The Harvard Joint Center for Housing Studies advances understanding of housing issues and informs policy through research, education, and public outreach.

Working Paper, March 2017

What is the Microelasticity of Mortgage Demand to Interest Rates?

Stephanie H. Lo
Ph.D. Candidate, Economics
Harvard University
Meyer Fellow 2016, JCHS

Abstract

What is the microelasticity of mortgage demand to interest rates? Despite the importance of this parameter for models of monetary policy efficacy, little is known about the intensive and extensive margins of mortgage demand to interest rates. I propose an identification strategy using novel microdata on mortgage rates. I exploit the fact that, due to regulatory factors, spreads in mortgage rates across borrowers exhibit a cutoff at certain FICO scores, and show using default and securitization data that a regression discontinuity design across mortgage pricing breakpoints isolates demand, not supply, margins. I show that the intensive and extensive margins of demand for mortgages are sensitive to interest rates and are economically large: a 25 basis point decrease in mortgage rates for high-FICO individuals is associated with a 50% increase in the likelihood of a potential borrower to demand a loan and an increase in loan size of approximately $15k, or approximately 10% of the average origination volume. I additionally find that for both the intensive and extensive margin, borrowers with high FICOs tend to be more sensitive to interest rate changes, elasticities are relatively constant over time, and the marginal responsiveness to interest rates is decreasing.

Direct correspondence about this paper to:
Stephanie H. Lo, PhD Candidate, Economics, Harvard University, shlo@fas.harvard.edu.

What is the Microelasticity of Mortgage Demand to Interest Rates? *

Stephanie H. Lo

January 19, 2017

Introduction

An important parameter for understanding the impact of macroeconomic policy is the elasticity of new mortgage borrowing to interest rates. Housing is a major component of the business cycle, and one channel of monetary policy transmission centers on the premise that decreases in interest rates will ultimately pass through to residential investment by decreasing the cost of mortgages and increasing the demand for housing. Yet, in the aftermath of the financial crisis, even as unconventional monetary policy put downward pressure on interest rates of varying maturities, the number of purchase mortgages hardly budged.

The measurement of the elasticity of mortgage demand to interest rates is not as straightforward as it may seem. Using macroeconomic data obscures the measurement of the mortgage elasticity since low interest rates tend to be driven by negative macroeconomic shocks, which in turn have large negative impacts on mortgage demand. Figure 1 shows the time series of the headline mortgage rate and the total purchase mortgage volume from 2000 to 2014. From 2008 onward, in the aftermath of the crisis, mortgage rates fell from 6 percent to 4 percent, yet mortgage originations also fell over this period. This is not surprising given that the financial crisis was accompanied by a macroeconomic slowdown and may have discouraged borrowers. Yet it

*shlo@post.harvard.edu. I thank Joshua Abel, Gabriel Chodorow-Reich, Benjamin Friedman, Andreas Fuster, Edward Glaeser, Nathan Hipsman, Greg Mankiw, Kenneth Rogoff, Paul Willen, and seminar participants at Harvard and the Joint Center for Housing Studies for helpful discussions on this project. I acknowledge financial support from the Harvard Graduate School of Arts and Sciences, the AEA CSWEP Summer Fellows program, and the John H. Meyer Fellowship from the Joint Center of Housing Studies at Harvard. Part of this work has been done at the Federal Reserve Bank of Boston; the views contained herein are those of the author and do not necessarily reflect those of the Board of Governors of the Federal Reserve System, its members, or its staff.

highlights the econometric challenge of identifying the true responsiveness of mortgage demand to interest rates, since estimation using broad macroeconomic data must make a number of structural assumptions for the impact of other macroeconomic factors, and in doing so may introduce a lot of uncertainty about the elasticity parameter estimate itself.

In this paper, I measure the mortgage microelasticity of demand to interest rates using a novel identification method that uses interest rate discontinuities across certain borrower credit scores. This empirical method measures the "local" or "micro" elasticity of mortgage demand–the responsiveness of borrowers to interest rates, holding all else constant, such as borrower wealth and house prices.[1] I show that, for my sample, these discontinuities in pricing are completely determined by regulation, namely Loan Level Pricing Adjustments (LLPAs), which cause breakpoints in mortgage pricing depending on credit scores and leverage, typically referred to in the mortgage context as the loan-to-value ratio (LTV). I use a novel proprietary dataset to derive the exact wholesale mortgage rates offered on a daily basis. I provide evidence that, for borrowers with the credit scores that I consider, the change in mortgage behavior across breakpoints is driven by mortgage demand rather than lender-driven supply.

I find large and statistically significant effects of interest rate changes on the demand for purchase mortgages. On the extensive margin, I find that a decrease in interest rates by 25 basis points results in an increase in the propensity to obtain a mortgage of about 50 percent. On the intensive margin, I find that the average borrower increases the amount of mortgage borrowing by approximately 10 percent for a 25 basis point decrease in interest rates. While the average estimate indicates the loan to value ratio increases as interest rates fall, for most of the estimates, a zero effect cannot be ruled out.

I find evidence of heterogeneity in the responsiveness to mortgage rates. Across credit scores, higher FICO borrowers seem to be more responsive, both by increasing selection into obtaining a mortgage and also by obtaining larger mortgages. Across mortgage rate changes, there appears to be concavity in borrower responsiveness, with a decreasing elasticity as the interest rate changes become larger.

The potential policy implications of my study are large. Purchase mortgages have fallen after the crisis despite decreases in mortgage rates. The total number of first-lien purchase mortgages was 2.74 million in 2012, a 44.4 percent decline since 2011 and a 54.5 percent decrease from

[1]The term "microelasticity" has been used to think about labor elasticity, where the micro elasticity is the partial equilibrium response and the macro elasticity is the general equilibrium response.

the peak volume of 2005.[2] This decrease in purchase mortgages has not translated one-for-one to lower sales activity; all-cash purchases have increased to partially offset the mortgage decline, and overall house sales have only decreased 20 percent from 2001 to 2012. Evidence points to a differential change in mortgage demand across FICO scores: Figure 2 shows the impact for high versus low credit scores. The percent of originations and refinances, both by count and by loan volume, rose for the strongest borrowers after the financial crisis. The fraction of loans going toward the best (720+ FICO) borrowers increased from 55 percent in early 2010 to 65 percent in 2015, while the share of loans from low-FICO borrowers (FICO 620-659) fell from over 18 percent in 2010 to less than 10 percent in 2015.

While some of the decrease in borrowing for the lowest-FICO (particularly subprime) individuals was likely driven by supply constraints, my results indicate that for higher FICO borrowers, the margin of adjustment was on the demand side.[3] Since my analysis focuses on relatively high credit score borrowers after the financial crisis, my estimates have direct implications for the efficacy of monetary policy after the crisis. The LLPAs essentially function as a credit surface, with borrowers facing different interest rates depending on their credit score and LTV. My estimates indicate that, holding fixed borrower characteristics, the responsiveness of mortgage borrowing to interest rates was relatively constant over the period, pointing to the role in regulatory-induced credit spreads in facilitating a decrease in overall purchase mortgage originations, particularly amongst low credit score borrowers.

My work contributes to many strands of literature. Empirical estimates of the elasticity of mortgage demand to interest rates are sparse. Glaeser and Shapiro (2003) investigate the elasticity of housing demand to interest rates by using state-level variation in the home mortgage interest deduction, but do not find a significant response in homeownership levels across states to the policy. Focusing on house prices rather than housing demand, Glaeser, Gottlieb, and Gyourko (2012) find that house prices are less responsive to interest rates than the standard pricing model used in housing market analysis would predict. Fuster and Zafar (2014) attempts to measure the sensitivity of housing demand using a survey that asks the respondents' willingness to pay under various financing conditions, including different mortgage rates. An increase in mortgage rates by two percentage points is found to change the willingness to pay for a home by only about five percent on

[2] Numbers from Goodman, Zhu, and George (2014).

[3] For low credit score borrowers, lenders may have been more sensitive to putback risk and therefore more cautious to make the mortgages at all.

average. The extensive margin choice of whether to purchase a home is not explored. Best et al. (2015) exploits quasi-experimental variation in interest rates due to notched mortgage contracts in the UK; that is, mortgage interest rates follow a step function of the loan-to-value ratio (LTV) at the time of loan originations. Examining bunching estimates at LTV breakpoints at time of refinancing (i.e. holding constant the purchased house), Best et al. find that the mortgage demand elasticity is about 0.3 on average and is strongly heterogeneous, in particular increasing in leverage. Best et al.'s study has important implications for the elasticity of intertemporal substitution–remortgagors are deciding how much consumption to give up now to lower interest payments in the future.

My paper is the first to study the impact of interest rates on purchase mortgage originations for a recent time period. The closest paper is DeFusco and Paciorek (2014), which uses bunching at an interest rate discontinuity at the jumbo-conforming spread to measure the elasticity of demand for pre-2007 loans.[4] In terms of broader trends in the elasticity of demand for loans to interest rates, Karlan and Zinman (2013) run an experiment in Mexico in which the researchers are able to exogenously impose lower interest rates.[5] By showing that there does not appear to be credit rationing for high-quality borrowers, my paper touches on themes in Li and Goodman (2014), and is consistent with the estimate in Anenberg et al. (2015) that credit supply was unchanged for high-FICO borrowers from about 2008 to 2015.

My paper also has important implications for the recent academic discussion of economic inequality after the crisis. Recent research has documented a fall in the number of purchase mortgages, alongside a rise in the average FICO score and average income of individuals acquiring purchase mortgages[6]

[4]My paper differs from DeFusco and Paciorek in several ways. First, I have direct pricing from lender rate sheets, whereas DeFusco and Paciorek must estimate the jumbo-conforming spread using rates for different borrowers and trying to condition on observables–a method which could be biased if unobservables drove the sorting and rates offered near the jumbo-conforming breakpoint. Second, I claim that borrowers just below and above the FICO thresholds are identical and hence interest rate variation from their perspective is exogenous, whereas variation in the jumbo-conforming spread could be endogenous. Third, one might believe the conforming loan limit is subject to supply thresholds, in the sense that lenders may be more likely to offer loans just below the threshold since these are considered less risky. Finally, my method has the benefit of estimating potentially heterogeneous elasticities across borrower types, time, and interest rate gaps.

