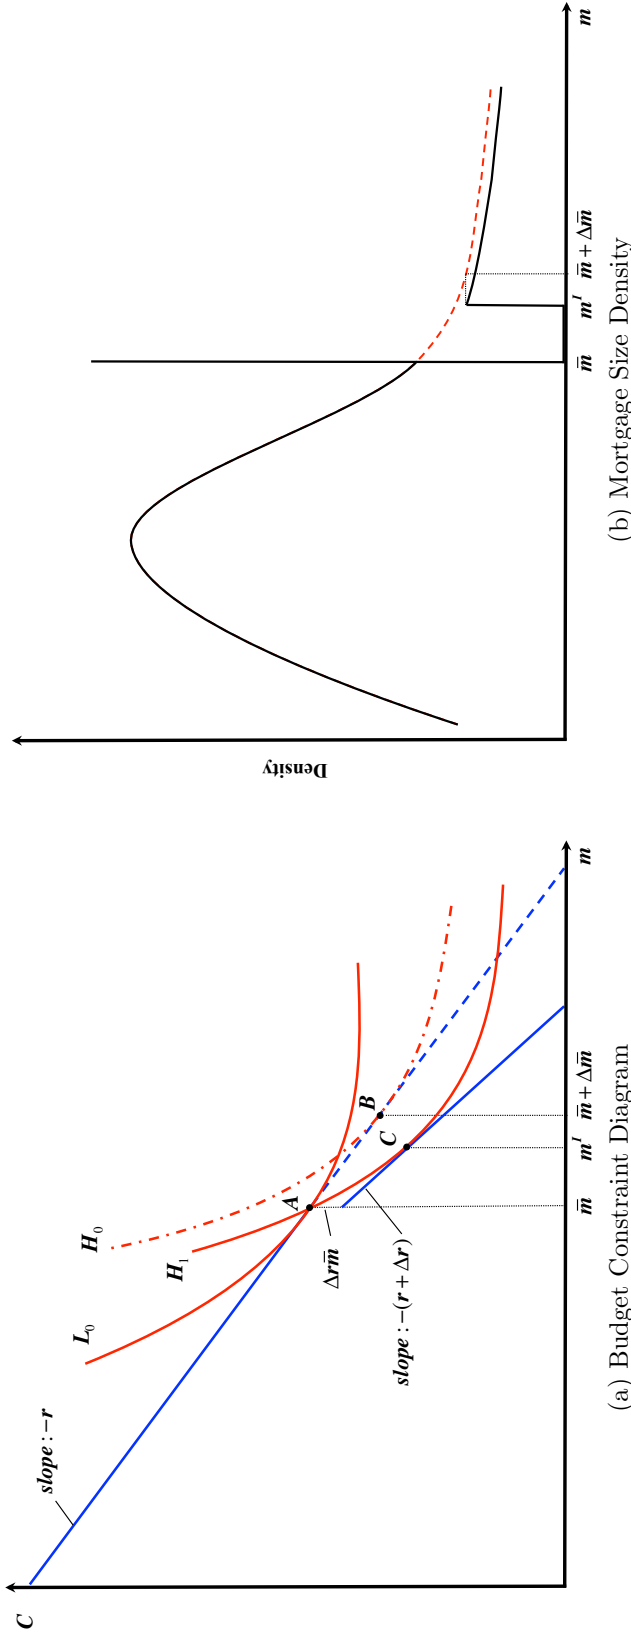


FIG. 4.—Behavioral Response to Conforming Loan Limit. This figure shows the effect of the conforming loan limit (CLL) on individual behavior (Panel a) and the aggregate loan size distribution (Panel b). The CLL generates a notch in the intertemporal budget constraint, which is characterized by a jump 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 distribution 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').




