

THE INTEREST RATE ELASTICITY OF MORTGAGE DEMAND:  
EVIDENCE FROM BUNCHING AT THE CONFORMING LOAN  
LIMIT (ONLINE APPENDIX)

**Anthony A. DeFusco\***  
Kellogg School of Management  
Northwestern University

**Andrew Paciorek†**  
Board of Governors  
of the Federal Reserve System

March 25, 2016

---

\*E-mail: [anthony.defusco@kellogg.northwestern.edu](mailto:anthony.defusco@kellogg.northwestern.edu)

†E-mail: [andrew.d.paciorek@frb.gov](mailto:andrew.d.paciorek@frb.gov)

# Appendix A Supplementary Material

## A.1 The GSEs and the Conforming Loan Limit

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

The GSEs play a large role and exert a substantial amount of influence in the mortgage market.<sup>1</sup> However, they are only allowed to purchase loans which satisfy a specific set of criteria as outlined by their regulator. These criteria include requirements for loan documentation, debt-to-income ratios, leverage, and a nominal cap on the dollar amount of any purchased loan. Loans which meet these criteria and are therefore eligible to be purchased by the GSEs are referred to as “conforming loans.” In this paper we are primarily interested in the cap on loan size, known as the “conforming limit.” Mortgages exceeding this limit are not eligible for GSE purchase and are referred to as “jumbo loans.”

Figure A.1 plots the conforming limit in nominal terms (solid line) and in real 2007 dollars (dashed line) for each year during our sample period. During this period, the GSEs were regulated by the Office of Federal Housing Enterprise Oversight (OFHEO), which set the limit each year based on changes in the national median house price. The limit was the same for all mortgages in a given year irrespective of local housing market conditions.<sup>2</sup> Following the trend in national house prices, the nominal limit increased from around \$215,000 in 1997 to its peak in 2006 and 2007 at approximately \$420,000. In real terms, the limit also rose sharply over this period, especially during the house price boom of the mid-2000s.

Interest rates on loans above the conforming limit are typically higher than those on comparable loans below the limit for two reasons. First, because the debt underlying the MBS

---

<sup>1</sup>As of 2010, the GSEs were responsible for nearly 50 percent of the approximately \$10.5 trillion in outstanding mortgage debt, either directly or through outstanding MBS (Jaffee and Quigley, 2012). More than 75 percent of all mortgages originated in 2011 passed through the hands of one of the GSEs (Kaufman, 2012).

<sup>2</sup>The only exceptions to this rule were Alaska, Hawaii, the Virgin Islands, and Guam, which were deemed to be high cost areas and had a 50 percent higher conforming limit prior to 2008. Since the housing crisis, the national conforming loan limit has been replaced by a more complicated series of limits set at the metropolitan level. All of the analysis in this paper pertains to the pre-2008 period.

issued by the GSEs is backed by an implicit government guarantee, investors are willing to accept lower yields in exchange for that guarantee.<sup>3</sup> Part of this savings is eventually passed on to borrowers in the form of lower interest rates on conforming loans.<sup>4</sup> Second, the GSEs are also granted several special privileges that private securitizers are not. These include access to a line of credit at the U.S. Treasury, exemption from disclosure and registration requirements with the Securities and Exchange Commission (SEC), and exemptions from state and local income taxes.<sup>5</sup> These advantages lower the cost of securitizing mortgages for the GSEs relative to private market securitizers, with some of the savings passed on to borrowers in the form of lower interest rates on loans below the conforming limit.

## A.2 Endogenous Housing Choice

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

Our model extends to cover endogenous housing choice, albeit at the cost of a closed-form solution. Starting with our original model, if we solve equation (2) from section 2 for  $c_1$  and substitute this, along with equation (3), into equation (1), we can rewrite the household's problem in terms of mortgage debt,

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

subject now only to the borrowing constraint (4).