[5]The researchers find that the price elasticity of demand for credit is quite elastic: outstanding loan balances and the number of loans each increase by more than 10 percent from the 10 percentage point reduction in the interest rate (on a base of roughly 100 percent APR). While their setting is obviously quite different–due in part to being situated in a developing country with less formal credit markets and higher baseline interest rates–the finding lends support that the extensive and intensive margins of borrowing increases can be quite "elastic", in the sense that the amount of credit demanded changes by a greater percentage than the percentage by which the price of credit changes when a shift in price occurs.

[6]See Swanson (2015) and Goodman, Zhu, and George (2014)

The structure of the paper is as follows. The first section discusses data and measurement, briefly introducing mortgage market mechanics as necessary. The second section outlines a simple model that describes how mortgage rates might be expected to respond to mortgage rates. The third section gives an overview of the specific regression discontinuity approach, discussing the baseline specification and multiple pieces of evidence suggesting that the approach is valid. The fourth section discusses the estimation results. The fifth section discusses empirical robustness. The sixth section discusses the economic implications of the estimations. The last section concludes.

Mortgage Background and Data

In this section, I give a brief overview of my regression discontinuity strategy. I provide a brief background of the mortgage market to give a sense for why this regression discontinuity design can be used, and give details on the underlying data and measurement.

Figure 3 shows an example of the regression discontinuity design at the heart of my paper. The plot shows the "mortgage propensity"–defined as the number of mortgages originated per individual in the population–per FICO score for a week in 2009 (January 29 - February 5) against the rate spread. The mortgage rate discretely jumps from almost 5.9 percent for FICO 719 to 5.4 percent for FICO 720. Linear fits for the mortgage propensity over the relevant ranges (700-719 and 720-739) are shown. The mortgage propensity increases as FICO increases, and the graph shows a discrete jump upward at FICO 720, just where the mortgage rate falls. In my empirical exercise, I would calculate the extensive elasticity by taking the jump size (approximately 20 mortgages per 10,000 individuals) divided by the rate spread (approximately 50bp). I then take the average of this estimate across all weeks in my sample.

The time period studied is October 2008 to December 2014. I study conforming purchase mortgages and restrict the baseline analysis to first-lien mortgages.[7] For my baseline regression discontinuity analysis, I construct a table with the count of the potential borrower population over time (from Equifax tables reflecting the population per credit score), the count of conventional mortgages originated, mortgage rates, origination amounts, and appraisal amounts, by week for each FICO score. For robustness, I create a similar table for FHA mortgages, as well as separate

[7] From October 2008 onward, second liens were only 27.8k of the 14.4 MM purchase mortgages made (including non-conforming loans, such as FHA). Because the second liens play such a small role in the sample, including them causes virtually no change in the results.

tables for default and securitization trends.

I discuss the details underlying each of these data in the following subsections.

Mortgage Rates

One of the most distinctive data sources used in this project is lender ratesheet data, from a vendor called LoanSifter (now part of Optimal Blue), available from October 2008 onward. This is a rare dataset, accessed through the Federal Reserve, that reflects mortgage rates being offered on the primary mortgage market conditional on borrower characteristics. That is, this database allows the user to pose as a loan officer, inputting desired loan size amount, loan-to-value (LTV), debt-to-income (DTI), MSA, and credit score. The database then searches through a collection of lender-uploaded rate sheets (typically updated at least once daily) and finds a menu of rate/point combinations to offer to the borrower. The data is collected from an actual software platform that loan officers use to search for mortgage rates, so misreporting is not an issue.

Rate sheets offer several combinations of points ("Yield spread premiums", or YSPs; also known as the Service Release Premium or negative discount points) and rates and reflect the willingness-to-pay of an investor for a given mortgage. The yield spread premium reflects the amount, as a percentage-point of the loan amount ("point"), transacted upon closing the loan, where 100 reflects no additional payment or rebate. YSPs above 100 reflect payments from the investor, and are often split equally toward the loan officer's commission and the borrower's closing costs on the loan. Lower YSPs correspond to lower mortgage rates and reflect that the borrower must compensate the investor for the lower cash flow.

In the estimates throughout this paper, I hold borrower characteristics (except for FICO) constant and YSP constant at 0. Summary statistics for the mortgage rates utilized are shown in Table 1.

The ability to access rate/point combinations is important for a few reasons.[8] First, perhaps contrary to popular belief, there is no single mortgage rate, even conditional on all borrower characteristics. Rather, the borrower has the option to pay points upfront, quoted as a percentage of the loan amount, to lower the ongoing rate; conversely, borrowers may actually choose to pay "negative" points to help cover the downpayment and closing costs in exchange for a higher mort-

[8] Fuster, Lo, and Willen (2017) discusses other important implications of using point-normalized mortgage rates, such as correctly evaluating the passthrough from mortgage backed securities prices to the effective prices of mortgages that borrowers see.

gage rate. Second, most datasets used in the academic literature only have the mortgage rate, but do not contain any information on points, and hence may misrepresent the actual trend in mortgage costs. Third, many papers try to control for pricing on characteristics by backing out the relationship of mortgage rates and borrower parameters such as the LTV and credit score, which is imperfect with a small sample size. [9] In contrast, I can input these parameters directly.

LLPAs

Over my sample, mortgage rates have only varied across borrowers due to regulatory Loan Level Price Adjustments (LLPAs). Even though default risk is insured against in securitized loans, pricing of mortgages across FICO scores has historically varied as lenders offered lower rates to higher credit score borrowers.[10] There was no systematic premium for having a low credit score until November 6, 2007, when the FHFA announced the implementation of loan-level price adjustments (LLPAs), applicable to all Fannie/Freddie loans, which over this period accounted for approximately 80 percent of all mortgages.[11] LLPAs were issued as additional fees, paid upfront by the lender to Fannie/Freddie, to compensate the perceived additional risk imposed in mortgages. LLPAs increase in leverage (LTV) and FICO, with discrete breakpoints that incentivize remaining just below certain LTV cutoffs. A brief history of LLPA changes is shown in Table 2.

I find that LLPAs, once instituted, completely determine mortgage spreads. By matching my proprietary rate sheet data with the time series of LLPAs, I test whether the wholesale mortgage rates include additional "overlays"–premiums charged to individuals with different credit scores. Even though some lenders may price differentially, on average, the gap between (say) a FICO 680 and FICO 740 loan, all else equal, is exactly equal to the LLPA charged by the GSEs.

A key implication of the fact that LLPAs explain the exact spread between borrowers with different FICO scores is that the actual mortgage rate obtained by the mortgage investor is constant across the cutoff, even though the mortgage rates that the borrowers face vary, with the LLPA "wedge" between the lender and borrower rates paid directly to the GSEs. Hence, lenders are acting optimally, in the sense that they charge the same mortgage rate to virtually identical borrowers

[9]See, for instance, the jumbo-conforming spread estimates in Sherlund (2008)

[10]One potential explanation for this is that lenders prefer to keep safer loans on portfolio (rather than securitize them) and were willing to pay more–or offer a lower rate–to attract the higher credit score mortgages. Also, servicing could be more profitable on higher credit score investors, who are less likely to default and therefore require less costly action on the part of the servicer.

[11]The announcement is available online; see Fannie Mae (2007)

with virtually identical default and putback risk.

The rate differences across FICO scores induced by LLPAs vary over time and depend both on the LLPA matrix and prepayment risk.[12] The rate spreads were the largest in the early part of the sample (October 2008 - October 2009), with spreads between FICO 680-700 and FICO 700-720 hovering above 40 basis points at times.

FICO scores

FICO scores are meant as a ranking of borrower credit quality, made by a company formerly called the Fair Isaac Corporation (and now simply called FICO). The possible range runs from 350 to 850, with a low score signaling a low credit quality borrower. Typically lenders call subprime borrowers those who fall below FICO 620.

There is a version of FICO score available at all three credit bureaus, Equifax, TransUnion and Experian, which use similar inputs but different models to determine credit scores. Mortgage lenders typically pull all three credit scores and use the median as the score at which to price a mortgage.[13] The exact FICO score recorded at origination is available in the McDash LLC data, and I use this value in the regression discontinuity to determine which mortgage rates borrowers were quoted. [14]

I have three major credit scores via the NY Fed CCP/Equifax and Equifax Credit Risk Insights Database (CRISM). Table 3 shows the variation of credit scores for borrowers in the sample. The mean and median credit score across the different metrics is similar. The summary statistics, which combine both cross-sectional and panel variation, also indicate how FICO scores vary quite a lot, with a standard deviation of about 50 credit score points. Much of this variation does exist for the same borrower across time; as shown in Figure 4, even in the space of 6 months, FICO

[12] As a back of the envelope, consider that the LLPA for a FICO 700 borrower with LTV 80 on April 30, 2011 was 1 percentage point of the loan amount. The median loan amount for my sample is approximately $200k, which would correspond to an LLPA of $2k. While this is an upfront payment paid by the lender, my ratesheet data indicates that lenders price the entire LLPA into mortgage rates by increasing mortgage rates to offset this fee. Exactly how much mortgage rates need to rise depends on projected prepayment rates, which can be valued through the mortgage backed securities market.

[13] In the case of joint loans, meaning two borrowers jointly applying for a mortgage, it is typically the lower of the two median credit scores, known as the “minimum FICO rule”. Joint applications for mortgages may help alleviate borrowing limits by documenting extra income, but may come at the cost of increasing the overall mortgage rate if the FICO scores fall in different bins.

[14] For the purpose of comparing the Equifax credit bureau data to loan origination data, I can use the merged CRISM data to create a mapping from the Equifax credit score to the FICO score. For the relevant range of credit scores, the mapping is actually one-for-one with a small positive adjustment.

scores can vary substantially. While the credit score formula changes over time and remains a closely guarded secret, discrete events such as defaults or the opening of new tradelines may be drivers of large changes in credit scores from month to month.

Mortgage Activity

My mortgage data comes from McDash LLC. The data cover approximately 70 percent of all mortgage originations in the United States. The data include characteristics of the mortgage (fixed or adjustable rate; the term; jumbo; conventional vs. FHA, etc.); relevant dates (origination, first payment, first appearance in data);and further loan-level data such as the loan-to-value ratio, origination amount; and outcome data such as foreclosure dates. The data also include unique individual IDs, which allows the mortgages to be linked to individuals in the New York Fed Consumer Credit Panel (discussed below), and unique loan IDs, which allows loans to be tracked over time.