FIG. 5.—Estimating Bunching and the Counterfactual Distribution. This figure provides a graphical illustration of the bunching estimation procedure. The solid black line represents the empirical loan size distribution, where loan amounts are centered at the conforming limit. The heavy dashed red line represents the estimated counterfactual distribution which is obtained by fitting a polynomial to the empirical distribution, omitting the contribution of the loans in the region between m_L and m_U . Bunching (\hat{B}) is estimated as the difference between the observed and counterfactual distributions in the excluded region at and to the left of the conforming limit. The upper limit of the excluded region is chosen to minimize the difference between the amount of missing mass in the excluded region to the right of the limit (\hat{M}) and the amount of bunching.




FIG. 6.—Bunching at the Conforming Limit, All Loans. This figure plots the empirical and counterfactual distribution 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 distribution. 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 distribution 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 5.1.




FIG. 7.—Bunching Adjusting Over Time. This figure plots the empirical distribution 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. 8.—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.




FIG. 9.—Bunching at the Conforming Limit, Fixed-Rate Mortgages Only. This figure plots the empirical and counterfactual distribution of (log) loan size relative to the conforming limit for fixed-rate mortgages only. Estimation was carried out in the full sample of DQ loans with fixed interest rates, but the figure shows only loans within 50 percent of the conforming limit. The connected black line is the empirical distribution. 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 distribution 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 5.1.







FIG. 10.—Bunching at the Conforming Limit, Adjustable-Rate Mortgages Only. This figure plots the empirical and counterfactual distribution of (log) loan size relative to the conforming limit for adjustable-rate mortgages only. Estimation was carried out in the full sample of DQ loans with adjustable interest rates, but the figure shows only loans within 50 percent of the conforming limit. The connected black line is the empirical distribution. 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 distribution 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 5.1.




(a) Non-Minority



(b) Minority




(c) High-Income




(d) Low-Income


FIG. 11.—Bunching at the Conforming Loan Limit by Borrower Type, Fixed Rate Mortgages Only. This figure plots the empirical and counterfactual distribution of (log) loan size relative to the conforming limit estimated separately for: (a) high-income borrowers, (b) low-income borrowers, (c) minority borrowers and (d) non-minority borrowers. The connected black line is the empirical distribution. 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 distribution 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 5.1.




(a) Interest Rate



(b) Monthly Payment



(c) Average Marginal Interest Rate Per Dollar (r^*)



(d) Average Marginal Payment Per Dollar (P^*)

FIG. 12.—Effective Marginal Interest Rate and Mortgage Payment Example. This figure shows how the effective marginal interest rate and marginal mortgage payment per dollar of the loan varies as a function of the distance between the loan amount and the conforming limit. The vertical dashed blue line in each panel is the conforming limit in 2006 and the solid vertical blue line denotes our estimate of the average behavioral response ($\Delta \hat{m}$). The conforming interest rate is calculated based on the average rate for conforming loans just below the limit in 2006. The jumbo-conforming spread is taken from our pooled OLS estimate in column 3 of table 5.

