

The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit[†]

By ANTHONY A. DEFUSCO AND ANDREW PACIOREK*

This paper provides novel estimates of the interest rate elasticity of mortgage demand by measuring the degree of bunching in response to a discrete jump in interest rates at the conforming loan limit—the maximum loan size eligible for purchase by Fannie Mae and Freddie Mac. The estimates indicate that a 1 percentage point increase in the rate on a 30-year fixed-rate mortgage reduces first mortgage demand by between 2 and 3 percent. One-third of this response is driven by borrowers who take out second mortgages, which implies that total mortgage debt only declines by 1.5 to 2 percent. (JEL D14, G21, R21, R31)

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 a large majority of total household debt, the elasticity plays a significant role in governing the degree to which monetary policy affects aggregate borrowing behavior (Mishkin 1995, Jordà, Schularick, and Taylor 2014). In public finance, the elasticity is important for understanding the effect of the home mortgage interest

*DeFusco: Kellogg School of Management, Northwestern University, 2001 Sheridan Road, 4215 Donald P. Jacobs Center, Evanston, IL 60208 (e-mail: anthony.defusco@kellogg.northwestern.edu), Paciorek: Board of Governors of the Federal Reserve System, Mailstop 80, 20th Street NW and Constitution Avenue NW, Washington, DC 20551 (e-mail: andrew.d.paciorek@frb.gov). We are grateful to Manuel Adelino, Neil Bhutta, Gilles Duranton, Fernando Ferreira, Joseph Gyourko, Raven Molloy, Christopher Palmer, Dan Sacks, Stephen Shore, Todd Sinai, James Vickery, Yiwei Zhang, two anonymous referees, and seminar participants at the Board of Governors, the 2014 AREUEA-ASSA conference, the 2014 spring HULM conference, the 2014 NBER Summer Institute workshop on Real Estate, and the 2014 UEA annual conference for helpful comments and discussion. 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.

[†]Go to <https://doi.org/10.1257/pol.20140108> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

deduction on both government tax revenue and consumer leverage (Poterba 1984, Poterba and Sinai 2008, Poterba and Sinai 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, Schoar, and Severino 2012; Kaufman 2012). This final consideration has become particularly salient recently in light of the ongoing debate over the future of the GSEs in the wake of the 2007–2008 financial crisis.

Despite these potentially important policy implications, there are essentially no existing causal estimates of the extent to which individual loan sizes respond to changes in interest rates.¹ This is due in large part to data limitations and a lack of plausibly exogenous variation in interest rates. 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. (See, for example, Zinman’s 2015 review, as well as Gross and Souleles 2002; Alessie, Hochguertel, and Weber 2005; Attanasio, Goldberg, and Kyriazidou 2008; Karlan and Zinman 2008; and Karlan and Zinman 2014.) In the spirit of these studies, we estimate the interest rate elasticity of mortgage demand using microdata on over 2.7 million mortgages and an identification strategy that leverages “bunching” at nonlinearities in household budget constraints.

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

While identifying the precise magnitude of this interest rate differential is challenging because borrowers may sort around the limit, some insight can still be gleaned from examining the raw data.² For example, Figure 1, panel A plots the interest rate as a function of the difference between the loan amount and the conforming limit for all fixed-rate mortgages in our analysis sample that were originated in 2006. There is a clear discontinuity precisely at the limit, with interest rates

¹One exception is in recent work by Fuster and Zafar (2015), who survey households about their mortgage choices under randomized hypothetical interest rate scenarios and find results similar to what we report below. In other related work, Martins and Villanueva (2006) estimate how the extensive margin probability of obtaining a mortgage responds to an interest rate subsidy for low-income households in Portugal but do not report direct estimates of the effects on loan size. Similarly, several others including Follain and Dunskey (1997), Ling and McGill (1998), Dunskey and Follain (2000), and Jappelli and Pistaferri (2007) have estimated how mortgage debt responds to changes in the rate at which interest expenses can be deducted from personal income but do not focus explicitly on effective interest rates themselves.

²Many papers have attempted to overcome this challenge using a variety of empirical methods. See, for example, Hendershott and Shilling (1989); Passmore, Sparks, and Ingpen (2002); Passmore, Sherlund, and Burgess (2005); Sherlund (2008); and Kaufman (2012). We address these issues in detail below in Section III.

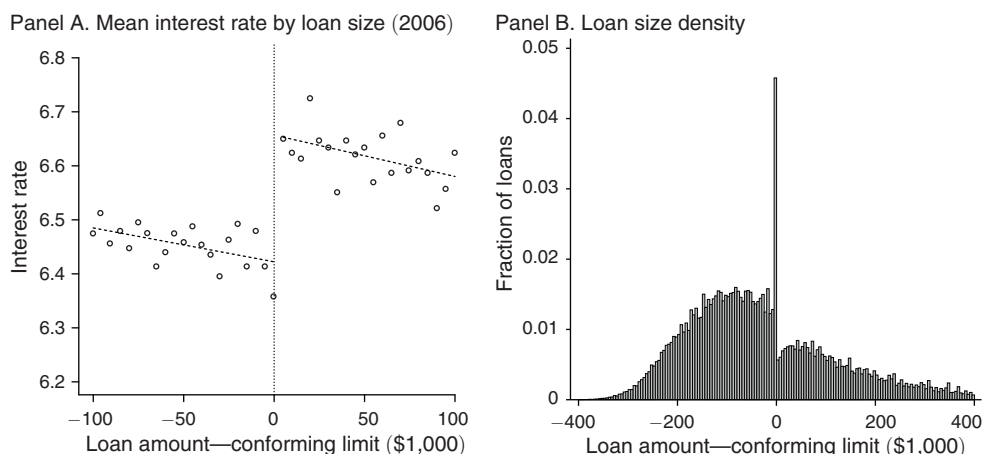


FIGURE 1

Notes: Panel A plots the mean interest rate as a function of the loan amount relative to the conforming limit for fixed-rate mortgages originated in 2006. Each dot represents the mean interest rate within a given \$5,000 bin relative to the limit. The dashed lines are predicted values from a regression fit to the binned data, allowing for changes in the slope and intercept at the conforming limit. Sample includes all loans in the LPS fixed-rate sample that fall within \$100,000 of the conforming limit. Panel B plots the fraction of all loans that are in any given \$5,000 bin relative to the conforming limit. Data are pooled across years and the sample includes all transactions in the primary DataQuick sample that fall within \$400,000 of the conforming limit. In both panels, each loan is centered at the conforming limit in effect at the date of origination so that a value of zero represents a loan at exactly the conforming limit.

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

The difference in interest rates between jumbo and conforming loans creates a substantial “notch” in the intertemporal budget constraint of households deciding how much mortgage debt to incur. This notch induces some borrowers who would otherwise take out loans above the conforming limit to instead bunch right at the limit. This behavior is confirmed by Figure 1, panel B, which shows the fraction of all loans in our analysis sample that fall into any given \$5,000 bin relative to the conforming limit in effect at the date of origination.³ 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 above the limit.

The intuition behind our empirical strategy is to combine reasonable estimates of the jumbo-conforming spread with a measure of the excess mass of individuals who bunch at the conforming limit to back out estimates of the interest rate (semi-) elasticity of demand for mortgage debt. Recent papers in public finance have developed methods for estimating behavioral responses to nonlinear incentives in similar settings (Saez 2010, Chetty et al. 2011, Kleven and Waseem 2013).⁴ We adapt

³During our sample period the conforming limit varied from around \$215,000 in 1997 to approximately \$420,000 in 2006 and 2007. Online Appendix Figure A.1 shows the full path of the limit in real and nominal terms during this period.

⁴Other recent applications of these and similar methods include Manoli and Weber (2011), Sallee and Slemrod (2012), Chetty, Friedman, and Saez (2013); Gelber, Jones, and Sacks (2015); Kopczuk and Munroe (2015); and

these methods to the case of mortgage choice in the face of a notched interest rate schedule. To the best of our knowledge, ours is the first application of these methods to the mortgage market or to a consumer credit market of any kind.⁵

Our preferred specifications indicate that the average size of a borrower's first (fixed-rate) mortgage declines by between 2 and 3 percent for each 1 percentage point rise in the first mortgage rate, holding constant all prices and interest rates that do not change at the conforming limit. Because both the bunching estimates and the jumbo-conforming spread estimates vary depending on the assumptions used in estimation, we also provide alternative estimates under a range of different scenarios. These estimates imply a decline of between 1.5 and 5 percent for a 1 percentage point increase in the mortgage rate.