Now we allow households to choose the quantity of housing services to purchase ( $h$ ), and this quantity has a direct effect on first-period utility, so that

$$V = \max_{m,h} \{u(y + m - ph, h) + \delta v(ph - (1 + r)m)\}, \quad (\text{A.2})$$

---

<sup>3</sup>The implicit guarantee became explicit in 2008 when the GSEs were placed under government conservatorship.

<sup>4</sup>Passmore et al. (2002) and Passmore et al. (2005) provide several theoretical explanations for how the savings from the guarantee are eventually passed down to mortgage borrowers.

<sup>5</sup>For a full description of the direct benefits conferred on the GSEs as a result of their special legal status see Congressional Budget Office (2001).

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

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

$$\frac{\partial V}{\partial m} = u_1 - \delta(1+r)v_1 = 0 \tag{A.3}$$

$$\frac{\partial V}{\partial h} = u_2 - (pu_1 - p\delta v_1) = 0. \tag{A.4}$$

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

While there are no obvious functional forms that allow us to derive equivalents to equation (5), the intuition remains the same. Under standard conditions, there are optimal  $m^*$  and  $h^*$ , both of which can shift in response to the notch in the interest rate schedule. Our bunching estimation will capture the shifts in  $m^*$ , which could result in part from changes in housing consumption ( $h^*$ ).

### A.3 Summary Statistics

Table A.1 presents summary statistics for our primary estimation sample from DataQuick as well as the sub-sample of transactions with first loan amounts within \$50,000 of the conforming limit that was in place in the year of the transaction. All dollar amounts here and throughout the analysis are converted to real 2007 dollars.

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

Table A.2 presents summary statistics from the LPS data for fixed-rate (FRM) and adjustable-rate (ARM) loans separately. Columns 1 and 3 report statistics for the full

analysis sample while columns 2 and 4 restrict the sample to loans within \$50,000 of the conforming loan limit. In general, the restricted samples for each loan type are quite similar to the full sample, suggesting that loans near the limit are reasonably representative of the entire sample, at least along these dimensions.

## A.4 Heterogeneity by Borrower Type

In addition to investigating bunching behavior by loan type, as in section 5.1.2, and over time, as in section 6.3, it is also interesting to examine whether bunching varies with the observable characteristics of borrowers. While the available information on borrower demographics is somewhat limited, we are able to provide several rough cuts of the data based on race and income by matching a subset of DQ transactions to loan application information made available through the Home Mortgage Disclosure Act (HMDA).<sup>6</sup> For this exercise we restrict attention to fixed-rate mortgages, where the sample sizes are largest. Using the race and ethnicity information in the HMDA data, we define a loan as belonging to a “minority” borrower if the primary loan applicant reports his race as black or his ethnicity as Hispanic and as belonging to a “non-minority” borrower otherwise. Similarly, we define a borrower as “low-income” if the income reported on the loan application was below the median income reported across all loans and “high-income” if the reported income is above the median.

Figure A.2 shows results from estimating bunching separately in each of these four subsamples. Panels (A) and (B) show results for non-minority and minority borrowers while panels (C) and (D) present the results for high-income and low-income borrowers, respectively. In each case there is substantial evidence of heterogeneous responses, with far less bunching among minority and low-income borrowers than among non-minority and high-income borrowers. For non-minority and high-income borrowers the estimated percentage reduction in loan size is roughly 7 to 8 percent while for minority and low-income borrowers it is closer to 4 to 5 percent.<sup>7</sup>

---

<sup>6</sup>The matching procedure uses information on the primary loan amount, lender name, Census tract, property type, and year. We successfully match about 60 percent of the larger DQ sample to observations in HMDA. Further details are available from the authors on request.