I focus on purchase mortgages. I restrict to 30-year fixed rate mortgages, which are by far the most common share of mortgages.[15] I use data for only the first-lien mortgage and consider only data that has recorded FICO at origination. I also aggregate the data to a weekly level. [16]

While the detailed loan data is available by origination date, the relevant decision date for borrowers is the application date, which happens well before the loan is approved and originated. Unfortunately, due to data limitations, the exact lag between application and origination date of any given loan is unknown. I therefore use the average lag across approved loans, which is approximately 35 days (5 weeks) in my sample. The lag across time varies slightly but is dwarfed by the heterogeneity within a day even for borrowers with similar characteristics. While this assumption adds noise to my exercise, it should not introduce any bias.

The mortgage data are linked on an individual level to each individual's credit report data, which allows for some demographic data (age, geography, household size) and the tracking of all credit trends of individuals who have a mortgage, on a monthly level, starting six-months before origination and ending six months after the loan is removed from the data set due to termination.

[15] 15-year mortgages and adjustable-rate mortgages take some share of the market. The rates on across different loan terms (durations) tend to move together. Previous literature has examined borrower choice of ARM versus FRM, most of which, to my knowledge, is not focused on a price mechanism.

[16] Loan officers tend to record rounded dates (typically the first and last of each month). I drop these observations, which causes some noise but should not result in any bias. Monthly robustness checks, which include all data points, confirm the estimated magnitudes are approximately the same. I use a weekly baseline since this gives me more granular mortgage rate spread measurement.

The data on default and securitization trends comes this from mortgage servicing data. The data track the status of the mortgage for each month after origination. I consider a loan securitized if it is securitized within 36 months of origination. I consider a loan to have defaulted if it is ever 60 days or more delinquent within 36 months of origination. Given that the data runs through 2014, the restriction to examining the first 36 months helps limit the potential bias resulting from the fact the last years of the data haven't yet existed for 36 months and are therefore censored.

Potential Borrowers and Credit Trends

Credit trends and other individual-level data come from one of the major credit reporting agencies, the credit bureau Equifax, via the New York Fed Consumer Credit Panel. The data contain a random 5 percent of the U.S. population with a Social Security Number. The Equifax data contain useful individual-level credit trends such as credit score, debt balances, debt payments, number of accounts, and a few demographic data such as zipcode and age of borrower. The Equifax data is available quarterly.

FICO distributions are "smooth", meaning that there is no discontinuity between the number of mortgage borrowers from one FICO score to the next, so comparing the number of loans across thresholds gives a good sense for whether a true breakpoint exists at that cutoff. Still, to facilitate comparison across thresholds, I normalize the number of loans received by the number of "potential borrowers" for the relevant credit score. Since the NY Fed CCP is a 5 percent sample of the population, I scale this up appropriately. I then aggregate the number by credit score. The Equifax data only contains "riskscore", which is not equivalent to FICO score, but historical regressions indicate that a linear offset can approximately correct for this, which I do.[17]

Figure 5 shows the distribution of total individuals in the sample in 2008 and 2013. The distribution of individuals over credit score changes over time in the sample. While the distribution of borrowers has remained roughly constant over time, in 2013, fewer borrowers were at the very low end of the distribution, and instead more mass was concentrated at good credit scores (around 700).

I refer throughout the paper to "mortgage propensity", which refers to the fraction of the population that actually originates a purchase mortgage in any given week. (While one could theoretically

[17]The McDash LLC mortgage data has the true FICO score recorded at time of mortgage origination. This exact measure is crucial to study the response of borrowers to the breakpoint in mortgage pricing. The approximation of the credit score is only relevant for the denominator of the mortgage propensity measure and mortgage shopping sections.

use the number of mortgage applicants as the denominator, this suffers from selection biases).[18] Over the breakpoints from 620-760, the average mortgage propensity in my sample is 0.00008, or 8 in 100,000. The higher FICO scores have much higher propensities, with the best borrowers in my baseline estimate group (740) having a mortgage propensity of approximately 0.00012 (12 in 100,000), which is an order of magnitude larger than the worst borrowers I consider (0.00001, or 1 in 100,000). [19]

Theory

A Basic Model

For this analysis I focus on the consumer demand for mortgages. For ease of analysis and to isolate the effects that are important for my purposes, I assume that all consumers have already chosen to purchase a house worth p. Hence my analysis focuses on the decision of how much to borrow (origination amount), which then determines the loan to value ratio. The extensive margin can be interpreted as borrowers deciding to pay cash rather than borrow mortgage funds.

The model follows that of Brueckner (1994) closely, except that I consider the payment ("debt-to-income") constraint binding rather than the loan-to-value constraint. The difference between the two constraints is essentially one of upfront collateral constraints (if the downpayment is prohibitive) or a longer-term affordability constraint. The rationale behind the choice of the payment, rather than upfront collateral constraint, is because we are studying breakpoints in conforming loans, which already have quite large downpayments (20 percent); individuals who are particularly collateral constrained are likely to pursue FHA loans, which require a lower down payment constraint.

Consider a two-period model, which is easily extended to multiple periods. Utility is a function of housing consumption h and the numeraire nonhousing good x. Decisions made in the current

[18] The reason that mortgage applicants are not considered as the "eligible borrower" pool is that there seems to be selection into formally applying for a mortgage. Loan officers may discourage a borrower from submitting a formal application if they think she will be rejected. This is consistent with the fact that rejection rates were actually highest in the housing boom, and fell in the aftermath. Moreover, individuals with a lot of cash on hand may choose not to apply for a mortgage, but may be willing to switch to mortgage financing if mortgage rates are sufficiently low.

[19] To give a sense that this mortgage propensity is reasonable, note that approximately 1.3 MM conforming purchase mortgages are given out per year. Considering the main mass of borrowers are between FICO 660-780, this means that there are about 10,000 mortgages per FICO score. Dividing by the 400k potential borrowers per FICO score, this is a 0.025 mortgage propensity per year, or 0.0005 per week. This aligns with Table 4.

period affect future wealth z, and the discounted value of future utility is given by $\delta V(z)$, where V is a strictly concave function and $\delta < 1$ is the discount factor.

The consumer enters the current period with wealth w, which is the sum of current income and assets and is taken as exogenous. w can be allocated toward housing consumption, non-housing consumption, or saving. Current wealth cannot be supplemented by borrowing against future income, but consumers may use a mortgage to make a house purchase. Let m denote the size of the mortgage, s denote the amount of saving, and p denote the purchase price per unit of housing. Denote $\alpha < 1$ be the maximum mortgage loan-to-value ratio, and suppose the maximum payment-to-income fraction is given by β.

The problem faced by the borrower is as follows:

$$\max U(x, h) + \delta V(z)$$

subject to the current-period budget constraint:

$$x = w - s - (ph - m)$$

and the constraints:

$$s \geq 0 \tag{1}$$

$$\frac{\beta y}{r_m} \geq m \tag{2}$$

$$m \geq 0 \tag{3}$$

Equation 1 is the liquidity constraint. Equation 2 is the payment-to-income constraint, which states that that mortgage payment as a percentage of income falls below some critical value, here β.[20] Equation 3 restricts mortgage borrowing to be positive.

We assume the mortgage rate (r_m) and the interest rates earned on savings (r_s) are non-stochastic. We assume future income y is known, house prices h are constant across periods,

[20] Fannie Mae currently restricts debt-to-income ratios to be 38 percent for manually underwritten loans or up to 45 percent for borrowers who hit very specific requirements. Given the graph of DTI in Figure 5, it's clear that most borrowers hover exactly at this DTI limit of 36 percent, with a slight decrease as credit scores get larger. See https://www.fanniemae.com/content/guide/selling/b3/6/02.html

and future wealth is given by

$$z = y + (1 + r_s)s + ph - (1 + r_m)m \tag{4}$$

Substituting Equation 4 into V and substituting in the current-period budget constraint, the consumer's objective function is given by:

$$\max_{h,s,m} U[w - s - (ph - m), h] + \delta V[y + (1 + r_s)s + ph - (1 + r_m)m] \tag{5}$$

subject to the constraints (1)-(3).

Letting the multipliers for these constraints to be denoted by λ, μ, and θ, the Kuhn-Tucker optimality conditions for the problem are given by:

$$s : -U_x + (1 + r_s)\delta V' + \lambda = 0 \tag{6}$$

$$m : U_x - (1 + r_m)\delta V' - \mu + \theta = 0 \tag{7}$$

$$h : -pU_x + U_h + p\delta V' = 0 \tag{8}$$

along with the constraints (1)-(3) and the additional conditions:

$$\lambda \geq 0, \lambda s = 0 \tag{9}$$

$$\mu \geq 0, \mu\frac{\beta y}{r_m} - m = 0 \tag{10}$$

$$\theta \geq 0, \theta m = 0 \tag{11}$$

Then the optimal house size is governed by :

$$U_h/U_x = p[1 - \delta V'/U_x] = 0$$

and the choice between savings and mortgage is given by:

$$(1 + r_s)\delta V' + \lambda = (1 + r_m)\delta V' + \mu - \theta \tag{12}$$

The implications of these equations depend on the relative magnitudes of r_s and r_m. Brueckner

(1994) argues that the case when $r_s > r_m$ is most representative of the U.S. economy, i.e. that pretax investment returns typically exceed the (pretax) mortgage rate. When this is the case, then Equation 12 requires that

$$\lambda < \mu - \theta$$

This constraint implies that we cannot have both $s > 0$ and $\beta y/r_m > m > 0$, since this would require $\lambda = \mu = \theta = 0$, violating the above. Then either constraint (2) or (3) needs to bind. Consider $m = 0$; then $\theta > 0$ and $\mu = 0$, which could not be possible since this implies $\lambda < 0$. Hence, it must be the case that $m = \beta y/r_m$. This implies that, holding all else constant, an increase in r_m (for instance, as in across our FICO breakpoint) results in a discretely lower level of mortgage borrowing m. Because we take h as given, this implies the LTV is also discretely lower under the FICO breakpoint than above.

Elasticity Estimates and Discussion

Variable Discontinuities Across FICO Score

The best check for the validity of the regression discontinuity design is a visual one. Figure 6 shows the relevant outcome variables plotted against the FICO score for 2009. There are a few interesting trends to notice. First, mortgage rates do vary quite a lot across FICO scores, with low FICO borrowers obtaining mortgage rates (adjusted to no points) of over 7 percent in 2009, while mortgage rates were under 5 percent for the borrowers with the best FICO scores. To control for potential supply effects that may be present in low-FICO borrowers, I focus on FICO 680 and above. While the rate variation here is smaller–about 25 basis points across the 719/720 cutoff–I argue that the supply constraints across this threshold are negligible (see upcoming subsection "Testing for Supply Side Factors: RD Tests on Defaults and Securitizations"). Both the count of originations and the origination percent of potential borrowers exhibit striking cutoffs at FICO breakpoints, including the 720/740 cutoff (and less so, due to the small mortgage rate variation, the 739/740 cutoff) that is relevant to our analysis.