While the mortgage demand elasticity is of innate interest, its interpretation depends in part on the channels through which borrowers adjust their first mortgage balance. Our second main contribution is to provide 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. As we discuss below, these estimates appear to be quite small and are in line with similar findings of small behavioral responses to interest rates for other types of household debt as well as survey evidence on hypothetical mortgage choice (Attanasio, Goldberg, and Kyriazidou 2008; Karlan and Zinman 2008; Fuster and Zafar 2015).

To further gauge the economic magnitude of the effects we estimate, we apply them to recently proposed increases in 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, while large by historical standards, would only reduce the total dollar volume of fixed-rate conforming mortgage originations by approximately one-fifth of 1 percent. When we apply our elasticity to similar increases in fees that have occurred in recent years, we estimate an effect on the order of one-half of 1 percent.

The remainder of the paper is organized as follows. Section I presents our conceptual framework. In Sections II and III, we discuss our data and empirical research design. We then present our main results in Sections IV–V. Section VI applies these results to changes in the GSE guarantee fees, and Section VII concludes by discussing avenues for future research.

Best and Kleven (2016).

⁵In work subsequent to ours, Best et al. (2015) have applied this approach in the context of the UK mortgage market to provide structural estimates of the intertemporal elasticity of substitution.

I. Theoretical Framework

We begin by considering a simple two-period model of mortgage choice similar to that in Brueckner (1994). 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, Fiedman, and Saez (2013); and Gelber, Jones, and Sacks (2015) study labor supply and earnings responses to kinked income tax and social security benefit schedules. Similar models have also been developed to study behavioral responses in applications more analogous to ours, where the budget constraint features a notch as opposed to a kink. Our analysis draws heavily on the framework developed by Kleven and Waseem (2013) and Best and Kleven (2016), who study behavioral responses to notched income and real estate transfer tax schedules, respectively. Kopczuk and Munroe (2015) have also used this framework to study real estate transfer taxes and others have used it to study fuel economy regulation (Sallee and Slemrod 2012) and retirement incentives (Manoli and Weber 2011).

A. Baseline Case: Linear Interest Rate Schedule

Households live for two periods. Since our primary focus is on the intensive margin choice of how much debt to incur conditional on purchasing a home, we shut down housing choice by assuming that each household must purchase one unit of housing services in the first period at an exogenous per unit price of p .⁶ Households can finance their housing purchase with a mortgage, m , that 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.

Since we impose the exogenous requirement that households consume one unit of housing services, we can suppress the argument for housing consumption. The household's problem is thus to maximize lifetime utility by choosing nonhousing consumption in each period, denoted by c_1 and c_2 . In general, the household solves:

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

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

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

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

⁶In the online Appendix, we relax the assumption that households cannot choose the quantity of housing services to consume. While the intuition is the same, extending the model in this way prevents us from deriving a closed-form solution, which makes it less useful for motivating the empirical work.

where $\delta \in (0, 1)$ is the discount factor and y is first period income.

To proceed, we make several simplifying assumptions. First, we assume that household preferences are given by the constant elasticity function $u(c) = \frac{1}{1-\xi} c^{1-\xi}$. Second, heterogeneity in the model is driven by the discount factor, which is assumed to be distributed smoothly in the population according to the density function $f(\delta)$. For illustrative purposes, we assume that y , p , and ξ are constant across households; however, this assumption is not crucial, and we discuss below how relaxing it affects the interpretation of our results. 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:

$$(5) \quad m^* = \frac{p - (\delta(1+r))^{1/\xi} (y-p)}{(\delta(1+r))^{1/\xi} + (1+r)}.$$

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

B. 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 ineligible to be purchased by the GSEs and are therefore subject to a higher interest rate for reasons we discuss in detail in the online Appendix. This leads to the new interest rate schedule $r(m) = r + \Delta r \cdot \mathbf{1}(m > \bar{m})$. Here, Δr is the difference in interest rates between jumbo and conforming loans and $\mathbf{1}(m > \bar{m})$ is an indicator for jumbo loan status. Combining equations (2) and (3) yields the lifetime budget constraint

$$(6) \quad C = y - m \cdot [r + \Delta r \cdot \mathbf{1}(m > \bar{m})],$$

where $C = c_1 + c_2$ is lifetime consumption. This budget constraint is plotted in Figure 2, panel A, along with indifference curves for two representative households.

The notch in the budget constraint induces some households to bunch at the conforming loan limit. In Figure 2, panel A, household L is the household with the lowest baseline mortgage amount—the largest δ —that 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 $\bar{m} + \Delta \bar{m}$.

Panel A. Budget constraint diagram



Panel B. Mortgage size density




FIGURE 2. BEHAVIORAL RESPONSE TO CONFORMING LOAN LIMIT

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

With the notch, however, the household is indifferent between locating at \bar{m} and the best interior point beyond the conforming limit, m' . 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' in the notch scenario.

The density when a notch exists, $g_1(m)$, will therefore be characterized by both a mass of households locating precisely at the conforming limit as well as a missing mass of households immediately to the right of the limit. The effect of the notch on the mortgage size distribution is shown in the density diagram in Figure 2, panel B. The solid line shows the density of loan amounts in the presence of the notch and the dashed line to the right of the notch shows the counterfactual density that would exist in the absence of the conforming loan limit.

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

$$(7) \quad B = \int_{\bar{m}}^{\bar{m} + \Delta\bar{m}} g_0(m) dm \approx g_0(\bar{m})\Delta\bar{m},$$

where the approximation assumes that the counterfactual no-notch distribution is constant on the bunching interval $(\bar{m}, \bar{m} + \Delta\bar{m}]$. The expression in equation (7) is the primary motivation for our empirical strategy. In essence, it tells us the width of the interval between the conforming limit and the point in the loan size distribution from which the marginal bunching borrower would have come under the counterfactual.

More formally, 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.⁷

C. Heterogeneous Intertemporal Elasticities, Incomes, and Prices

The derivation of equation (7) rests upon the assumption that ξ , y , and p are constant across households. In that case, it is possible to back out the exact change in mortgage amount for the marginal bunching individual. When intertemporal elasticities, incomes, and prices are allowed to vary across households, the amount of bunching instead identifies the *average* response among the marginal bunching individuals associated with each intertemporal elasticity, income, and price level. To see this, let the joint density of discount factors, intertemporal elasticities, incomes, and prices be given by $\bar{f}(\delta, \xi, y, p)$, where $y \in (0, \bar{y}]$, $\xi \in (0, \bar{\xi}]$, and $p \in (0, \bar{p}]$ for some upper bounds, \bar{y} , $\bar{\xi}$, and \bar{p} . For a fixed (ξ, y, p) triple, the bunching interval is determined in exactly the same way as in the baseline model. Denote this interval $(\bar{m}, \bar{m} + \Delta \bar{m}_{\xi, y, p})$, where $\Delta \bar{m}_{\xi, y, p}$ is the behavioral response of the marginal bunching individual among those with intertemporal elasticity $1/\xi$, income y , and who face price p .⁸ Further, let $\bar{g}_0(m, \xi, y, p)$ denote the joint density of mortgage sizes, intertemporal elasticities, incomes, and prices in the pre-notch scenario and $g_0(m) \equiv \int_{\xi} \int_y \int_p \bar{g}_0(m, \xi, y, p) dp dy d\xi$ the unconditional mortgage size density. The amount of bunching can then be expressed as

$$(8) \quad B = \int_{\xi} \int_y \int_p \int_{\bar{m}}^{\bar{m} + \Delta \bar{m}_{\xi, y, p}} \bar{g}_0(m, \xi, y, p) dm dp dy d\xi \approx g_0(\bar{m}) E[\Delta \bar{m}_{\xi, y, p}].$$

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.⁹

II. Data

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

⁷ Much of the structure in the model above is not needed for this result to hold. For example, the assumption of a locally constant no-notch distribution merely simplifies the discussion. In the empirical application we allow for curvature in the counterfactual distribution. 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.

⁸ For simplicity, we assume here that p is exogenous from the borrower's perspective. In the online Appendix, we present suggestive evidence that borrowers do not buy cheaper houses because of the presence of the limit.

⁹ Kleven and Waseem (2013) show a directly analogous result in the context of earnings responses to notched income tax schedules.