<sup>7</sup>Table A.3 confirms the visual impressions given by figure A.2, reporting the point estimates and standard errors for the bunching parameters we estimate. For each of the reported parameters, the standard errors are small enough to reject the null that bunching behavior is the same across the high-income and low-income samples, as well as across the non-minority and minority samples. Part of the smaller response among low-income borrowers could be driven by the fact that there are virtually no borrowers in the low-income sample who take out jumbo loans. That is, almost all of the low-income borrowers who would locate anywhere to the right of the limit have chosen to bunch, which limits the possible magnitude of the estimated response. Consequently, the bunching borrowers may all be infra-marginal, and the bunching estimate cannot be

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

We also examined variation in bunching behavior and mortgage rate spreads across zip codes with varying incomes and price-to-income ratios. This evidence is generally consistent with the individual-level results, in that the behavioral response implied by bunching is monotonically and significantly increasing in zip code income quartile, from a minimum of 0.04 to a maximum of 0.07. We also found a monotonic relationship between bunching and zip code price-to-income quartile (with implied behavioral responses ranging from 0.05 to 0.09), but it is difficult to know whether this finding reflects unobserved characteristics of the borrowers—for example, higher non-housing wealth in high price areas—or a different mix of mortgage products available in different areas or to different people. In contrast, we found no stable evidence of heterogeneity in the jumbo-conforming spread, again suggesting that either preferences or adjustment costs drive the bunching. We also found that the fraction of borrowers who took out second mortgages was close to constant at about 35 percent, so these differences in bunching imply variation in both the first mortgage and total demand elasticities.

As discussed in the text, unlike Kleven and Waseem (2013), we cannot distinguish between the role of preferences and adjustment costs, because there is no dominated region in which no borrower, regardless of preferences, would prefer not to locate in the absence of adjustment costs. While it seems more likely that these differences are driven by adjustment costs than by differences in underlying preferences for first mortgage debt, we leave a full analysis of this issue for future research.

## A.5 More Cash or Cheaper Houses?

In section 6.3 we argue that roughly 65 percent of borrowers who bunch at the conforming limit do so without the use of a second mortgage. In order to lower their first mortgage amount, these borrowers must either be putting up more cash or spending less on housing than they otherwise would. If they are spending less on housing while holding their leverage

---

interpreted as an average marginal response, as in the theory.

roughly constant, then both their first and combined LTVs should be little changed relative to the counterfactual in which the conforming limit does not exist.<sup>8</sup> Therefore, if all of the bunching borrowers were either taking out second mortgages or buying cheaper houses, then we would expect the average combined LTV at the conforming limit to be about the same as in nearby bins. Figure A.3, which plots the combined LTV against the first mortgage amount, makes clear that this (admittedly extreme) scenario is far from true. The combined LTV at the limit is about 75 percent, well below the 80 to 85 percent that would be predicted based on the dashed line, which is a polynomial fit using the data outside of the same excluded region that was used to estimate bunching in figure 7A.<sup>9</sup> Thus, a significant portion of the 65 percent of borrowers who bunch without a second mortgage must be doing so by putting up more cash as opposed to spending less on housing.

In the other extreme case, all of these borrowers are putting up more cash. Since both putting up more cash and taking out a second mortgage reduce a borrower’s first-mortgage LTV while spending less on housing does not, we can gauge the plausibility of this extreme case by examining the relationship between *first-mortgage* LTV ratios and loan size near the limit. To do so, figure A.4 plots the first-mortgage LTV against the first mortgage amount. This figure is analogous to figure A.3 except for the bold “X”, which is a first-mortgage LTV calculated under the assumption that none of the bunching borrowers adjust their housing expenditure.

To calculate this LTV, we first generate a “counterfactual” mean house price for each loan size bin by fitting a 5th degree polynomial to the observed mean price in each bin omitting the bins in the excluded region used to estimate bunching. We then take the weighted average of these mean counterfactual house prices in the bin containing the limit and in each bin to the right of the limit in the excluded region. In calculating this average, the weight assigned to the mean price at the limit is the estimated counterfactual bin count from the bunching procedure, and the weights assigned to each of the mean prices to the right of the limit are equal to the difference between the counterfactual and observed bin counts from figure 7A. To calculate the LTV plotted in the figure, we then compute the average conforming limit

---

<sup>8</sup>Some borrowers may buy a cheaper house but target an ideal monthly payment, rather than an ideal LTV. Such borrowers would slightly reduce their LTV when they bunch, but by much less than buying the same house at the same price.