The characteristics of the loans vary discretely across the threshold as well. Origination amounts and loan-to-value ratios jump at each FICO threshold, becoming higher just above the breakpoint. This is consistent with a lower cost of borrowing makes larger loans more affordable. Appraisal amounts also jump at breakpoints, indicating that a lower cost of borrowing may induce an income effect (relatively richer, so more debt) in addition to the aforementioned substitution effect (debt relatively cheaper than consumption).

Table 4 shows the baseline propensity to get a mortgage, origination amount, LTV, and debt-to-income ratio. All numbers are per-week and reflect averages for the *exact* FICO score shown. The number of mortgages averages 20-50 for FICOs 680 to 760, with the number increasing as FICO increases. The propensity to get a mortgage increases with FICO score, starting from 0.000014 (approximately 1 in 100,000) for FICO 620 and increasing to 0.000189 (approximately 19 in 100,000) for FICO 760. Origination amounts and appraisal amounts of mortgages also increase with FICO. LTV is roughly constant across all FICO scores, and the debt to income ratio tends to fall with FICO score.

Baseline Specification

I run, for cutoffs $\Gamma \in \{660, 680, 700, 720\}$:

$$Y_i = \beta_0 + \beta_1 1_{FICO_i \geq \Gamma} + f(FICO_i) + 1_{FICO_i \geq \Gamma} \times g(FICO_i) + \varepsilon_i \tag{13}$$

where i indexes individual loans, the dependent variable Y_i indicates whether loan i is one of three outcome variables: a) the percent of the "potential borrower" population that obtained a loan (extensive margin); b) the origination amount of the mortgage (intensive margin); or c) the appraisal amount of the house being purchased; and both $f(FICO_i)$ and $g(FICO_i)$ are local polynomial regressions.

The RD is run using the rectangular kernel, but the results do not change qualitatively for different choices of specification.[21] Recall each FICO breakpoint is 20 points away (i.e. *ceterus parabis*, the same mortgage rate is given for all borrowers in FICO 680-699, and a different score is given to borrowers with origination FICOs between 700-719). To avoid capturing multiple breakpoints in any given regression, I only use data for loans with FICO scores $+/-19$ points from the FICO breakpoint (e.g. 701 to 739 for FICO threshold 720). The baseline range is chosen to be $+/-9$ from the FICO breakpoint, but to show robustness of the results to the choice in FICO range, I show also the result for the full potential range $(+/-19)$ and half of the baseline range $(+/-5)$.

[21] This is as recommended by Imbens and Wooldridge (2007), which notes that more sophisticated kernels only make a difference when the results are not credible anyway since the sensitivity of the kernel likely implies too much sensitivity to the choice of bandwidth. Gelman and Imbens (2014) discusses potentially undesirable features of using higher-order polynomials in RD and instead recommend a linear or quadratic.

Standard errors are run by bootstrapping 10000 times, with clustering on the monthly level. All baseline results are run restricting to rate spreads of at least 5 basis points; since elasticities are measured with rate spreads in the denominator, this prevents the elasticity measure from becoming unreasonably large due to random variation in low-spread weeks. Untabulated results indicate that the results hold even without restrictions on rate spreads.

For each week of the data, I run this regression and collect a sequence of β_1. I then divide by the relevant difference in mortgage rates across the threshold to derive the semi-elasticity of mortgage demand. The estimation is run using interest rate *level* changes, i.e. a one percentage point change from 5 percent to 4 percent, rather than in percent changes (where a 1 percent change would correspond to a change in mortgage rates from 5 percent to 4.95 percent).[22]

Elasticity Measurements

Table 5 shows the base case regression discontinuity results, which constistute the main contribution of this paper. For the FICO cutoffs 660, 680, 700, and 720, the regression discontinuity is run for a baseline range of +/- 10 FICO points. The loan amount, appraisal amount, loan-to-value ratio, and mortgage propensity (extensive margin) results are shown, alongside the "base" level of those variables just above each discontinuity. [23]

The first two rows show the intensive elasticity of demand, in units of dollars per percentage point decrease in mortgage rate. The results indicate that the intensive margin response is large; origination amounts increase about $52k-$79k for FICO 680 and 700 borrowers for each percentage point change in mortgage rate. This is relative to a base amount of about $200k, amounting to a 25 percent increase in the mortgage amount per 100 basis point change in interest rates.

The third row shows that the overall impact on the loan-to-value ratio is often statistically indistinguishable from zero. The only case when this is not true is for the cutoff 700, where the positive jump in the loan-to-value ratio is significantly different from zero. Mechanically, another way to think of this result is that the origination amount increased statistically significantly more than the

[22] the toy model presented earlier in this paper implies we should be thinking in percent terms, although the elasticity would vary by the level of interest rates... ($\partial m / \partial r_m = -\beta y(r_m)^{-}2$)

[23] Later robustness tests indicate that the results for FICO 700 are likely the most reliable since there is no concern about supply-side constraints. While the other FICOs shown pass most tests, such as continuity in default rates across cutoffs, there is a small discontinuity in the probability of securitization for FICO 660 and FICO 680. This does not seem to induce differential screening, since the probability of default across these cutoffs is smooth (economically or statistically indistinguishable from 0). Because of this, I argue the estimates for all of these cutoffs are demand-, rather than supply-, driven.

appraisal amount rose across the threshold, per basis point of interest rate spread.

The last row shows the extensive margin estimate over the entire sample period, defined as the break in the propensity to get a loan divided by the absolute rate spread on a weekly level. The elasticity for FICO 700 borrowers is greater than that of FICO 680 borrowers, indicating that higher FICO borrowers are more sensitive to interest rates on the extensive margin. The results for FICO 720 are noisy, likely because the rate spread was historically low, and at times zero. The magnitudes suggest that, for a group of FICO 700 borrowers facing a 25 basis point drop in interest rates, the increase in individuals getting a mortgage at all will be about 9.5 per 100,000 potential borrowers.

These magnitudes are economically large, in light of the baseline propensity to obtain a mortgage as highlighted in Table 4. The estimates indicate that a 25 basis point decrease in interest rates induces a FICO 700 borrower to be 50 percent more likely to demand a loan, and a 720 borrower 75 percent more likely to demand a loan.

Across FICO scores, the magnitudes are generally increasing: higher FICO borrowers seem *more* responsive to interest rates. This is true on both the intensive and extensive margins of adjustment. The most straightforward interpretation is that higher-FICO individuals may be less liquidity-constrained, so that part of the extensive margin response may be individuals who would have otherwise bought the property in cash.

Concavity with Respect to Mortgage Rate Changes

Does the response of borrowers to mortgage rate changes seem to be linear, in that a 100 basis point change in mortgage rate induces four times the response relative to a 25 basis point change in mortgage rate? The answer depends on exactly how consumer mortgage demand is modeled. If there were a large fixed cost to obtaining a mortgage, one might believe the response to a large mortgage rate change would be larger than a response to a small mortgage rate change. If the results are instead driven by discrete switching of borrowers into mortgages from cash when the mortgage rate falls below a certain point (for instance, if borrowers were willing to get a mortgage at 4 percent but prefer to pay cash when the mortgage rate hits 4.25 percent), then additional decreases in the mortgage rate may not induce too much extra demand.

To test the shape of the response to mortgage rate changes of different magnitudes, for each FICO, I estimate separate RDs for mortgage rate changes that are 5-20bp, 20-40bp, and 40-60bp. Figures 7 and 8 show the results. The graphs show that the largest change in demand per interest rate spread, both on the intensive and extensive margin, is induced by small (5-20bp) changes in interest rates. While the *per-basis-point scaled responsiveness* is decreasing in interest rate gaps, the aggregate responsiveness is increasing, as expected.

Elasticity Over Time

Table 6 shows the results for running the regression discontinuity over two-year increments. Broadly, we cannot reject that the intensive responsiveness of borrowers to interest rates was constant over the period. On the extensive margin, the estimates are increasing over time, with the responsiveness in 2013-2014 statistically and economically significantly larger than the estimates for 2008-2010. This may be due to a changing macroeconomic environment, i.e. a change in housing sentiment, that is not captured in the RD exercise in this paper.

Potential Channels

The observed magnitudes suggest that as debt becomes cheaper, individuals substitute away from other forms of saving or consumption and spend more on mortgages. This can be demonstrated using a simple heuristic: suppose for a moment that agents held fixed their desired mortgage payment and simply shifted the size of the house appraisal rather in response to interest rate changes. This approximation could proxy for a scenario in which individuals are credit constrained to only spend a fixed amount of money per period going forward.

As a back of the envelope for this rule of thumb, consider a $200,000 30-year fixed rate mortgage at an interest rate of 5 percent. The monthly payment is $1074. Decreasing the interest rate to 4 percent leads to a monthly payment of $955, a $119 (or 11 percent) savings. Or, over the life of the loan, the borrower saves $42,773.

Empirically, borrowers increase their loan size and their appraisal amount by about the same amount, so that the LTV remains roughly constant but they are able to borrow more for a more expensive house. For FICO 700 borrowers and a $200k origination amount and $290k appraisal amount, this intensive margin adds up to approximately $80,000 higher origination and $116,000 larger appraisal amount. If borrowers had instead desired to keep their interest payment constant at $1074, the origination amount would instead have increased to $225,000. This implies borrowers are willing to borrow more at lower interest rates.

Since I measure local elasticities using cross-sectional variation, I abstract from general equilibrium effects. My exercise is meant to isolate the substitution effect from consumption or non-mortgage saving to mortgage debt in the face of relatively cheaper mortgage debt, all else equal. There is still a Hicksian elasticity (income effect) in play, however. In reality, mortgage rates often fall due to easy monetary policy, which has further implications on the stock market, house prices, and overall income.