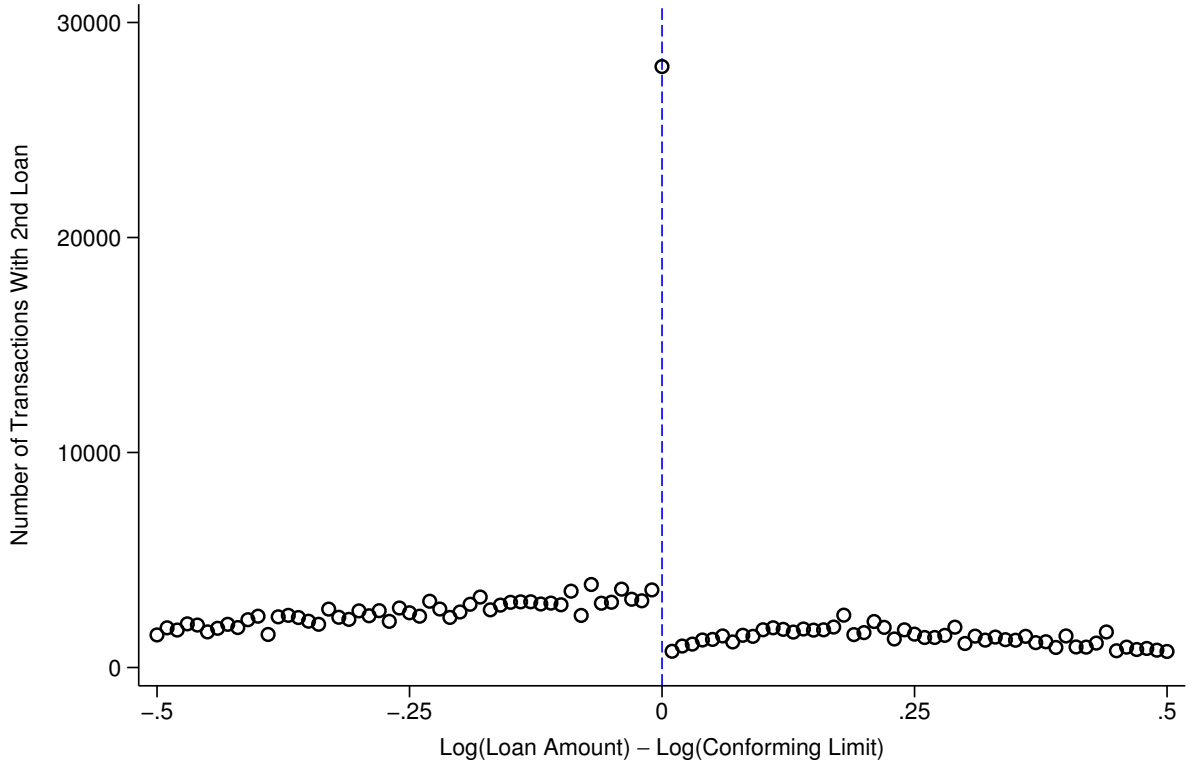


FIG. 13.—Number of Second Mortgages by First Mortgage Amount. This figure plots the number of transactions financed using two loans as a function of the first loan amount relative to the conforming limit. Each dot represents the number of transactions in a given 1-percent bin relative to the limit in effect at the time of origination. Sample includes only transactions with a fixed-rate first mortgage.




FIG. 14.—Combined Loan-to-Value Ratio by First Mortgage Amount. This figure plots the average combined loan-to-value ratio (CLTV) as a function 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. 15.—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 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.

TABLE 1
SUMMARY STATISTICS, DATAQUICK SAMPLE

	Full Sample		Within \$50k of CLL	
	(1) Mean	(2) SD	(3) Mean	(4) SD
<i>Transaction and Loan Characteristics</i>				
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)
<i>Housing Characteristics</i>				
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)
<i>Borrower Characteristics</i>				
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 2
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 (0.91)	6.70 (0.90)	5.05 (2.13)	5.04 (2.09)
Jumbo	0.17 (0.38)	0.14 (0.35)	0.50 (0.50)	0.39 (0.49)
First Loan Amount (\$1,000)	278.03 (174.40)	322.55 (56.81)	452.13 (283.42)	376.95 (47.29)
Appraisal Amount (\$1,000)	396.50 (281.31)	447.16 (156.74)	611.28 (438.51)	497.01 (105.31)
Loan-to-Value Ratio	74.09 (16.35)	76.09 (13.63)	76.31 (10.07)	77.33 (9.17)
Debt-to-Income Ratio	35.62 (12.47)	36.79 (11.96)	35.51 (11.96)	36.50 (11.46)
Missing DTI Ratio	0.70 (0.46)	0.75 (0.44)	0.45 (0.50)	0.45 (0.50)
FICO Score	731.18 (51.81)	731.70 (49.98)	719.70 (52.82)	717.97 (52.52)
Missing FICO Score	0.32 (0.47)	0.35 (0.48)	0.20 (0.40)	0.22 (0.41)
Term (Months)	345.91 (52.58)	350.14 (46.44)	365.78 (30.73)	365.63 (28.83)
30-Year	0.90 (0.30)	0.93 (0.26)	0.93 (0.25)	0.94 (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 3
BUNCHING ESTIMATES BY LOAN TYPE