<sup>9</sup>The 85 to 90 percent combined LTVs to the right of the limit, which are higher than any of the other points on the plot, also stand out. One possibility is that the borrowers in these bins who do not bunch have different characteristics than those who do and are thus a selected sample. Indeed, these borrowers are clearly somewhat “abnormal,” in that they do not bunch despite the seemingly large gains from doing so. However, as we know from figure 7A, there are relatively few borrowers remaining in these bins.

for all loans observed at the limit and divide by the weighted average counterfactual house price.

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

Although these calculations suggest that there is no direct impact of the interest rate differential on house prices and the demand for housing itself, we do not necessarily want to draw that inference, for three reasons. First, the LTV calculations above are “back of the envelope” and there are several untested assumptions involved. Second, the interest rate differential at the limit is relatively small and may not be as informative about larger changes in rates over time. Finally, it is likely that other methods, in particular those used by Adelino et al. (2012) and Kung (2014), are better suited to studying the effects of rates on house prices. They estimate the elasticity of house prices to interest rates by following similar houses over time as they become more or less difficult to finance with a conforming loan at a constant LTV of 80 percent, essentially using a difference-in-difference approach rather than the bunching approach we follow. This approach has advantages over ours for this particular question, especially because we cannot observe borrowers who drop out of the market entirely (the extensive margin).

## Additional References

**Congressional Budget Office**, “Federal Subsidies and the Housing GSEs,” 2001.

**Jaffee, Dwight and John M. Quigley**, “The Future of the Government Sponsored Enterprises: The Role for Government in the U.S. Mortgage Market,” in Edward Glaeser and Todd Sinai, eds., *Housing and the Financial Crisis*, University of Chicago Press, 2012.

**Kung, Edward**, “The Effect of Credit Availability on House Prices: Evidence from the Economic Stimulus Act of 2008,” November 2014. Working Paper.

**Passmore, Wayne, Roger Sparks, and Jamie Ingpen**, “GSEs, Mortgage Rates, and the Long-Run Effects of Mortgage Securitization,” *Journal of Real Estate Finance and Economics*, 2002, *25* (2), 215–242.

– , **Shane M. Sherlund, and Gillian Burgess**, “The Effect of Housing Government Sponsored Enterprises on Mortgage Rates,” *Real Estate Economics*, 2005, *33* (3), 427–463.