Alternative Explanations

One might worry that the results are due to lender supply rationing rather than price-rationing. Given that securitizations and default rates are consistent across breakpoints, this concern would be consistent with lender screening on default rates and tolerating a set threshold for defaults that is consistent across thresholds. One could then interpret the jump in the extensive margin as reflecting the fact that fewer low-FICO borrowers qualify for the default cutoff. The only reason to believe this might be discrete is that the higher interest rate causes these borrowers to default more. Similarly, one could interpret the intensive margin jump as lenders being more tolerant of higher origination volumes from higher-credit score borrowers.

My results are inconsistent with these possible concerns. Recall that the empirical results suggest a roughly constant elasticity for any given FICO score over time, meaning that the change in loan demand tends to change by the same factor as the interest rate across a credit score gap changes. There are two main institutional details that support the idea that lenders are not changing screening in lockstep with rate changes, which would be necessary to argue that the effects I measure are supply-side rather than demand-driven. First, LLPAs change discretely, and the reason the rate gap changes is the change in the valuation of the upfront cost of the LLPA, driven in part by the mortgage stack. It is unclear why these mechanisms should induce a roughly constant response from lenders. Second, LLPAs are imposed by Fannie/Freddie and go to the GSEs, so the nature of lender rationing would have to exactly move in line with the LLPAs to obtain consistent elasticity measurements over time. That is, higher LLPAs (higher rate spreads) would have to induce greater relative screening below any given threshold. While LLPAs are typically raised to protect Fannie and Freddie against perceived default risk, this default risk should not directly affect lender behavior, since the lenders only assume “putback risk” on the loans.

Ideally, one could test whether lenders are rationing supply by examining borrower characteristics for applications and for accepted loans. If the characteristics across the cutoff are discontinuous, the change in observed loan count and size might be attributed to differential screening across the cutoff. Unfortunately, the data available do not contain much information about selection. HMDA mortgage applications are not linked to FICO score, and most data sets do not contain additional demographic or income information linked to specific mortgages.

RD Validity and Empirical Robustness

The interpretation of discrete jumps in originations and defaults at certain FICO score thresholds has been the subject of some academic debate. Keys et al. (2010) and Keys, Seru, and Vig (2012) argue that the discontinuity in default rates from FICO 619 to FICO 620 can be attributed to moral hazard induced by the increased likelihood of FICO 620 loans to be securitized by the GSEs. Bubb and Kaufman (2014) instead argue that the cutoffs are driven by lender screening, as evidenced by the discontinuous number and default rate of loans at these same credit score cutoffs, so that the exclusion restriction of using this cutoff as an instrument for securitization is not valid. For the empirical approach of this paper, I use higher FICO scores than those used in this previous literature. I show that for these cutoffs, there does not appear to be evidence of differential securitization or default trends, indicating that there is not differential lender screening across these thresholds. This result is key for interpreting my other measurements as the elasticity of demand for interest rates.

Testing for Supply Side Factors: RD Tests on Defaults and Securitizations

One potential concern on interpreting the regression discontinuity results is whether the jumps in origination amount across mortgage rate breakpoints is driven by supply-side factors (such as lender rationing and screening) rather than demand. In this section, I show that there does not appear to be evidence that lenders screen differentially across our cutoffs, and argue they do not have incentive to screen differentially. This provides further support that the regression discontinuity is isolating the demand elasticity of borrowers to interest rates.

The default and securitization trends are shown in Figure 9, and Table 7 shows the results for defaults at the various cutoffs. The coefficients for the discontinuity at the first two cutoffs–619/620 and 639/640–are positive and significant, with economically meaningful magnitudes, indicating that there is likely differential screening across these breakpoints. While FICO 619 and FICO 620 borrowers have essentially identical default risk *ex-ante*, “bad” borrowers are more likely to get approved at FICO 620 given lender screening rules-of-thumb that causes more intense screening below cutoffs. Hence, there is a discrete jump upwards in the default risk–i.e. better FICO score individuals just above the cutoff default more than those individuals just below.

Promisingly, these default discontinuities disappear after the 620 and 640 breakpoints. The results for the higher FICO breakpoints are statistically insignificant, and the point estimates are economically close to zero. This holds regardless of the bandwidths tested. These results indicate that there is no evidence of differential lender screening across breakpoints 660 and higher.

Table 7 also shows the results for the discontinuity in securitization rates across breakpoints. At FICO breakpoints up to and including 680, there is evidence that there is increased securitization just above the cutoff. This may induce differential lender screening: since otherwise identical borrowers are more likely to be securitized at 680 than at 679, the lenders may have an incentive to screen 679 borrowers more carefully, as there is a greater chance that the loan will be held in portfolio and therefore that the lenders will be directly exposed to default risk. These trends indicate that for the purposes of my analysis, only FICO breakpoints 700 and above can be used.

For robustness, I rerun this analysis for an earlier time period. Because my mortgage rate data are limited, I do not use this time period for my demand elasticity measurement. However, the results indicate that even in earlier periods used by previous literature, there do not appear to be default or securitization cutoffs at higher FICO scores.

Smoothness Across Thresholds for FHA

This paper estimates the impact of conventional mortgage rates on conventional mortgage loans. One potential concern about this identification strategy is that the propensity to obtain FHA loans and FHA loan sizes might also respond to the mortgage rate differentials across these cutoffs. To test this, I estimate a regression discontinuity in the count of FHA mortgages over each threshold. If there were a substantial breakpoint in the propensity to get an FHA loan at my cutoffs of interest, then my estimates would not be valid: some individuals might still obtain mortgages but under the FHA program instead of Fannie/Freddie, but I would count these borrowers as selected out of mortgages on the extensive margin.

My estimates suggest that FHA switching is not a major concern in my setting. Table 7 shows the RD exercise performed on the propensity of borrowers to get an FHA loan. There is no change across thresholds to get FHA loans: the estimates are close to zero, with the 720 breakpoint indicating that being above 720 FICO amounts to a change in 0.42 individuals *per 100,000* getting an FHA mortgage per week. This is small relative to the total of 10.4 (per 100,000) individuals getting a conventional loan per week and small relative to the estimated discontinuity of 9.5 individuals (per 100,000) per 25 basis point of rate change. Hence, while FHA switching might add some noise to the general estimates, the RD exercise is still valid. [24]

Testing "Search-and-Wait" Behavior

Another potential concern for the regression discontinuity validity is that borrowers query for their credit score multiple times, only receiving a mortgage if their credit score is above the desired threshold. If this were true, it would mean that my method does not pick up an extensive margin of mortgage borrowing, but rather the same borrowers "timing" their loans to attempt to get the best possible mortgage rate.

In this subsection, I show that borrowers do not seem to time their mortgages to their credit score. I justify this in a few ways. First, I discuss that why this sort of timing would be difficult, if not impossible, given how the mortgage market is structured. Second, I show that credit scores move somewhat randomly before acquisition of a mortgage. Finally, I study shopping behavior using mortgage inquiries in my credit bureau data.

First, it is difficult, if not impossible, for borrowers to acquire real-time data on credit scores. The Fair Credit Reporting Act (FCRA) grants consumers free access once each year to their credit report from each major credit bureau, although take-up seems to be low; the CFPB estimates that only 10 percent of the eligible population receives a free credit report each year.[25] Even though some credit cards offer "free FICO" reports, the credit score displayed tends to be at least a couple weeks lagged, and is therefore not incredibly useful for individuals trying to time their mortgage purchases. Moreover, previous studies have found that mortgage shopping hardly occurs. A recent CFPB report finds that about 77 percent of purchase mortgage-borrowers apply to only one lender, and are therefore unlikely to be querying their credit score by visiting loan officers multiple times.[26]

[24] Given that FHA mortgage insurance premia have changed over this sample, it is important also to test the smoothness across time. The reason this could be a concern is that the tradeoff between taking a conforming vs. an FHA loan could change over time as the rate spreads between the two vary. When FHA loans are much more expensive than conforming loans, as they have been historically, borrowers may only resort to FHA when they cannot afford a down payment, and this may not vary discontinuously across the FICO breakpoints we consider. In contrast, if FHA loans are less expensive for some FICO borrowers but not others, this may vary with the FICO-induced breakpoint in mortgage pricing. Untabulated results indicate that there are no significant FHA breakpoints even breaking the sample into periods of two years each.

[25] See Consumer Financial Protection Bureau (2015a)

[26] See Consumer Financial Protection Bureau (2015b)

If credit scores systematically moved upward before a mortgage, then we would worry that some borrowers query their credit score, attempt to manipulate it by closing accounts or reducing other loan balances, and then re-query their credit score, repeating the process until their credit score is high enough to take advantage of the discretely lower mortgage rate. I show that borrowers' credit scores do not systematically increase before mortgages are originated, indicating that at large, manipulation does not happen. In Figure 10, Panel A shows the initial FICO time at the start of the search against the origination FICO score. Panel B shows a similar graph, using instead median credit score both at start of search and origination (which reflects the fact that lenders use the median of 3 credit scores). Both graphs are roughly linear. There is no breakpoint at any given FICO score, indicative that there is no differential mortgage timing that causes bunching at higher credit score breakpoints; rather, the distribution of credit scores at origination is as smooth as the credit scores at the start of the mortgage search.

To study shopping behavior, I examine the number of queries for mortgages for each borrower up to 6 months before the mortgage was originated. The number of mortgage inquiries is smooth across credit scores at origination, as seen in Panel C of Figure 10. This provides evidence that individuals just below the cutoff are not inclined to continue inquiring about their credit score to facilitate mortgage timing. The number of months searched across FICO score at origination is shown in Panel D of Figure 10, with 95 percent confidence intervals shaded in grey. The number of months searched is roughly smooth across the breakpoints I consider. If anything, the number of months searched is higher just below a given cutoff than just above. This is potential evidence of failed manipulation: one could imagine that a borrower has a fixed number of months before she must buy a house, and she buys if her FICO score falls above a desired threshold or at her deadline. The fact that there is a slight uptick of the number of months searched *under* the FICO thresholds indicates that these individuals may not have been able to wait sufficiently long before originating a mortgage.[27]

[27]Consider the expected stopping time of a Brownian motion with drift. The decision rule of the agent is to stop at $\min\{\bar{T}, T_{stop}\}$. We can solve for the expected stopping time T_{stop}. But there is noise in the process, and the presence of a strict deadline causes the distribution of realized stopping times to be truncated on the right by $\bar{T}$. In this case, conditional on realizing the desired FICO score, the average actual stopping time is less than $\bar{T}$, whereas conditional on being below the cutoff, the actual stopping time is exactly $\bar{T}$. This rough sketch shows why we might expect the average shopping time for mortgages to be higher just under a FICO pricing cutoff.