	(1) Combined	(2) FRMs	(3) ARMs
Bunched Loans (\hat{B})	84183.8 (2687.0)	68306.3 (1561.6)	16174.3 (1446.7)
Behavioral Response ($\Delta\hat{m}$)	0.0378 (0.0018)	0.0627 (0.0025)	0.0144 (0.0014)
Excess Mass ($\hat{B}/\sum_{j=L}^0 \hat{n}_j$)	3.781 (0.175)	6.266 (0.253)	1.436 (0.141)
Upper Limit (m_H)	0.160 (0.028)	0.160 (0.021)	0.080 (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 5.1.

TABLE 4
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 5.1.

TABLE 5
JUMBO-CONFORMING SPREAD ESTIMATES, PERCENTAGE POINTS

	(1)	(2)	(3)	(4)
	Baseline	Splines	Within \$50k of CLL	Within \$10k of CLL
<i>Fixed-Rate Mortgages</i>				
OLS	0.179 (0.002)	0.182 (0.002)	0.171 (0.014)	0.172 (0.038)
IV	0.107 (0.010)	0.121 (0.009)	0.118 (0.028)	0.126 (0.056)
Observations	1,061,738	1,061,738	263,641	87,617
<i>Adjustable-Rate Mortgages</i>				
OLS	-0.076 (0.009)	-0.090 (0.009)	-0.299 (0.037)	-0.362 (0.083)
IV	-0.001 (0.020)	0.004 (0.019)	0.074 (0.054)	0.040 (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 the “appraisal limit” used as an instrument for the CLL, as described in the text. Controls for 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 splines in LTV and DTI ratios, as well as FICO score. Columns 3 and 4 limit the sample to loans near the CLL.

TABLE 6
JUMBO-CONFORMING SPREAD ESTIMATES, LOG POINTS

	(1)	(2)	(3)	(4)
	Baseline	Splines	Within \$50k of CLL	Within \$10k of CLL
<i>Fixed-Rate Mortgages</i>				
OLS	0.026 (0.000)	0.027 (0.000)	0.023 (0.002)	0.025 (0.006)
IV	0.017 (0.002)	0.018 (0.001)	0.017 (0.005)	0.017 (0.010)
Observations	1,061,738	1,061,738	263,641	87,617
<i>Adjustable-Rate Mortgages</i>				
OLS	-0.051 (0.003)	-0.053 (0.003)	-0.168 (0.015)	-0.210 (0.034)
IV	0.005 (0.008)	0.009 (0.008)	0.014 (0.022)	0.002 (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 the “appraisal limit” used as an instrument for the CLL, as described in the text. Controls for 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 splines in LTV and DTI ratios, as well as FICO score. Columns 3 and 4 limit the sample to loans near the CLL.

TABLE 7
 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)	
	$\Delta \hat{r}$		$\Delta \log \hat{r}$	
$\Delta \hat{m}$	0.171 (0.014)	0.118 (0.028)	0.023 (0.002)	0.017 (0.005)
0.052 (0.003)	-0.015 (0.002)	-0.022 (0.005)	-0.188 (0.022)	-0.246 (0.064)
0.063 (0.003)	-0.022 (0.002)	-0.031 (0.008)	-0.267 (0.028)	-0.351 (0.091)
0.083 (0.004)	-0.037 (0.004)	-0.053 (0.013)	-0.445 (0.051)	-0.589 (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 ($\Delta \hat{m}$), 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 5.1. Standard errors for the elasticities were calculated using the delta method.