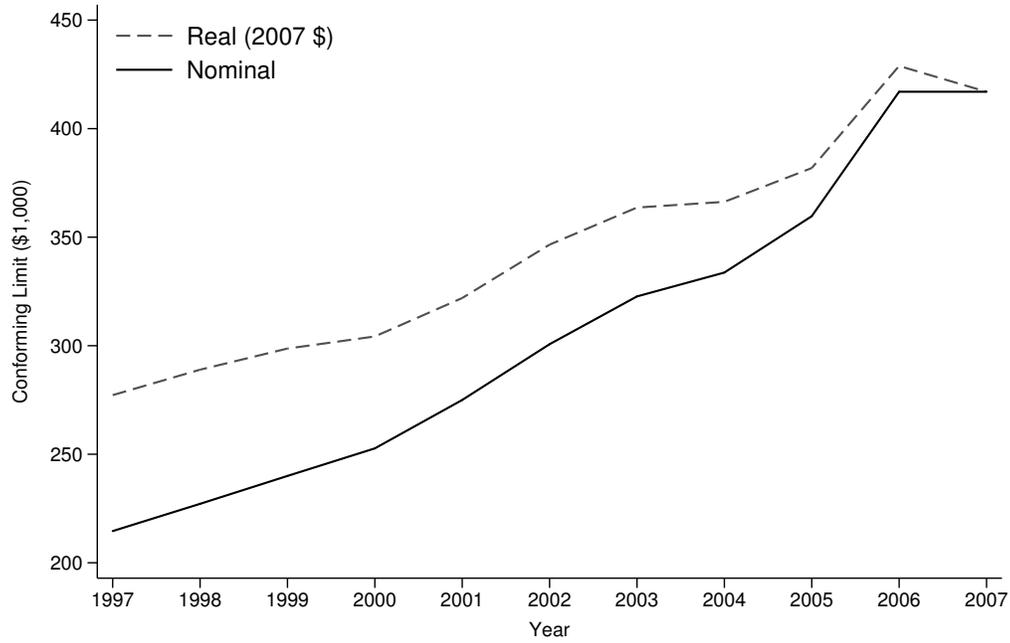
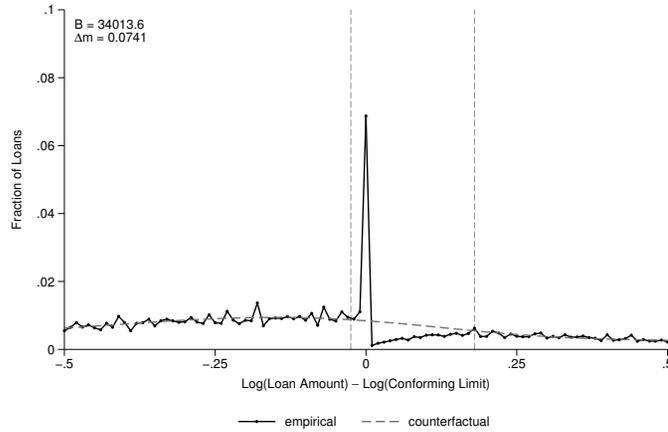
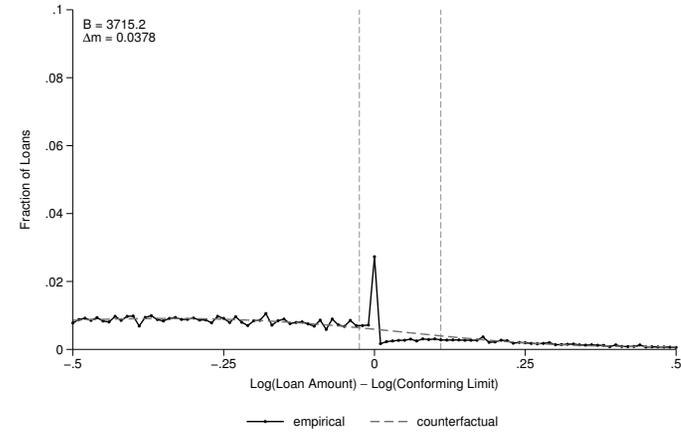


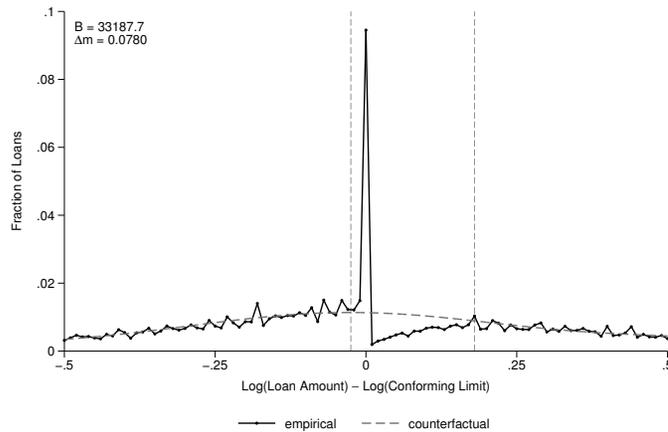
FIG. A.1.—Conforming Loan Limit Over Time. This figure plots the annual conforming loan limit for single family homes in nominal and real 2007 dollars for each year between 1997 and 2007. Historical conforming limits were taken from the Office of Federal Housing Enterprise Oversight’s (OFHEO) 2007 annual report to congress. Nominal dollars are inflated using the Consumer Price Index for all Urban Consumers (CPI-U).