These data serve as our primary source of information on loan size. The second data source consists of loan-level records collected by Lender Processing Services (LPS) and contains extensive information on interest rates, borrower characteristics, and loan terms, which we use to estimate the jumbo-conforming spread.¹⁰

A. DataQuick

Each record in the DQ dataset 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 also drop transactions where the initial loan-to-value ratio or transaction price was a clear outlier. We use data from California because that is where the information from DataQuick is most reliable, particularly for identifying when multiple loans were used to finance a purchase. In addition, because average house prices in California are higher than in other states, we expect that the differences between the typical transaction and one financed with a loan near the conforming limit will be less stark in California than in other parts of the country.

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

The final reason we avoid using post-2007 data is that there were significant changes to the structure of the mortgage market during the financial crisis that could potentially confound our analysis. For example, the jumbo securitization market almost completely dried up during this period, which led to a sharp reduction in the number of jumbo loans originated and a large rise in the jumbo-conforming spread (Fuster and Vickery 2015). 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 2015). These restrictions leave us with a primary estimation sample of approximately 2.7 million transactions across 26 MSAs. Summary statistics for this sample and the subsample with first loan amounts within \$50,000 of the conforming limit are reported in online Appendix Table A.1.

¹⁰In 2014, Lender Processing Services became Black Knight and DataQuick was acquired by CoreLogic, another large property information service provider. We refer to the datasets by the prior names and initialisms.

B. LPS

The primary disadvantage of the DQ dataset 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.¹¹ 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 MSAs in California and the first nine months of each year between 1997 and 2007. Descriptive statistics for the full LPS sample and the subsample of loans with first loan amounts within \$50,000 of the conforming limit are reported in online Appendix Table A.2.

III. Empirical Methodology

A. Estimating the Behavioral Response to the Conforming Limit

In Section I, 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 dataset 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 $j = -J, \dots, L, \dots, 0, \dots, U, \dots, J$, and count the number of loans in each bin, n_j . To obtain estimates of bunching and the counterfactual loan size distribution, we define an excluded region around the conforming limit, $[m_L, m_U]$, such that $m_L < 0 < m_U$ and fit the following regression to the count of loans in each bin

$$(9) \quad n_j = \sum_{i=0}^p \beta_i (m_j)^i + \sum_{k=L}^U \gamma_k \mathbf{1}(m_k = m_j) + \epsilon_j.$$

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

¹¹ Although data are available from earlier years, they are less comprehensive and the loans have higher average “seasoning,” meaning that it takes longer after origination for them to appear in the dataset (Fuster and Vickery 2015). If loans that are quickly prepaid or foreclosed on never appear, seasoned data may be less representative of the universe of loans.

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

$$(10) \quad \hat{n}_j = \sum_{i=0}^p \hat{\beta}_i(m_j)^i.$$

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,

$$(11) \quad \hat{B} = \sum_{j=L}^0 (n_j - \hat{n}_j) = \sum_{j=L}^0 \hat{\gamma}_j,$$

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

The parameter of primary interest is $\Delta \hat{m}$, the empirical analogue of $\Delta \bar{m}$ from equation (7). 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

$$(12) \quad \Delta \hat{m} = \frac{\hat{B}}{\hat{g}_0(\bar{m})},$$

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. We calculate standard errors for all estimated parameters using a bootstrap procedure, as in Chetty et al. (2011).¹²

There are two key identifying assumptions necessary for equation (12) to provide a valid estimate of the behavioral response to the conforming limit. The first is that the counterfactual loan size distribution that would exist in the absence of the limit is smooth. That is, any spike in the loan size distribution at the conforming limit must be solely attributable to the existence of the limit and not some other factor. We test for violations of this assumption below by examining how the distribution of loan sizes changes when the conforming limit moves from one year to another. The second assumption is that households can be uniquely indexed by their counterfactual choice of mortgage size in the absence of the limit—that is, there is a well-defined marginal buncher. While this assumption is fundamentally untestable, most reasonable models of mortgage choice would imply such a result.

In order to estimate the components of equation (12), there are several free parameters that we must choose: the bin width $\left(\left| \frac{m_0 - m_L}{L} \right| \right)$, the order of the

¹² At each iteration (k) of the bootstrap loop we draw with replacement from the estimated errors, $\hat{\varepsilon}_j$, in equation (9) to generate a new set of bin counts, n_j^k . We then re-estimate bunching using these new counts. Our estimate of the standard error for $\Delta \hat{m}$ 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.

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}).¹³ For the other three parameters, our preferred specification uses 1-percent bins, a thirteenth-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.¹⁴

B. Estimating the Jumbo-Conforming Spread

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

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 flexibly for all other relevant determinants of interest rates. In particular, we estimate variants of the following equation

$$(13) \quad r_{i,t} = \alpha_{z(i),t} + \beta J_{i,t} + f^{J=0}(m_{i,t}) + f^{J=1}(m_{i,t}) + s^{LTV}(LTV_{i,t}) + s^{DTI}(DTI_{i,t}) \\ + s^{FICO}(FICO_{i,t}) + PMI_{i,t} + PP_{i,t} + g(TERM_{i,t}) + \epsilon_{i,t},$$

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

¹³This 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.

¹⁴It is worth noting, however, that the estimated missing mass from the right of the limit need not be exactly equal to the number of bunched loans since the procedure we use to estimate bunching ignores both extensive margin responses and the leftward shift of the distribution outside of the excluded region generated by intensive margin responses among those who do not bunch. If these types of responses have a sizable effect on the observed loan size distribution, then choosing parameters to minimize the difference between bunching and missing mass could lead to bias in the estimated behavioral response. To account for this, we explore robustness to various choices of the underlying parameters, which often yield estimates of \hat{B} that are smaller than \hat{M} but give very similar estimates of $\Delta \hat{m}$ to our preferred specification.

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 for whether the borrower took out private mortgage insurance (PMI) and if the mortgage had a prepayment penalty (PP). Finally, we also control flexibly for the length of the mortgage ($TERM$).¹⁵ 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.

Instrumenting for Jumbo Loan Status.—Of course, in a finite sample, it is not possible to control completely for all observed determinants of interest rates, and there may be some unobserved characteristics which our controls are unable to capture. If these unobserved characteristics are also correlated with jumbo loan status, then estimates of β based on equation (13) 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 (13) in which we instrument for jumbo loan status using a discontinuous function of the value of the home, following Kaufman (2012).¹⁶

When making a loan, lenders typically require an independent appraisal as a check that the agreed transaction price accurately reflects the value of the home. The “value” in the denominator of the LTV ratio is then set as the lesser of the appraised value and the transaction price. Many buyers choose a loan size that would yield an LTV of exactly 80 percent, both because it is a longstanding norm and because exceeding 80 percent typically requires purchasing private mortgage insurance. Consequently, if the value of the home is just under the conforming loan limit multiplied by 1.25, then a buyer is substantially less likely to end up with a jumbo loan than if the value is just over this limit. Figure 3, panel A confirms this, showing that the fraction of loans that are jumbos jumps discontinuously as the value crosses 125 percent of the conforming limit.

This fact suggests an approach in which we instrument for $J_{i,t}$ in equation (13) with an indicator for whether the value of the home falls above or below $1.25 \bar{m}$, while controlling flexibly for home value on either side of this cutoff. The key to the exogeneity of this instrument is that borrowers have essentially no control over the exact outcome of their appraisal and are also unlikely to be able to finely control their transaction price. As a result, some borrowers who purchase homes with values close to $1.25 \bar{m}$ —and who either do not want to or are unable to deviate from the 80 percent LTV norm—may be induced into or out of jumbo loan status based in part on factors that are outside their control. As supporting evidence for this assumption, Figure 3, panel B shows that the distribution of house values around

¹⁵ The exact specifications are described in the results section below.

¹⁶ Adelino, Shoar, and Severino (2012) and Fuster and Vickery (2015) employ similar strategies to look at the effects of the conforming limit on house prices and on mortgage supply, respectively.