Attempted manipulation of credit score

One might worry that manipulation of credit scores could invalidate my regression discontinuity. In this section, I explain both that because of institutional structure, this is theoretically impossible, but even if manipulation could occur, under some fairly loose assumptions the results would still be valid.

In theory, it should be impossible to manipulate credit scores. The credit score agencies change their credit scoring algorithms over time, and these algorithms are never revealed to the public. Still, individuals on online forums speculate that they can take certain actions, such as carrying less debt and closing extra tradelines, to improve their credit score.

Even if individuals are aware of credit score thresholds and attempt to manipulate their scores, my regression discontinuity strategy is still valid and is still as good as randomized as long as there is continuous noise in the ability of individuals to affect their credit score ("partial manipulation", rather than "complete manipulation").[28] Since credit scoring formulae are "black boxes" to the public, with uncertainty regarding how any given chosen action, such as closing a credit account or paying off debt, might affect one's credit score, it seems likely that the assumptions necessary for "as good as randomized" to hold.[29]

The noisiness of credit scores can be seen in Figure 4. The distribution of the 6-month lagged FICO scores for borrowers with a 699 or 700 FICO score at origination is quite wide, with many borrowers falling below 680 or above 720. Moreover, the distributions of previous credit scores for borrowers who ultimately originate at FICO 699 and 700 are virtually indistinguishable. The histogram looks similarly noisy even for one-month lagged FICO score.[30]

[28]This was shown in Lee (2008). Formally, each individual is assigned a FICO score V, which is influenced partially by (a) the individual's choices and characteristics, and (b) by random chance. Treatment (in this case, a lower mortgage rate) is given to an individual if and only if V is greater than the known FICO threshold v_0. Lee (2008) shows that if, conditional on the individual's choices and characteristics, the probability density of V is continuous–even if this density function of V varies across individuals–in the neighborhood of $V = v_0$, variation in the treatment status is *as good as randomized by an experiment*, and satisfies the minimal assumptions for RDDs. McCrary (2008) formalizes an econometric test based on this argument to test for a discontinuity in the running variable. This relies on manipulation to be monotonic, meaning that manipulation happens only in one direction, which is intuitive in our setting. This test would not apply to our count of actual loans since this is a choice variable that responds to mortgage rate changes. Moreover, the test relies on a continuous forcing variable, whereas the underlying potential borrower FICO scores are discrete.

[29]The CFPB held focus groups and found that "many consumers said they were not sure how to improve their scores and were confused by conflicting advice about which actions to take. See Consumer Financial Protection Bureau (2015a).

[30]The noise in lagged FICO scores means that a fuzzy RD–that is, trying to instrument for FICO at origination using previous FICO scores–has a weak first stage and cannot be performed.

Continuity of Characteristics in Underlying Population

Since credit scores are assigned by credit bureaus, one might worry that there is a mechanical breakpoint in individual characteristics at the FICO scores I study. In this section, I discuss trends in individual-level credit data across credit scores and show that the underlying population is continuous across a variety of credit variables.

I run regression discontinuities using a similar methodology to my base results, except that the population is now the entire Equifax sample rather than restricted to mortgage borrowers. To transform the Vantage credit scores provided by Equifax to FICO scores, I use the merged Equifax-McDash LLC data to derive the historical relationship between Equifax score and FICO score and use this adjustment so that the RD can be performed across FICO thresholds. Untabulated results indicate that using raw Equifax-prvided credit scores does not change the qualitative results.

Table 8 shows the results of these RDs for a select set of credit variables. The base bankcard balance increases over credit scores, but I cannot reject that the bankcard balance is smooth across credit score cutoffs (i.e. the 95 percent confidence intervals of the regression discontinuity comfortably contain 0). Car debt decreases over credit scores, but again the results indicate we cannot reject a smooth distribution of car debt across credit scores. Credit utilization, measured as the total balance held by an individual divided by the total credit available to that individual, tends to fall as credit scores increase, in line with the typical intuition that richer and higher-credit individuals are less likely to be credit constrained.

The results indicate that the underlying population is not substantially different at the FICO breakpoints used for the elasticity measurement. This provides further support for the validity of the regression discontinuity.

Economic Implications

While the estimates in this paper are meant to capture the partial equilibrium response of mortgage demand to interest rates, one way to gauge the potential economic magnitude of these estimates is to calculate the implied counterfactual mortgage demand had LLPAs not existed.

Suppose the mortgage rates were the same for borrowers between FICO 680 and 720 (i.e., LLPAs were the same across these borrowers). Consider FICO borrowers from FICO 680-699. These mortgage borrowers would have seen a decrease in their mortgage rate of about 25 basis points on average. The estimates imply that borrowers would have demanded about 5 more mortgages per week (per 100,000 borrowers), an increase of just over 25 percent. Moreover, borrowers would have demanded an origination amount of about $13k more. Using these numbers indicates an increase in new mortgage demand of about $5.3 B and an increase in origination amount from existing mortgages of $1.2 B.

A similar exercise can be performed for FICO 700-719 borrowers. While their rate spread was smaller, around 10 basis points, there are more borrowers who fell within these FICO scores over the sample than borrowers in the 680-699 bin. Back-of-the-envelope calculations imply an increase in new mortgage demand of about $5.6 B and an increase in origination amount from existing mortgages of $2 B.

Overall, this counterfactual suggests that had mortgage rates been the same for borrowers with FICO scores between 680-719 as those with FICO scores above 720, an additional $14.1 B in mortgage demand could have been created between October 2008 and December 2014, with the majority of the increase coming from new mortgage originations. This is a large potential increase relative to the actual mortgage demand of approximately $43.5 B for these borrowers over this period.

Conclusion

In this paper, I have developed a novel methodology to measure the demand elasticity of purchase mortgages to interest rates. My estimates suggest that individuals who are buying a new home are sensitive to interest rates both on the extensive (choosing whether to get a mortgage) and intensive (the size of the mortgage) margins. The magnitude of the estimates is large, indicating that a 25 basis point decrease in interest rates corresponds to a 50 percent increase in the likelihood of a potential borrower to demand a mortgage and an increase in loan size of approximately $15k. I show further evidence that borrowers with high FICOs are more sensitive to interest rate changes than those with smaller (but still high) FICO scores, elasticities are approximately constant over time, and the marginal responsiveness to interest rates is decreasing.

While the estimated elasticity of mortgages to interest rates is larger than recent literature that focuses on mortgages, the results are not economically surprising in light of the general view that durable goods tend to have higher price elasticities than nondurables. For instance, Mankiw (1983) derives model conditions that imply that the sensitivity of expenditure on consumer durables to the interest rate is much larger than that of nondurables and services. Attanasio, Koujianou Goldberg, and Kyriazidou (2008) investigates the demand for auto loans and finds that when credit constraints are not binding, individuals' demand for auto loans is quite sensitive to interest rates. Finally, Monacelli (2009) highlights the macroeconomic importance of the larger sensitivity of durable goods to monetary shocks and develops a variant of the New Keynesian model that is consistent with this observation.

The large responsiveness of borrowers to interest rates has policy implications. For instance, if policymakers wanted to encourage homeownership, it may be too blunt of a tool to lower the risk-free rate, since low-FICO borrowers are often subject to additional fees or higher interest rates in *all* borrowing markets they engage in. Rather, one could imagine a government subsidy to cover the LLPAs which induce the interest rate spreads in my paper. While this would be controversial for political reasons, it may well induce a greater marginal responsiveness of borrowing per dollar than (say) the outright purchase of mortgage backed securities by the Federal Reserve. This is all speculative, and would benefit from a more rigorous framework, which is beyond the scope of this paper.

There may be further implications of this study on the compositional change in lending after the crisis, with a larger share of lending going toward higher-FICO individuals. While this paper has focused mainly on high FICO scores due to the identification strategy, determining the impact of constrained supply (lender screening) versus decreased demand from mortgage borrowers for the lower FICO individuals is also of importance.

Finally, there are implications for my estimate beyond those I developed in the paper. For instance, under certain assumptions, one could theoretically back out the elasticity of intertemporal substitution (EIS). Also, while I have established the large magnitude of demand responsiveness, I have not investigated the drivers of the responsiveness. By calibrating a simple model, one could also investigate whether the quantities found in this paper are consistent with a frictionless world, or whether credit constraints (e.g. a binding payment-to-income constraint) are an important driver to the demand response observed. It may also be possible to examine further micro data on cash purchases or loan applications to further understand the underlying drivers of demand.

References

Anenberg, Elliot, Aurel Hizmo, Edward Kung, and Raven Molloy. 2015. "The Effect of Mortgage Credit Availability on House Prices and Construction: Evidence from a Frontier Estimation Approach."

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): 401–436.

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."

Brueckner, Jan K. 1994. "The demand for mortgage debt: some basic results." *Journal of Housing Economics* 3(4): 251–262.

Bubb, Ryan, and Alex Kaufman. 2014. "Securitization and moral hazard: Evidence from credit score cutoff rules." *Journal of Monetary Economics* 63: 1–18.

Consumer Financial Protection Bureau. 2015a. "Consumer voices on credit reports and scores." Available at http://files.consumerfinance.gov/f/201502_cfpb_report_consumer-voices-on-credit-reports-and-scores.pdf.

Consumer Financial Protection Bureau. 2015b. "Consumers' mortgage shopping experience." Available at http://files.consumerfinance.gov/f/201501_cfpb_consumers-mortgage-shopping-experience.pdf.

DeFusco, Anthony A, and Andrew Paciorek. 2014. "The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit."

Fannie Mae. 2007. "Announcement 07-16." fanniemae.com [Online; posted 6-November-2007].

Fuster, Andreas, Stephanie Lo, and Paul Willen. 2017. "The Time-Varying Price of Financial Intermediation in the Mortgage Market." *FRB of New York Staff Report* (805).

Fuster, Andreas, and Basit Zafar. 2014. "The sensitivity of housing demand to financing conditions: evidence from a survey." *FRB of New York Staff Report* (702).

Gelman, Andrew, and Guido Imbens. 2014. "Why high-order polynomials should not be used in regression discontinuity designs." Technical Report. National Bureau of Economic Research.

Glaeser, Edward L, Joshua D Gottlieb, and Joseph Gyourko. 2012. "Can cheap credit explain the housing boom?" In *Housing and the Financial Crisis*, 301–359. University of Chicago Press.