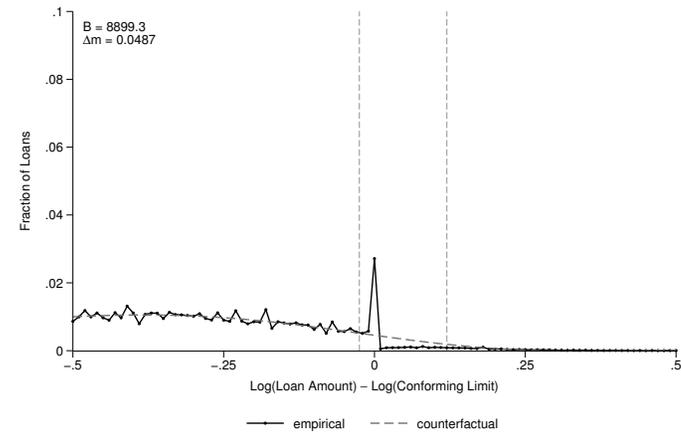
*Panel A. Non-Minority*



*Panel B. Minority*



*Panel C. High-Income*



*Panel D. Low-Income*

FIG. A.2.—Bunching at the Conforming Loan Limit by Borrower Type, Fixed Rate Mortgages Only. This figure plots the empirical and counterfactual density of (log) loan size relative to the conforming limit estimated separately for: (a) non-minority borrowers, (b) minority borrowers, (c) high-income borrowers and (d) low-income borrowers. The solid black line is the empirical density. Each dot represents the fraction of loans in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed gray line is the estimated counterfactual density obtained by fitting a 13th degree polynomial to the bin counts, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The figure also reports the estimated number of loans bunching at the limit ( $B$ ) and the average behavioral response among marginal bunching individuals ( $\Delta m$ ), calculated as described in section 4.1.