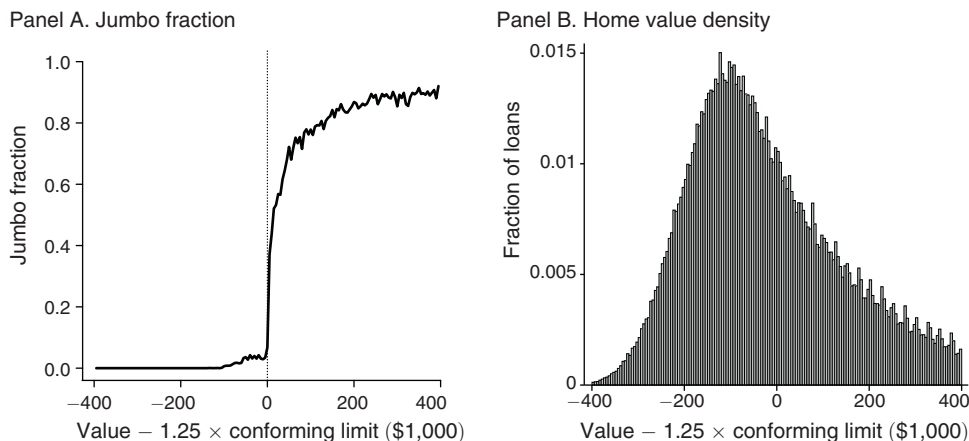


FIGURE 3. JUMBO LOAN FRACTION AND HOME VALUE DENSITY AROUND 125 PERCENT OF THE CONFORMING LOAN LIMIT

Notes: Panel A plots the fraction of loans in each bin with loan amounts greater than the conforming limit (jumbo loans). Panel B plots the fraction of all loans with values that are in any given \$5,000 bin relative to 125 percent of the conforming limit. Value is measured as the minimum of the appraisal amount and the contract price. Data are pooled across years and each loan is centered at 125 percent of the conforming limit in effect at the date of origination, so that a value of zero represents a loan at exactly 125 percent of the conforming limit. Sample includes all transactions in the primary LPS sample that fall within \$400,000 of the conforming limit. See text for details on sample construction.

125 percent of the conforming limit is quite smooth, in contrast with the distribution of loan amounts around the conforming limit itself, as in Figure 1, panel B.

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

IV. 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 V to calculate elasticities. In Sections VC and VD, we discuss the ways in

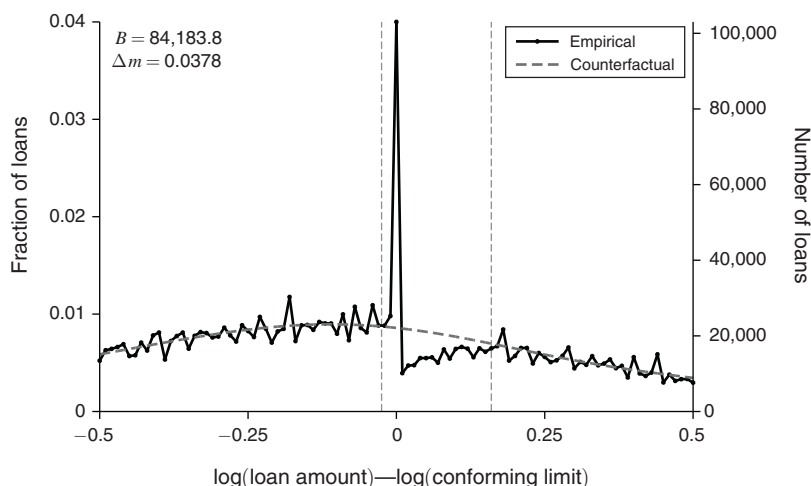


FIGURE 4. BUNCHING AT THE CONFORMING LIMIT, ALL LOANS

Notes: This figure plots the empirical and counterfactual density of (log) loan size relative to the conforming limit for all loans. Estimation was carried out in the full sample of DQ loans, but the figure shows only loans within 50 percent of the conforming limit. The solid black line is the empirical density. Each dot represents the count (fraction) of loans in a given 1 percent bin relative to the limit in effect at the time of origination. The heavy dashed gray line is the estimated counterfactual density obtained by fitting a thirteenth 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 IIIA.

which borrowers appear to be adjusting their loan sizes and provide context for the magnitude of the estimates.

A. Bunching at the Conforming Limit

Results for all Borrowers.—As a starting point for our empirical analysis, Figure 4 plots both the observed (log) loan size distribution and the counterfactual distribution estimated from the bunching procedure using all available loans in the DQ sample. Although our estimation is carried out in the full sample, in this figure and in subsequent figures we have narrowed our focus to the range of loans which fall within 50 percent of the conforming limit. The x-axis shows the difference between the log loan amount and the log conforming limit in the year the loan was originated, so that zero is the limit itself, and each bin represents roughly a 1 percent incremental deviation from the limit. The y-axis on the right indicates the number of loans in each bin, while the y-axis on the left indicates the fraction of all loans represented by that number.

The solid black connected line plots the histogram of (log) loan size, which exhibits a sharp peak at the limit. This bin contains approximately 100,000 loans, representing 4 percent of the entire sample, which is roughly 4 times as many loans as in the bin immediately to the left. This distribution 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 dashed gray line shows the fitted polynomial that we take as our counterfactual loan size distribution. The

TABLE 1—BUNCHING ESTIMATES BY LOAN TYPE

	Combined (1)	FRMs (2)	ARMs (3)
Bunched loans (\hat{B})	84,183.8 (2,687.0)	68,306.3 (1,561.6)	16,174.3 (1,446.7)
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)

Notes: 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 IIIA.

vertical dashed lines represent the lower (m_L) and upper (m_H) limits of the excluded region, as defined in Section IIIA.

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 solid and dashed 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 1 repeats these estimates along with their standard errors and several other parameters estimated during the bunching procedure. As another way of gauging the magnitude of the response, the third row of Table 1 reports a measure of the “excess mass” at the conforming limit. We calculate this as the ratio of the number of loans bunching at the limit to the number of loans that 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 1, we also report the upper limit of the excluded region used in estimation (m_H). If there were no extensive margin responses (borrowers leaving the market entirely), then this number would provide an estimate of the largest percent reduction in mortgage size among bunching individuals. That is, no individual with a counterfactual loan size more than m_H percent larger than the conforming limit would be induced to bunch. Given extensive margin responses, it is possible that our estimate of m_H differs from the true cutoff value. Nonetheless, it provides a useful gauge of the magnitude of behavioral responses among those who reduce their mortgage sizes the most. The estimate implies an upper bound on behavioral responses of roughly 16 percent, meaning that nearly all of the borrowers

bunching at the conforming limit would have had mortgages that were less than 16 percent larger than the limit had it not existed.

While the bunching evidence reported in Figure 4 is quite stark, it is interesting to note that there are still a number of borrowers who locate just above the conforming limit despite the large marginal interest rate associated with doing so. There are three potential explanations for this behavior: inelastic preferences, heterogeneous costs of adjusting first mortgage balances, or borrower-level differences in the magnitude of the jumbo-conforming spread. In our analysis of the magnitude of the jumbo-conforming spread, we are able to largely rule out the third explanation; we find no robust evidence of heterogeneity in the spread across borrower credit scores or the distribution of zip code-level median income or price-to-income ratios. This suggests that the borrowers who locate in the region just above the conforming limit are likely doing so as a result of either inelastic preferences or adjustment costs.

In the context of income taxes, Kleven and Waseem (2013) are able to distinguish between the role of preferences and adjustment costs using the fact that a notch can sometimes create a dominated region in which no wage earner, regardless of tax elasticity, would choose to locate in the absence of adjustment costs. By counting the number of wage earners observed in the dominated region they are able to back out an estimate of adjustment costs. Unfortunately, we cannot perform a similar exercise here because there is no such dominated region in our setting. In the terminology of Kleven and Waseem (2013), the jumbo-conforming spread creates a “downward notch” where, for any finite loan amount above the limit, there exists a first mortgage demand elasticity sufficiently close to zero, such that some borrower would be willing to take out that loan. It is therefore not possible to estimate the magnitude of any potential adjustment costs in this setting. Because of this, it is important to note that the elasticities we estimate are necessarily “reduced form,” in the sense that they incorporate the effect of adjustment costs and are not driven entirely by the intertemporal elasticity of substitution alone.

Finally, as noted above, one of the key assumptions necessary for bunching to identify behavioral responses is that the counterfactual loan size distribution be smooth. While there is no way to directly test this assumption, one way to evaluate its plausibility is to examine what happens to the empirical distribution of loan sizes as the conforming limit moves from one year to the next. If the conforming limit is the only thing causing bunching, then bunching should track the movement of the conforming limit, and the distribution should be smooth at previous and future conforming limits. This is exactly the outcome we find. At any given nominal loan amount, the loan size distribution appears to be smooth except in the year for which that loan amount serves as the conforming limit. For example, Figure 5 plots the empirical loan size distribution separately for the years 2000, 2002, and 2004. While not definitive, this result strongly supports the counterfactual smoothness assumption.