Glaeser, Edward L, and Jesse M Shapiro. 2003. "The Benefits of the Home Mortgage Interest Deduction." In *Tax Policy and the Economy, Volume 17*, 37–82. MIT Press.

Goodman, Laurie S, Jun Zhu, and Taz George. 2014. "Where Have All the Loans Gone? The Impact of Credit Availability on Mortgage Volume." *Journal of Structured Finance* 20(2): 45.

Imbens, Guido, and Jeffrey Wooldridge. 2007. “Regression Discontinuity Designs, NBER Lecture Notes 3.”
Available at http://www.nber.org/WNE/lect_3_rd_fig.pdf.

Karlan, Dean, and Jonathan Zinman. 2013. “Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in Mexico.” Technical Report. National Bureau of Economic Research.

Keys, Benjamin J, Tanmoy Mukherjee, Amit Seru, and Vikrant Vig. 2010. “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans.” *The Quarterly Journal of Economics* 125(1): 307–362.

Keys, Benjamin J, Amit Seru, and Vikrant Vig. 2012. “Lender screening and the role of securitization: evidence from prime and subprime mortgage markets.” *Review of Financial Studies* 25(7): 2071–2108.

Lee, David S. 2008. “Randomized experiments from non-random selection in US House elections.” *Journal of Econometrics* 142(2): 675–697.

Li, Wei, and Laurie Goodman. 2014. “A Better Measure of Mortgage Application Denial Rates.” *Washington: Urban Institute*.

Mankiw, N Gregory. 1983. “Consumer durables and the real interest rate.”

McCrary, Justin. 2008. “Manipulation of the running variable in the regression discontinuity design: A density test.” *Journal of Econometrics* 142(2): 698–714.

Monacelli, Tommaso. 2009. “New Keynesian models, durable goods, and collateral constraints.” *Journal of Monetary Economics* 56(2): 242–254.

Sherlund, Shane M. 2008. “The Jumbo-Conforming Spread: A Semiparametric Approach.”

Swanson, Jann. 2015. “CoreLogic: Income Inequality and the Housing Market.” MortgageNewsDaily.com [Online; posted 28-January-2015].

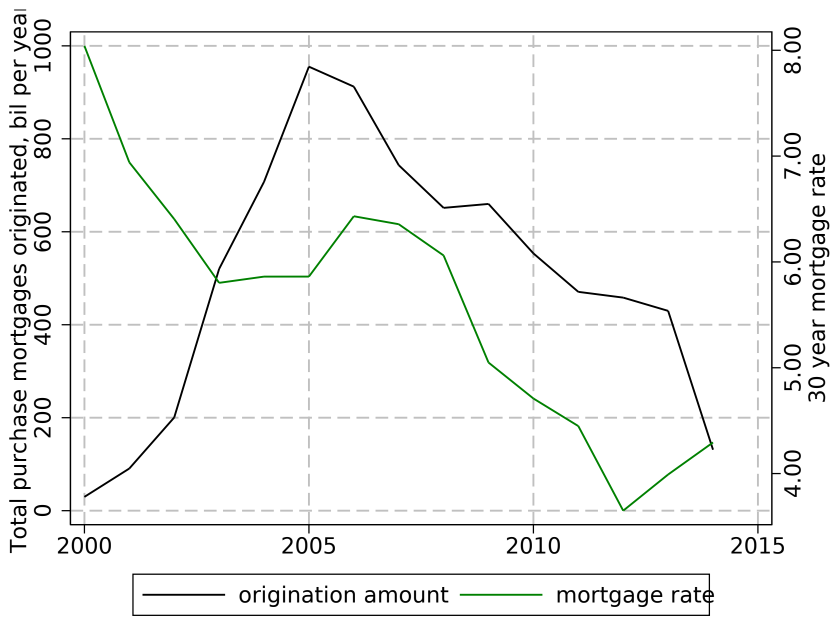

Figure 1: Total mortgage originations and mortgage rates over time. The mortgage rate series comes from the Freddie Mac Primary Mortgage Rates survey. Mortgage originations data is calculated as the total recorded origination amount for purchase mortgages by year, using the proprietary McDash LLC data.

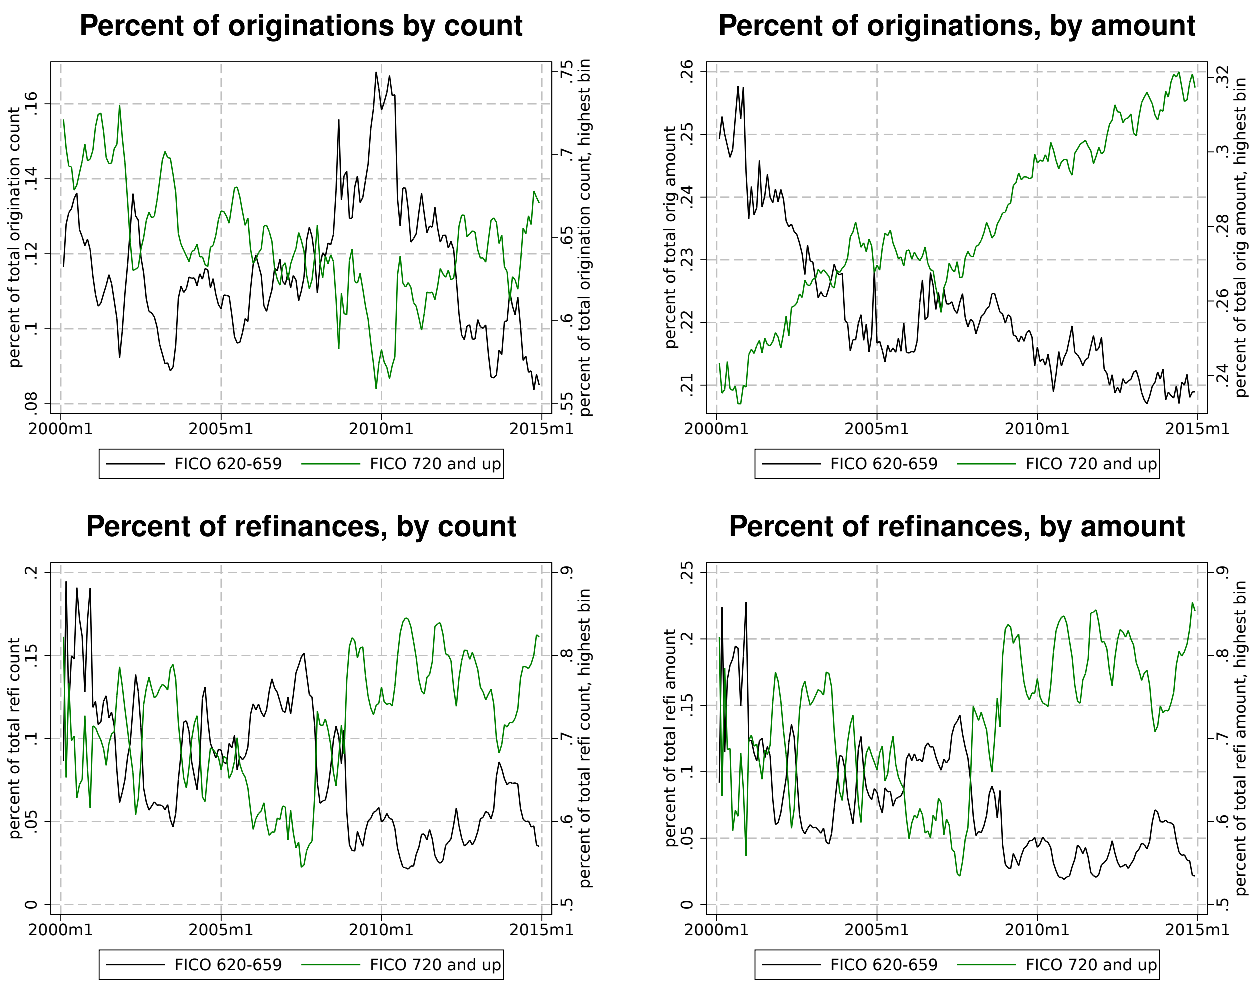

Figure 2: Originations (top) and refinances (bottom) for low and high-credit consumers. Count on left panel and amount or right panel, both as percent of total across all FICOs. The other categories not shown (FICO 660-719) have approximately constant percentage shares over this period. Source: McDash LLC, Author's calculations.