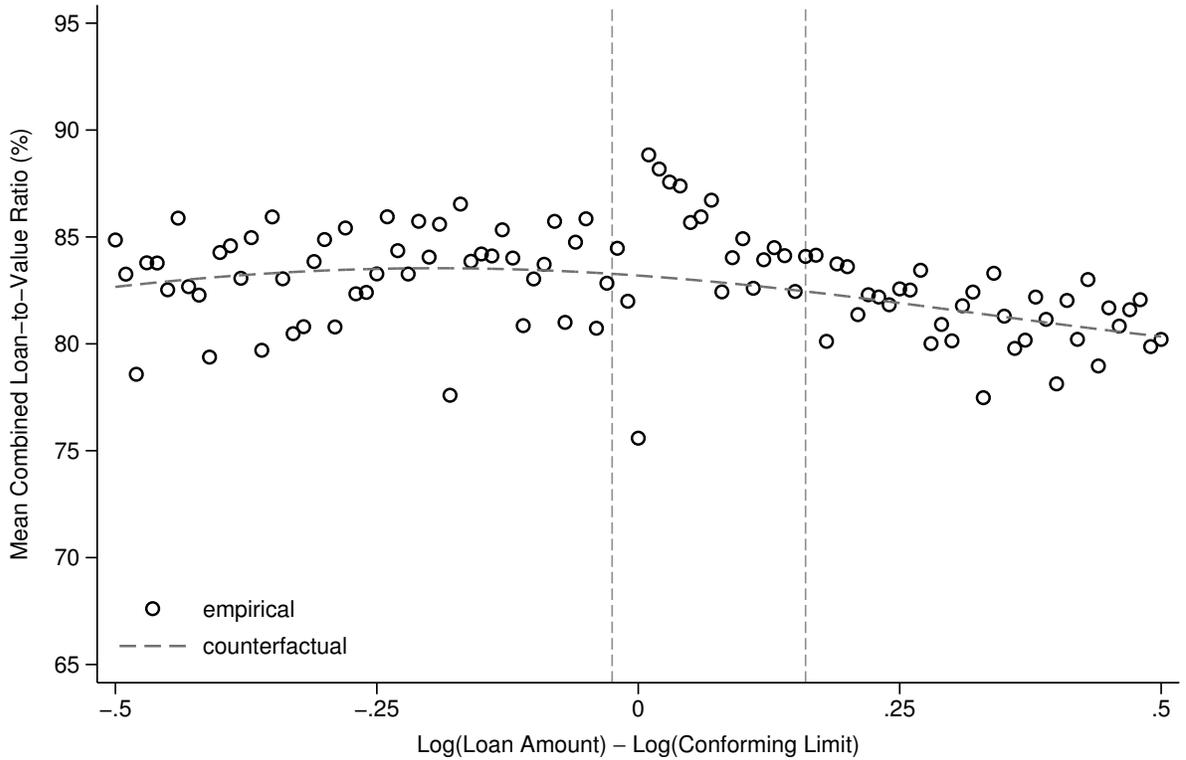


FIG. A.3.—Combined Loan-to-Value Ratio by First Mortgage Amount. This figure plots the average combined loan-to-value ratio (CLTV) as a function of the first loan amount relative to the conforming limit. Each dot represents the average CLTV in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed gray line is the counterfactual mean CLTV obtained by fitting a 5th degree polynomial to the bin averages, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The excluded region is the same region used to estimate bunching for the sample of fixed-rate mortgages. CLTV is calculated as the ratio of the sum of up to three mortgages used to finance a transaction to the recorded purchase price. Sample includes only transactions with a fixed-rate first mortgage.