Fixed versus Adjustable Rate Mortgages.—In addition to looking at the effect of the conforming limit on overall loan size, recent work by both Fuster and Vickery (2015) and Kaufman (2012) draws attention to several stylized facts that make it particularly interesting to investigate heterogeneity in the response by *type* of

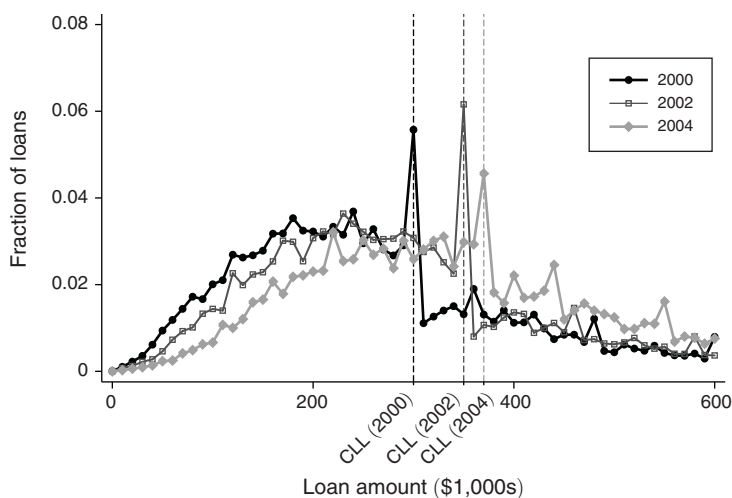


FIGURE 5. BUNCHING ADJUSTING OVER TIME

Notes: This figure plots the empirical density of loan sizes separately for three different years. Each dot represents the fraction of loans originated in a given \$10,000 bin in the indicated year. The vertically dashed lines mark the conforming limit that was in place in the indicated year. All dollar amounts are nominal.

loan.¹⁷ 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 6 using our own sample of loans from DataQuick. Using the same 1 percent bins as before, this figure plots the share of loans that are FRMs as a function of loan size relative to the conforming limit. To the left of the limit, the FRM share declines gradually as the loan amount increases, reaching roughly 55 percent just below the limit. It then spikes to about 75 percent at the limit before falling to 20 percent immediately to the right. Beyond the conforming limit, the share then rises, eventually reaching a plateau of about 35 percent.

This drop in the FRM share is not a coincidence. Fixed-rate mortgages are generally estimated to have a larger jumbo-conforming spread relative to ARMs due to the fact that their returns are much more vulnerable to interest rate risk.¹⁸ Since the FRM share well to the right of the limit is substantially lower than the FRM share to the left of the limit, a quick glance at Figure 6 might suggest an extensive margin response. That is, in response to the higher jumbo spread for FRMs, some jumbo borrowers may choose to substitute toward ARMs.

In contrast, we argue that the change in the FRM share at the conforming limit occurs because more FRM borrowers than ARM borrowers bunch at the limit,

¹⁷In the online Appendix, we also examine heterogeneity by classifying borrowers by race and income. We find that low-income or minority borrowers are substantially less prone to bunching at the limit than high-income or non-minority borrowers.

¹⁸We replicate this well-documented difference in spreads using our own sample of loans from LPS in Section IVB. 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.

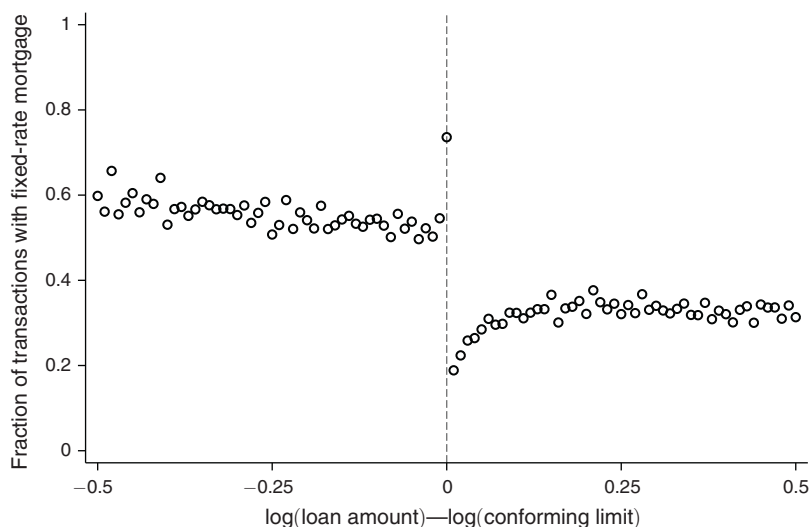


FIGURE 6. FIXED-RATE MORTGAGE SHARE RELATIVE TO THE CONFORMING LIMIT

Notes: 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.

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

Figure 7, panels A and B show the results from this exercise for FRMs and ARMs, respectively. The standard errors and additional bunching parameters are also reported separately for each loan type in columns 2 and 3 of Table 1. While Figure 7, panel A shows a substantial downward shift in the FRM distribution to the right of the limit, Figure 7, panel B 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.¹⁹

While we do not believe that our results invalidate any of the conclusions drawn by Fuster and Vickery (2015) 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 the differential

¹⁹Of 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.

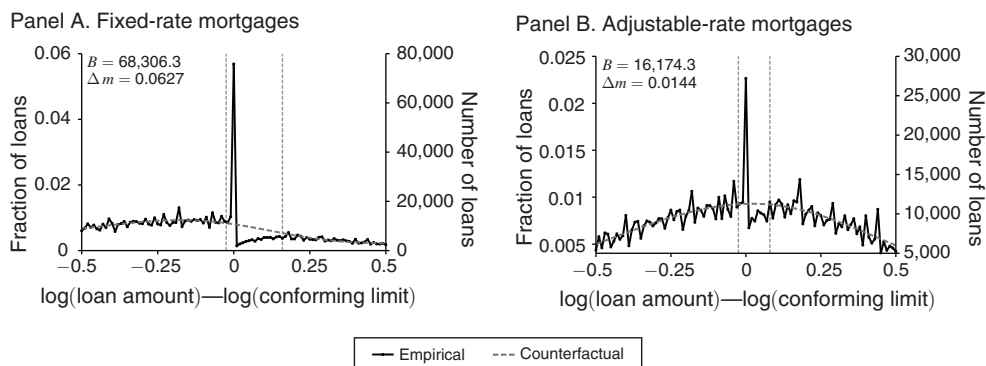


FIGURE 7. BUNCHING AT THE CONFORMING LIMIT, FRMs AND ARMs

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

bunching responses of FRM and ARM borrowers in mind, in the next section we present estimates for FRMs and ARMs separately.

B. Jumbo-Conforming Spread

To convert the behavioral responses estimated from bunching into elasticities, we next need to obtain an estimate of the interest rate differential at the limit. Table 2 presents estimates of the jumbo-conforming spread, following the strategies discussed in Section IIIB. 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 (or distance to 125 percent of the limit in the IV results) interacted with the jumbo loan indicator variable, as well as controls for the loan-to-value (LTV) ratio, debt-to-income (DTI) ratio, missing DTI ratio, FICO credit score, missing FICO score, whether the loan includes private mortgage insurance (PMI), and whether the loan has a prepayment penalty. They also include zip-code by month fixed effects and fixed effects for standard loan lengths, such as 15, 30, and 40 years, as well as a linear term to capture the effects of nonstandard lengths.

Across the columns of the table, the four specifications are: (1) a baseline, using all available data, with linear controls for the LTV and DTI ratios and the FICO score; (2) the same specification replacing the linear controls for LTV, DTI, and FICO with more flexible cubic B-splines; (3) the same specification as in (2) but with a sample limited to loans within \$50,000 of the conforming limit; and (4) the same specification in (2) but with a sample limited to loans within \$10,000 of the limit.

TABLE 2—JUMBO-CONFORMING SPREAD ESTIMATES, PERCENTAGE POINTS

	Baseline (1)	Splines (2)	Within \$50k of CLL (3)	Within \$10k of CLL (4)
<i>Fixed-rate mortgages</i>				
OLS	0.178 (0.002)	0.181 (0.002)	0.164 (0.015)	0.141 (0.039)
IV	0.108 (0.010)	0.123 (0.009)	0.122 (0.029)	0.147 (0.057)
Observations	1,011,345	1,011,345	257,280	85,664
<i>Adjustable-rate mortgages</i>				
OLS	−0.064 (0.009)	−0.079 (0.009)	−0.287 (0.037)	−0.357 (0.083)
IV	−0.001 (0.020)	0.005 (0.019)	0.075 (0.054)	0.034 (0.109)
Observations	676,718	676,718	155,740	39,112

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

For FRMs, applying least squares yields estimates of the jumbo-conforming spread that start around 18 basis points and then shrink slightly in columns 3 and 4 as the window around the limit narrows. These estimates are a bit smaller than Sherlund's (2008) estimate of 22 basis points, but some difference should be expected given that Sherlund's (2008) data are derived from a nationally representative survey, whereas our data come from actual mortgage originations and only cover California.

As discussed above, the OLS specifications do not control for borrower selection on unobservables around the conforming limit, which is a particular concern given the substantial bunching that we have just highlighted. To address this potential selection issue, the row labeled "IV" presents estimates of the spread when we instrument for the jumbo indicator with an indicator for whether the value of the house exceeds 125 percent of the conforming limit.²⁰ As would be expected, these estimates are somewhat less precise than the OLS results but are all still significantly different than zero at conventional levels. The estimates run from 11 to 15 basis points. In the first two columns they are smaller than the comparable OLS estimates, 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. The IV and OLS estimates in columns 3 and 4 are statistically indistinguishable. All of the IV estimates are similar to Kaufman's (2012) estimate of 10 basis points, using essentially the same technique but a different sample of loans.

²⁰ Regressing the jumbo indicator on the indicator for whether the value exceeds 125 percent of the conforming limit, including all of the other covariates as controls—that is, the first stage of the IV—yields coefficients of around 0.2, with modest variation depending on the specification. The instrument is quite strong; in all cases, the standard errors on this coefficient are tiny and the *F*-statistics very large.

Our estimates of the jumbo-conforming spread for initial ARM rates, presented in the lower half of Table 2, are considerably more noisy. 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. The “IV” row uses the same instrument as in the FRM results. Here we see coefficients that range from about 0, in the first two columns, to 8 and 3 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 for ARMs that are also close to zero (Kaufman 2012, Fuster and Vickery 2015).

Given that bunching at the limit is less prevalent but still apparent among ARM borrowers, it is surprising to find no strong support for a positive jumbo-conforming spread among ARM loans and, indeed, large negative spreads in some specifications. 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. For example, as suggested by Kaufman (2012), the lower initial rates for jumbo ARMs may be more than offset by other unfavorable terms, like the rates to which the ARMs eventually reset. Regardless, given the difficulty of explaining the negative spreads, the remaining discussion in this paper—including the calculation of elasticities—focuses on the results for FRMs.

V. Elasticities

A. Calculating Elasticities with a Notched Budget Constraint

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

Specifically, for $m > \bar{m}$, define $r^*(m)$, such that

$$(14) \quad (m - \bar{m}) \cdot r^*(m) = m \cdot (\hat{r} + \Delta\hat{r}) - \bar{m} \cdot \hat{r}.$$

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

$$(15) \quad r^*(m) = \hat{r} + \Delta\hat{r} + \Delta\hat{r} \cdot \frac{\bar{m}}{m - \bar{m}}.$$

Equation (15) makes clear that $r^*(m)$ is equal to the jumbo rate ($\hat{r} + \Delta\hat{r}$) plus a term that is increasing in the jumbo-conforming spread ($\Delta\hat{r}$) and decreasing in the size of the loan relative to the conforming limit ($m - \bar{m}$). For loans just above the limit this additional term is very large, reflecting the fact that the higher interest rate on jumbo loans is applied to the full balance of the loan.²¹ Given an estimate of $r^*(m)$, we can then calculate the (semi-)elasticity of (first) mortgage demand implied by our estimate of $\Delta\hat{m}$,

$$(16) \quad \epsilon_s^r = \frac{\Delta\hat{m}}{r^*(\bar{m} + \Delta\hat{m}) - \hat{r}}.$$

As before, $\Delta\hat{m}$ is estimated in logs, so it represents the approximate percentage change in mortgage demand induced by the conforming limit, while the denominator measures the level change in interest rates. We present our estimates as semi-elasticities because it is a bit more intuitive to consider changes in interest rates in basis or percentage points.

B. Estimated Elasticities of First Mortgage Demand

Table 3 reports 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 1 (0.063) is shown in the middle row. The other two estimates (0.052 and 0.083) are the smallest and largest estimates of $\Delta\hat{m}$ across a range of different options for the three parameters chosen ex ante: the bin width, the polynomial order, and the lower limit of the excluded region.²² They provide reasonable bounds on the variation in the elasticity implied by these parameters. The jumbo-conforming spread estimates are taken from column 3 of Table 2 and correspond to the OLS (column 1) and IV (column 2) estimates, respectively.

The estimated semi-elasticities range from -0.016 to -0.052 , with our preferred estimates in the middle row at -0.023 and -0.030 . The associated standard errors are relatively small, although those using the less precise 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

²¹ For example, in 2006 the conforming limit was \$417,000 and the average interest rate on loans made just below that was 6.37 percent. Given our OLS estimate of the jumbo-conforming spread of 16 basis points, this implies an average rate of 6.53 percent for a loan made just above the limit. Thus, the marginal interest rate on the last \$1,000 of a \$418,000 loan originated in 2006 was $r^*(418,000) = 6.53 + 0.16 \times 417 = 73.3$ percent.

²² We considered bin widths of 0.01, 0.025, and 0.05; polynomials of order 7, 9, 11, and 13; and lower limits of 0.025, 0.05, 0.075, and 0.1.

TABLE 3—INTEREST RATE SEMI-ELASTICITY OF MORTGAGE DEMAND (ϵ_s^r),
FRMs ONLY

$\Delta \hat{m}$	$\Delta \hat{r}$	
	(1)	(2)
	0.164 (0.015)	0.122 (0.029)
0.052 (0.003)	−0.016 (0.002)	−0.021 (0.005)
0.063 (0.003)	−0.023 (0.002)	−0.030 (0.007)
0.083 (0.004)	−0.039 (0.004)	−0.052 (0.013)

Notes: Table reports estimates and standard errors (in parentheses) of the interest rate semi-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 ($\Delta \hat{r}$). Standard errors for the bunching estimates were calculated using the bootstrap procedure described in Section IIIA. Standard errors for the elasticities were calculated using the delta method.

5 percent to 6 percent—100 basis points—would lead to a decline in first mortgage demand of 2 to 3 percent.

C. Accounting for Second Mortgages

The elasticities reported in Table 3 tell us how much borrowers reduce their first mortgage balance in response to the jumbo-conforming spread, but they do not tell us how the borrowers adjust. There are three primary channels through which a borrower can reduce the size of her first mortgage, each of which have different implications for the interpretation of our main results. First, a borrower could simply bring more cash to the table, making a larger down payment and taking out a smaller loan.²³ Second, she could take out an additional mortgage for the amount of debt desired in excess of the conforming limit. Finally, she could spend less on housing either by substituting to a lower-quality home or negotiating a lower price with the seller. In either case, she would end up with a smaller mortgage (holding leverage constant).

To measure the extent to which borrowers are using second mortgages to lower their first mortgage balance, Figure 8 plots the number of transactions financed using second loans as a function of the associated fixed-rate first mortgage value relative to the conforming limit. Consistent with the notion that many of the bunching borrowers take out second mortgages, there is a sharp spike in the number of transactions which are financed with two loans precisely at the limit. The plot suggests that roughly 25,000 more second loans were taken out in the bin at the conforming limit relative to the bin just below it, which provides a reasonable counterfactual.

²³ “Bringing more cash to the table” could be accomplished in many ways, including taking money out of savings, reducing current consumption, or taking out non-mortgage debt. We cannot distinguish between these cases.

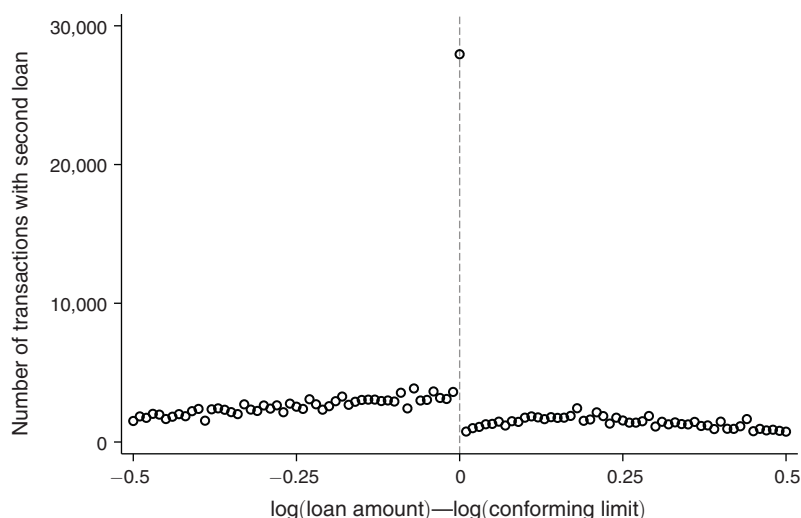


FIGURE 8. NUMBER OF SECOND MORTGAGES BY FIRST MORTGAGE AMOUNT

Notes: 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.