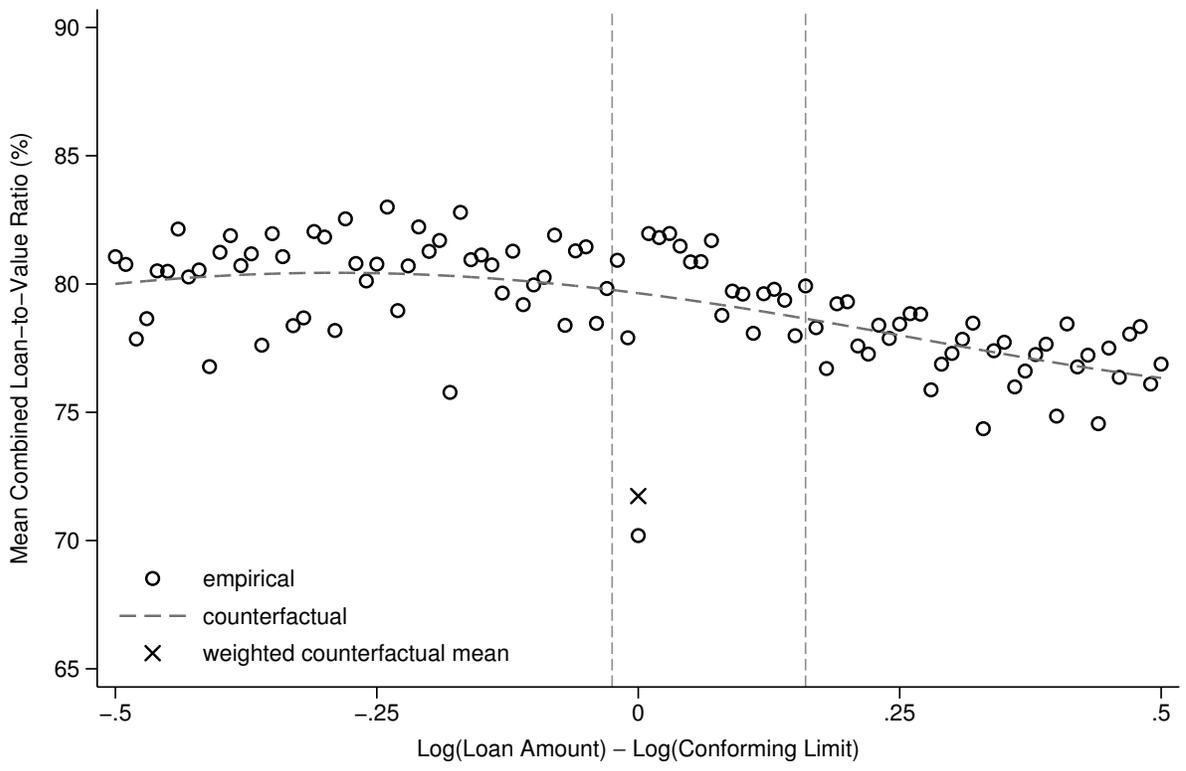


FIG. A.4.—First Mortgage Loan-to-Value Ratio by First Mortgage Amount. This figure plots the average first mortgage loan-to-value ratio (LTV) as a function of the first loan amount relative to the conforming limit. Each dot represents the average LTV in a given 1-percent bin relative to the limit in effect at the time of origination. The heavy dashed gray line is the counterfactual mean LTV obtained by fitting a 5th degree polynomial to the bin averages, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The excluded region is the same region used to estimate bunching for the sample of fixed-rate mortgages. The bold “X” is an LTV calculated assuming that borrowers who bunch at the limit do so without adjusting their housing price. See the text for the details of this calculation. LTV is calculated as the ratio of the first loan amount to the recorded purchase price. Sample includes only transactions with a fixed-rate first mortgage.

TABLE A.1  
SUMMARY STATISTICS, DATAQUICK SAMPLE

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

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

TABLE A.2  
SUMMARY STATISTICS, LPS SAMPLE

	FRMs		ARMs	
	(1) Full Sample	(2) Within \$50k of CLL	(3) Full Sample	(4) Within \$50k of CLL
Interest Rate (Initial for ARMs)	6.68 (0.92)	6.70 (0.91)	5.04 (2.13)	5.03 (2.09)
Jumbo	0.18 (0.38)	0.14 (0.35)	0.50 (0.50)	0.39 (0.49)
First Loan Amount (\$1,000)	281.56 (174.78)	322.96 (56.68)	453.52 (280.50)	377.15 (47.18)
Appraisal Amount (\$1,000)	401.54 (281.92)	447.75 (156.65)	612.49 (433.27)	496.99 (104.84)
Loan-to-Value Ratio	74.09 (16.34)	76.08 (13.61)	76.35 (10.01)	77.36 (9.14)
Debt-to-Income Ratio	35.77 (12.36)	36.88 (11.89)	35.64 (11.89)	36.59 (11.42)
Missing DTI Ratio	0.70 (0.46)	0.74 (0.44)	0.45 (0.50)	0.45 (0.50)
FICO Score	731.14 (51.74)	731.66 (49.94)	719.61 (52.76)	717.83 (52.50)
Missing FICO Score	0.33 (0.47)	0.35 (0.48)	0.20 (0.40)	0.22 (0.41)
Term (Months)	346.08 (52.46)	350.26 (46.38)	365.91 (30.75)	365.75 (29.07)
30-Year	0.90 (0.30)	0.93 (0.26)	0.93 (0.25)	0.94 (0.24)
Observations	1,011,753	257,349	926,272	221,203

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

TABLE A.3  
BUNCHING ESTIMATES BY BORROWER TYPE, FRMs ONLY

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

NOTE.—Each column reports the estimated number of loans bunching at the conforming limit ( $\hat{B}$ ), the average (log) shift in mortgage balance in response to the conforming limit among marginal bunching individuals ( $\Delta\hat{m}$ ), the excess mass at the conforming limit ( $\hat{B}/\sum_{j=L}^0 \hat{n}_j$ ), and the upper limit of the excluded region used in estimation ( $m_L$ ). Estimates are reported separately for high- and low-income borrowers and for minority and non-minority borrowers. High-income borrowers are those who report an income on their loan application that is higher than the median in the pooled sample. Low income borrowers are those below the median. Minority borrowers are those who identify as either black or Hispanic on their loan applications. Sample includes only transactions with a fixed-rate first mortgage which could be successfully matched to a mortgage application in the HMDA data and for which the borrower reported their income as well as both a race and an ethnicity. Standard errors (in parentheses) were calculated using the bootstrap procedure described in section 4.1.