Combining this with our estimate from Table 1 that about 70,000 FRM borrowers bunch at the limit suggests that roughly 35 percent of FRM borrowers who bunch do so by taking out a second mortgage. These borrowers are presumably shifting debt from their first mortgage onto their second, holding combined LTV roughly constant while reducing their first-mortgage LTV.

The remaining 65 percent of “extra” borrowers must be either putting up more cash, or spending less on housing than they otherwise would. In the online Appendix, we attempt to further distinguish between these two margins using back-of-the-envelope calculations involving the observed first mortgage LTVs and combined LTVs at the conforming limit. These calculations suggest that the 65 percent who do not take out second mortgages are likely putting up cash rather than buying cheaper houses. Regardless of how they adjust, however, the roughly two-thirds of borrowers who bunch at the conforming limit without the use of a second mortgage must be reducing their total mortgage debt.

Our bunching measure provides estimates of the mortgage demand elasticity for the marginal bunching borrower, while the estimate of second mortgage use is for the average bunching borrower. Since the fraction of all bunching borrowers who take out a second mortgage is not necessarily the same as the fraction of marginal bunching borrowers who do so, we cannot definitively solve for the response of total mortgage debt among the latter group. If we assume, however, that the use of second mortgages is the same among the average and marginal bunchers, then we can scale down our elasticity estimates by a factor of one-third to provide a rough estimate of the effect of a change in rates on total mortgage debt. Specifically, multiplying our preferred first mortgage semi-elasticity estimates of -0.023 and -0.030 by

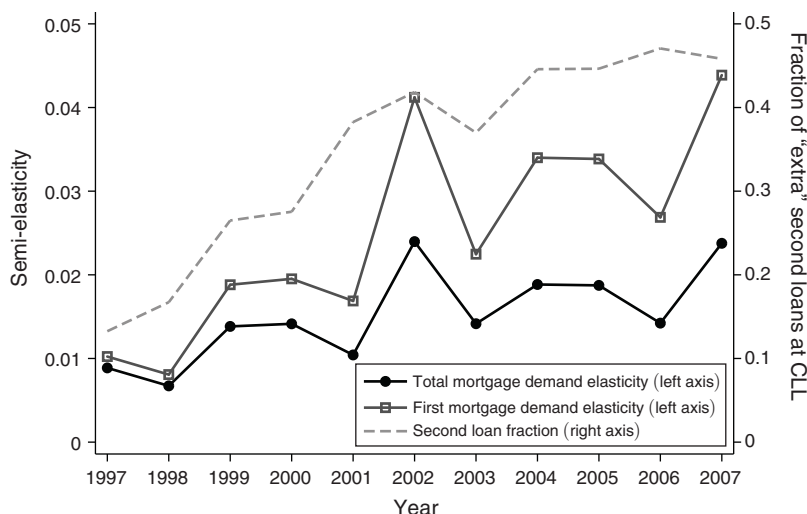


FIGURE 9. FIRST MORTGAGE AND TOTAL MORTGAGE DEMAND ELASTICITIES BY YEAR

Notes: This figure plots annual estimates of the first mortgage and total mortgage (semi-)elasticity of demand for each year in the analysis sample. The solid gray line with hollow squares reports the first mortgage demand elasticity, which is estimated using only fixed-rate mortgages originated in the indicated year. The yearly bunching and jumbo-conforming spread estimates underlying each elasticity are generated using the same parameters and specifications used to estimate bunching and the jumbo-conforming spread in the overall sample as reported in column 2 of Table 1 and column 2 of Table 2, respectively. The dashed line reports the fraction of “extra” second loans observed at the conforming loan limit. This fraction is calculated as the ratio of the number of transactions financed using two loans in the 1-percent bin containing the conforming limit to the number of transactions financed using two loans in the bin just below the limit. The line with solid circles reports the total mortgage demand elasticity, which is estimated by scaling down the first mortgage demand elasticity by the fraction of “extra” second loans observed at the conforming limit in the relevant year.

two-thirds yields total debt semi-elasticities of about -0.015 and -0.020 . That is, a 1 percentage point increase in rates should reduce total mortgage debt by between 1.5 and 2 percent.

Interestingly, the fraction of bunchers who use second mortgages is not constant over time. Figure 9 plots yearly estimates of this fraction (right axis) as well as estimates of the first mortgage demand elasticity and the implied elasticity of total mortgage demand (left axis) for each year in our analysis sample. As shown by the dashed line, between 1997 and 2007, the fraction of bunching borrowers making use of a second loan nearly quadrupled from a base of just over 0.1 to about 0.45. This rising second loan fraction was accompanied by an equally large increase in the first mortgage demand elasticity, plotted as a solid gray line (hollow squares). One possible explanation for the growing second mortgage share and first mortgage demand elasticity could be eroding credit standards, which made second mortgages a more available means for adjusting first mortgage balances. Regardless of the reason for this pattern, however, when the rising second mortgage share is taken into account, the implied total mortgage demand elasticity plotted by the solid black line (solid circles) is more stable and generally hovers in the same range as our baseline pooled estimate of 0.015 to 0.02.

D. Economic Magnitudes and Interpretation

Are our baseline estimates large or small? In an absolute sense, they appear to be quite small. For example, during our sample period the average interest rate on a loan taken out just below the conforming limit was approximately 6.5 percent.²⁴ Our semi-elasticity estimates imply that a 1 percentage point increase from that baseline (a 15 percent increase) leads to a reduction in total mortgage debt of only 1.5 to 2 percent.

While we are wary of applying our estimate too far outside the context in which it is estimated, we can nonetheless get a rough sense of its magnitude by comparing it to macro time series variation in rates and mortgage origination volume. For example, since 1990, the standard deviation of the change in the average annual rate on a 30-year fixed-rate mortgage has been a bit more than 0.5 percentage points, while the standard deviation of the percentage change in the dollar volume of purchase mortgage originations has been about 17 percentage points.²⁵ Given these figures, a naïve multiplication of our elasticity by the standard deviation in interest rate changes would suggest that rate changes can explain less than one-tenth of the variation in new purchase mortgage origination volume.

Our finding of a small intensive-margin behavioral response to changes in interest rates is in line with several other findings from the literature. For example, Fuster and Zafar (2015) survey households about their mortgage choices under randomized hypothetical interest rate scenarios and find that their chosen down payment fractions respond in a way that would imply even smaller intensive-margin semi-elasticities ranging between 0.6 and 1.8. Several studies of behavioral responses to interest rate changes in other consumer credit markets have found similarly small effects. For example, Attanasio, Goldberg, and Kyriazidou (2008) and Karlan and Zinman (2008) document small responses to interest rate changes for auto loans and micro-credit, respectively.

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

²⁴This figure is the average interest rate across all fixed-rate mortgages falling less than \$10,000 below the conforming limit in the LPS data. We get essentially the same number if we use our models to predict interest rates for everyone in our sample as if they had taken out a loan just below the conforming limit.

²⁵The interest rate calculation uses data from the Freddie Mac Primary Mortgage Market Survey. The origination calculation uses data on purchase mortgages collected under the Home Mortgage Disclosure Act.

substitutes (Gross and Souleles 2002). For these and other reasons, our elasticity estimate may not be applicable outside of the specific institutional setting in which we estimate it.²⁶

VI. Policy Application: GSE Guarantee Fee Increases

With that caveat in mind, we now illustrate how our elasticity estimates could be used to gauge the potential effects of one recently proposed change to the parameters of this institutional setting.

In February 2012, the Federal Housing Finance Agency (FHFA) published its “Strategic Plan for Enterprise Conservatorships,” outlining the steps that the agency planned 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” (Federal Housing Finance Agency 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. 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 2012, g-fees from single-family mortgages generated roughly \$12.5 billion in revenue for the GSEs, up 12 percent from the previous year (Federal Housing Finance Agency Office of Inspector General 2013).

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. 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 (Federal Housing Finance Agency 2013).²⁷ 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.” (Federal Housing Finance Agency Office of Inspector General 2013).²⁸

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

²⁶ Other things held constant in our setting, but not in most macro settings, include the availability of second loans with rates that do not change at the limit, as well as the more general term structure of interest rates and other rates of return.

²⁷ FHFA reports this 11 basis point number as the combined overall effect of an approximate 14 basis point increase in the g-fee for 30-year mortgages and a 4 basis point increase for 15-year mortgages, on average across all Freddie and Fannie loans.

²⁸ This plan was delayed by incoming director Mel Watt in January 2014 and its future remains uncertain as of early 2015.

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 nonconforming 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 (Federal Housing Finance Agency Office of Inspector General 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.²⁹

VII. Conclusion

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

We conclude by pointing to three potentially useful avenues for future research. First, our estimates are necessarily limited by their context. A large number of salient factors, especially the presence of adjustment costs and the availability of second mortgages, affect how borrowers respond to the limit and, in turn, our estimates of the demand elasticity. A better understanding of the importance of these factors for our estimates is required before they can be applied to more general policy questions. Second, our estimates abstract away from the potential effect of

²⁹ While our elasticity is well suited to examining the direct intensive margin effects of these relatively small changes in fees, we cannot draw any inference on more general questions, such as how high the fees would need to rise to draw private capital back into this part of the MBS market, or what the effects of the fee increases would be on extensive margin responses.

interest rates on the extensive margin choice of whether to purchase a home or not. A full accounting of the effect of interest rates on mortgage demand would need to incorporate this margin as well. Finally, although we show in the online Appendix that there seem to be differential responses of minority versus non-minority and high-income versus low-income borrowers, data limitations prevent us from being able to fully investigate this heterogeneity. Painting a fuller picture of heterogeneity in elasticities and adjustment costs may be as important as pinning down the overall average elasticity of demand.

REFERENCES

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino.** 2012. "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation." National Bureau of Economic Research (NBER) Working Paper 17832.
- Alessie, Rob, Stefan Hochguertel, and Guglielmo Weber.** 2005. "Consumer Credit: Evidence from Italian Micro Data." *Journal of the European Economic Association* 3 (1): 144–78.
- Attanasio, Orazio P., Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou.** 2008. "Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans." *International Economic Review* 49 (2): 401–36.
- Best, Michael Carlos, James Cloyne, Ethan Ilzetzki, and Henrik Jacobsen Kleven.** 2015. "Interest rates, debt and intertemporal allocation: Evidence from notched mortgage contracts in the United Kingdom." Bank of England Staff Working Paper 543.
- Best, Michael Carlos, and Henrik Jacobsen Kleven.** 2016. "Property Transaction Taxes and the Housing Market: Evidence from Notches and Housing Stimulus in the UK." https://web.stanford.edu/~mbest/best-kleven_landnotches_feb2016.pdf.
- Brueckner, Jan K.** 1994. "The Demand for Mortgage Debt: Some Basic Results." *Journal of Housing Economics* 3 (4): 251–62.
- Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri.** 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *Quarterly Journal of Economics* 126 (2): 749–804.
- Chetty, Raj, John N. Friedman, and Emmanuel Saez.** 2013. "Using Differences in Knowledge across Neighborhoods to Uncover the Impact of the EITC on Earnings." *American Economic Review* 103 (7): 2683–2721.
- DeFusco, Anthony A., and Andrew Paciorek.** 2017. "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit: Dataset." *American Economic Journal: Economic Policy*. <https://doi.org/10.1257/pol.20140108>.
- Dunsky, Robert M., and James R. Follain.** 2000. "Tax-Induced Portfolio Reshuffling: The Case of the Mortgage Interest Deduction." *Real Estate Economics* 28 (4): 683–718.
- Federal Housing Finance Agency.** 2012. *A Strategic Plan for Enterprise Conservatorships: The Next Chapter in a Story that Needs an Ending*. Federal Housing Finance Agency. Washington, DC, February.
- Federal Housing Finance Agency.** 2013. *FHFA Takes Further Steps to Advance Conservatorship Strategic Plan by Announcing an Increase in Guarantee Fees*. Federal Housing Finance Agency. Washington, DC, December.
- Federal Housing Finance Agency Office of Inspector General.** 2013. "FHFA's Initiative to Reduce the Enterprises' Dominant Position in the Housing Finance System by Raising Gradually Their Guarantee Fees." https://www.fhfa.ig.gov/Content/Files/EVL-2013-005_4.pdf.
- Follain, James R., and Robert M. Dunsky.** 1997. "The Demand for Mortgage Debt and the Income Tax." *Journal of Housing Research* 8 (2): 155–99.
- Fuster, Andreas, and James Vickery.** 2015. "Securitization and the Fixed-Rate Mortgage." *Review of Financial Studies* 28 (1): 176–211.
- Fuster, Andreas, and Basit Zafar.** 2015. "The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey." Federal Reserve Bank of New York Staff Report 702.
- Gelber, Alexander M., Damon Jones, and Daniel W. Sacks.** 2015. "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test." http://home.uchicago.edu/~jls/Jones_ET_Updated.pdf.

- Gross, David B., and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1): 149–85.
- Hendershott, Patric H., and James D. Shilling. 1989. "The impact of the agencies on conventional fixed-rate mortgage yields." *Journal of Real Estate Finance and Economics* 2 (2): 101–15.
- Jappelli, Tullio, and Luigi Pistaferri. 2007. "Do people respond to tax incentives? An analysis of the Italian reform to the deductibility of home mortgage interest." *European Economic Review* 51 (2): 247–71.
- Jordà, Òscar, Moritz H. P. Schularick, and Alan M. Taylor. 2014. "Betting the House." National Bureau of Economic Research (NBER) Working Paper 20771.
- Karlan, Dean S., and Jonathan Zinman. 2008. "Credit Elasticities in Less-Developed Economies: Implications for Microfinance." *American Economic Review* 98 (3): 1040–68.
- Karlan, Dean, and Jonathan Zinman. 2014. "Long-Run Price Elasticities of Demand for Credit: Evidence from a Countrywide Field Experiment in Mexico." https://www.dartmouth.edu/~jzinman/Papers/LongRunCreditElasts2014_01.pdf.
- Kaufman, Alex. 2012. "The Influence of Fannie and Freddie on Mortgage Loan Terms." Federal Reserve Board Finance and Economics Discussion Series (FEDS) Paper 2012-33.
- Kleven, Henrik J., and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan." *Quarterly Journal of Economics* 128 (2): 669–723.
- Kopczuk, Wojciech, and David Munroe. 2015. "Mansion Tax: The Effect of Transfer Taxes on the Residential Real Estate Market." *American Economic Journal: Economic Policy* 7 (2): 214–57.
- Ling, David C., and Gary A. McGill. 1998. "Evidence on the Demand for Mortgage Debt by Owner-Occupants." *Journal of Urban Economics* 44 (3): 391–414.
- Manoli, Dayanand S., and Andrea Weber. 2011. "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions." National Bureau of Economic Research (NBER) Working Paper 17320.
- Martins, Nuno C., and Ernesto Villanueva. 2006. "The impact of mortgage interest-rate subsidies on household borrowing." *Journal of Public Economics* 90 (8–9): 1601–23.
- Mishkin, Frederic S. 1995. "Symposium on the Monetary Transmission Mechanism." *Journal of Economic Perspectives* 9 (4): 3–10.
- Passmore, Wayne, Shane M. Sherlund, and Gillian Burgess. 2005. "The Effect of Housing Government Sponsored Enterprises on Mortgage Rates." *Real Estate Economics* 33 (3): 427–63.
- Passmore, Wayne, Roger Sparks, and Jamie Ingpen. 2002. "GSEs, Mortgage Rates, and the Long-Run Effects of Mortgage Securitization." *Journal of Real Estate Finance and Economics* 25 (2): 215–42.
- Poterba, James M. 1984. "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." *Quarterly Journal of Economics* 99 (4): 729–52.
- Poterba, James M., and Todd Sinai. 2008. "Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income." *American Economic Review* 98 (2): 84–89.
- Poterba, James M., and Todd Sinai. 2011. "Revenue Costs and Incentive Effects of the Mortgage Interest Deduction for Owner-Occupied Housing." *National Tax Journal* 64 (2, Pt. 2): 531–64.
- Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy* 2 (3): 180–212.
- Sallee, James M., and Joel Slemrod. 2012. "Car notches: Strategic automaker responses to fuel economy policy." *Journal of Public Economics* 96 (11–12): 981–99.
- Sherlund, Shane M. 2008. "The Jumbo-Conforming Spread: A Semiparametric Approach." Federal Reserve Board Finance and Economics Discussion Series (FEDS) Paper 2008-01.
- Zinman, Jonathan. 2015. "Household Debt: Facts, Puzzles, Theories, and Policies." *Annual Review of Economics* 7 (1): 251–76.