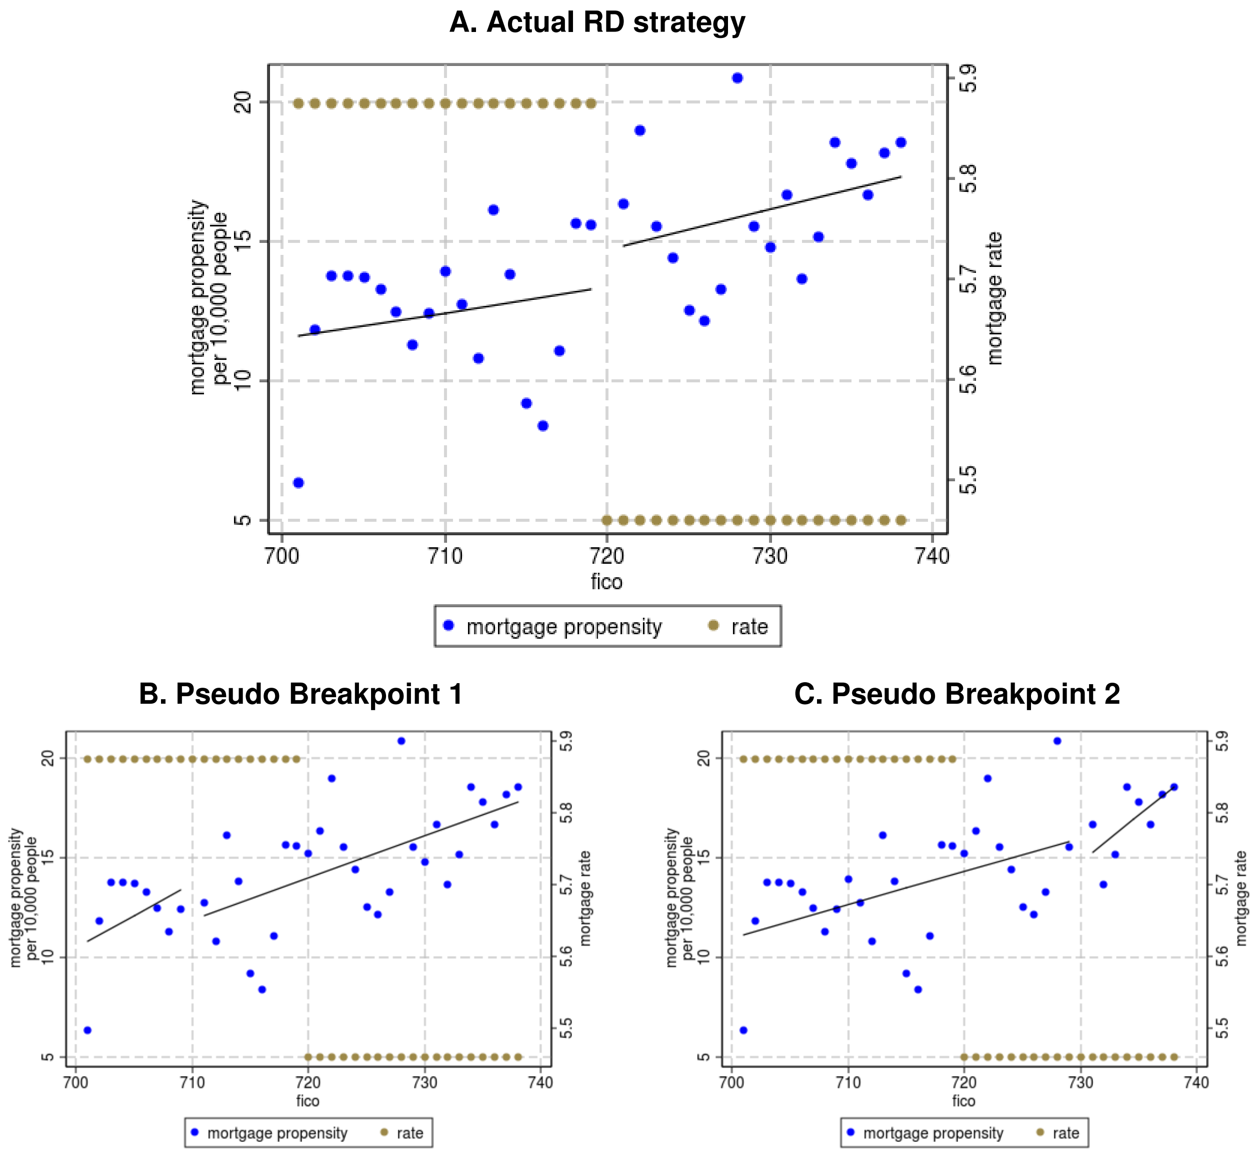

Figure 3: Example of RD strategy. Primary y-axis shows the measure of mortgage propensity (number of mortgages obtained in the week relative to the estimate of potential borrowers). The secondary y-axis shows the estimated mortgage rate for a borrower with the given FICO score; note the discontinuous jump at FICO 720. For 2009 week 5 purchase mortgages only. Graph shows simple linear regression on each side of the breakpoint. Panels B and C show different "pseudo" RD breakpoints. Source: New York Fed Consumer Credit Panel/Equifax, McDash LLC, Optimal Blue, Author's calculations.

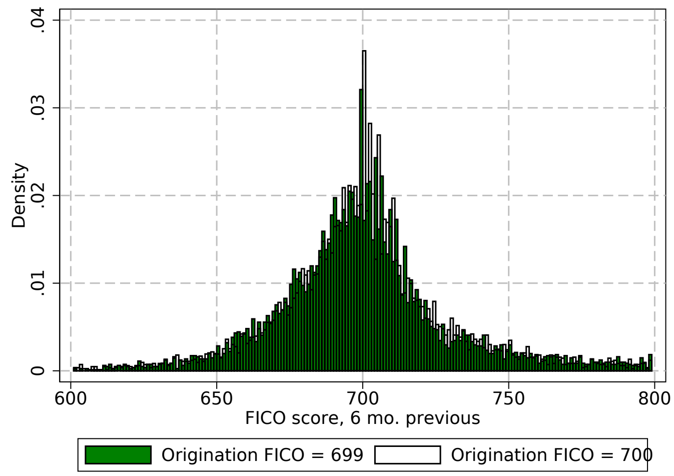

Figure 4: FICO 6 months previous to origination, conditional on FICO at origination being 699 or 700. Only about 3.5% of borrowers who originate loans at FICO 700 were at FICO 700 six months before origination; the rest of the borrowers are distributed both above and below the cutoff with quite a bit of noise. Source: Author's calculations and New York Fed Consumer Credit Panel/Equifax.

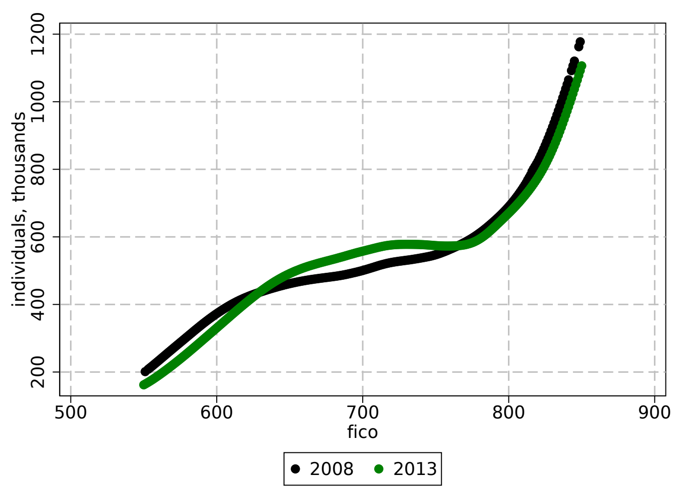

Figure 5: Total borrowers, as measured using the NY Fed CCP. All primary (5% sample) borrowers are pulled from the data, then observations missing risk score are removed (approx. 11% of data). The quarterly data is collapsed to a yearly frequency using a mean of the count over all quarters for a given year. To facilitate comparison with the loan data, which contains FICO score at origination, the Equifax riskscore is roughly converted to FICO using historical relationship using borrower-level data. Source: New York Fed Consumer Credit Panel/Equifax, Author's calculations.

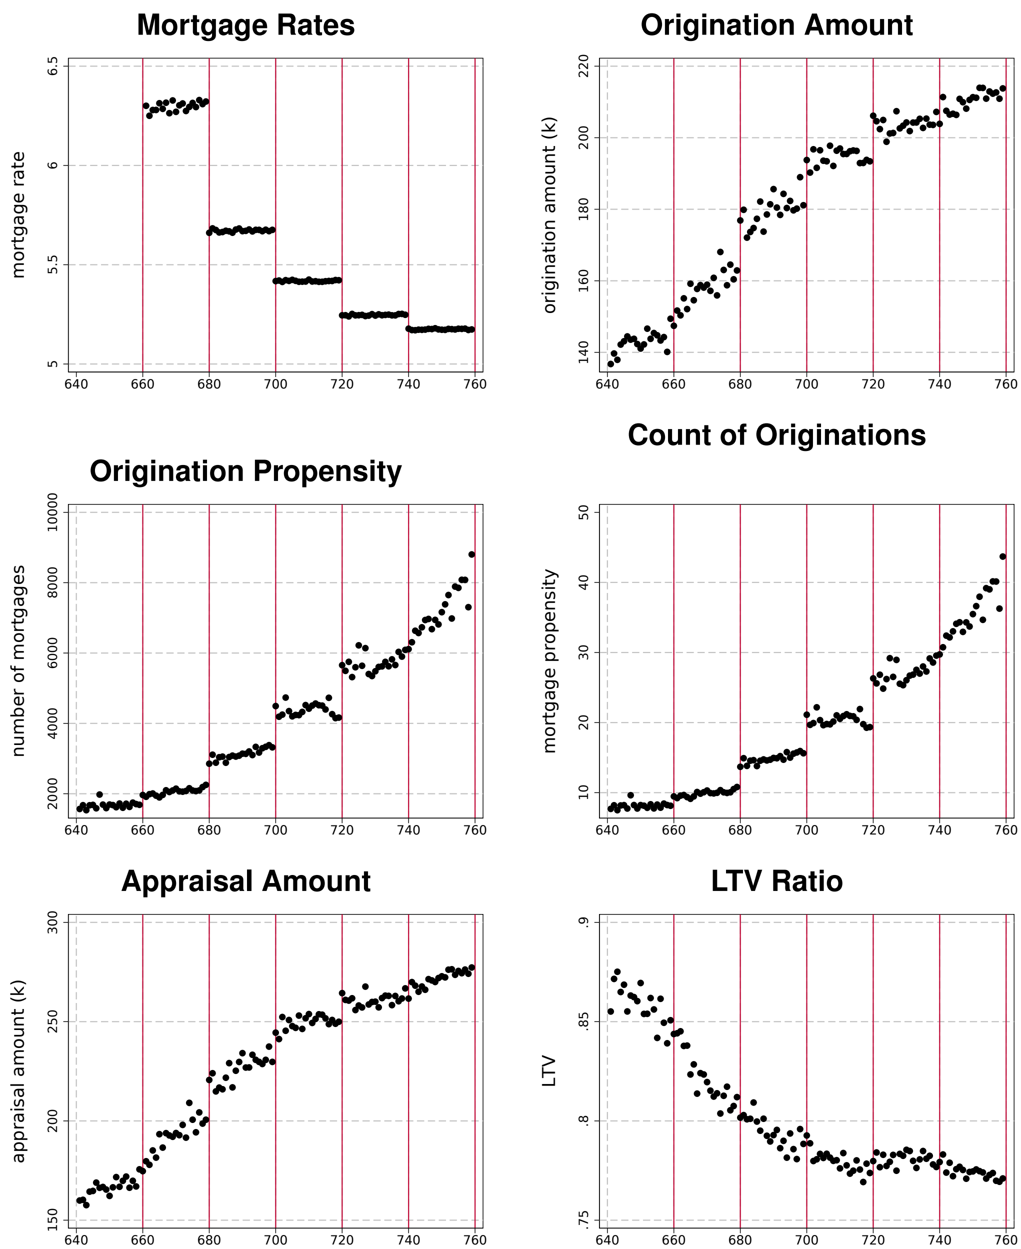

Figure 6: Mortgage characteristics by FICO score. Graphs show the averages in 2009 by individual FICO bin. Source: McDash LLC, Author's calculations.

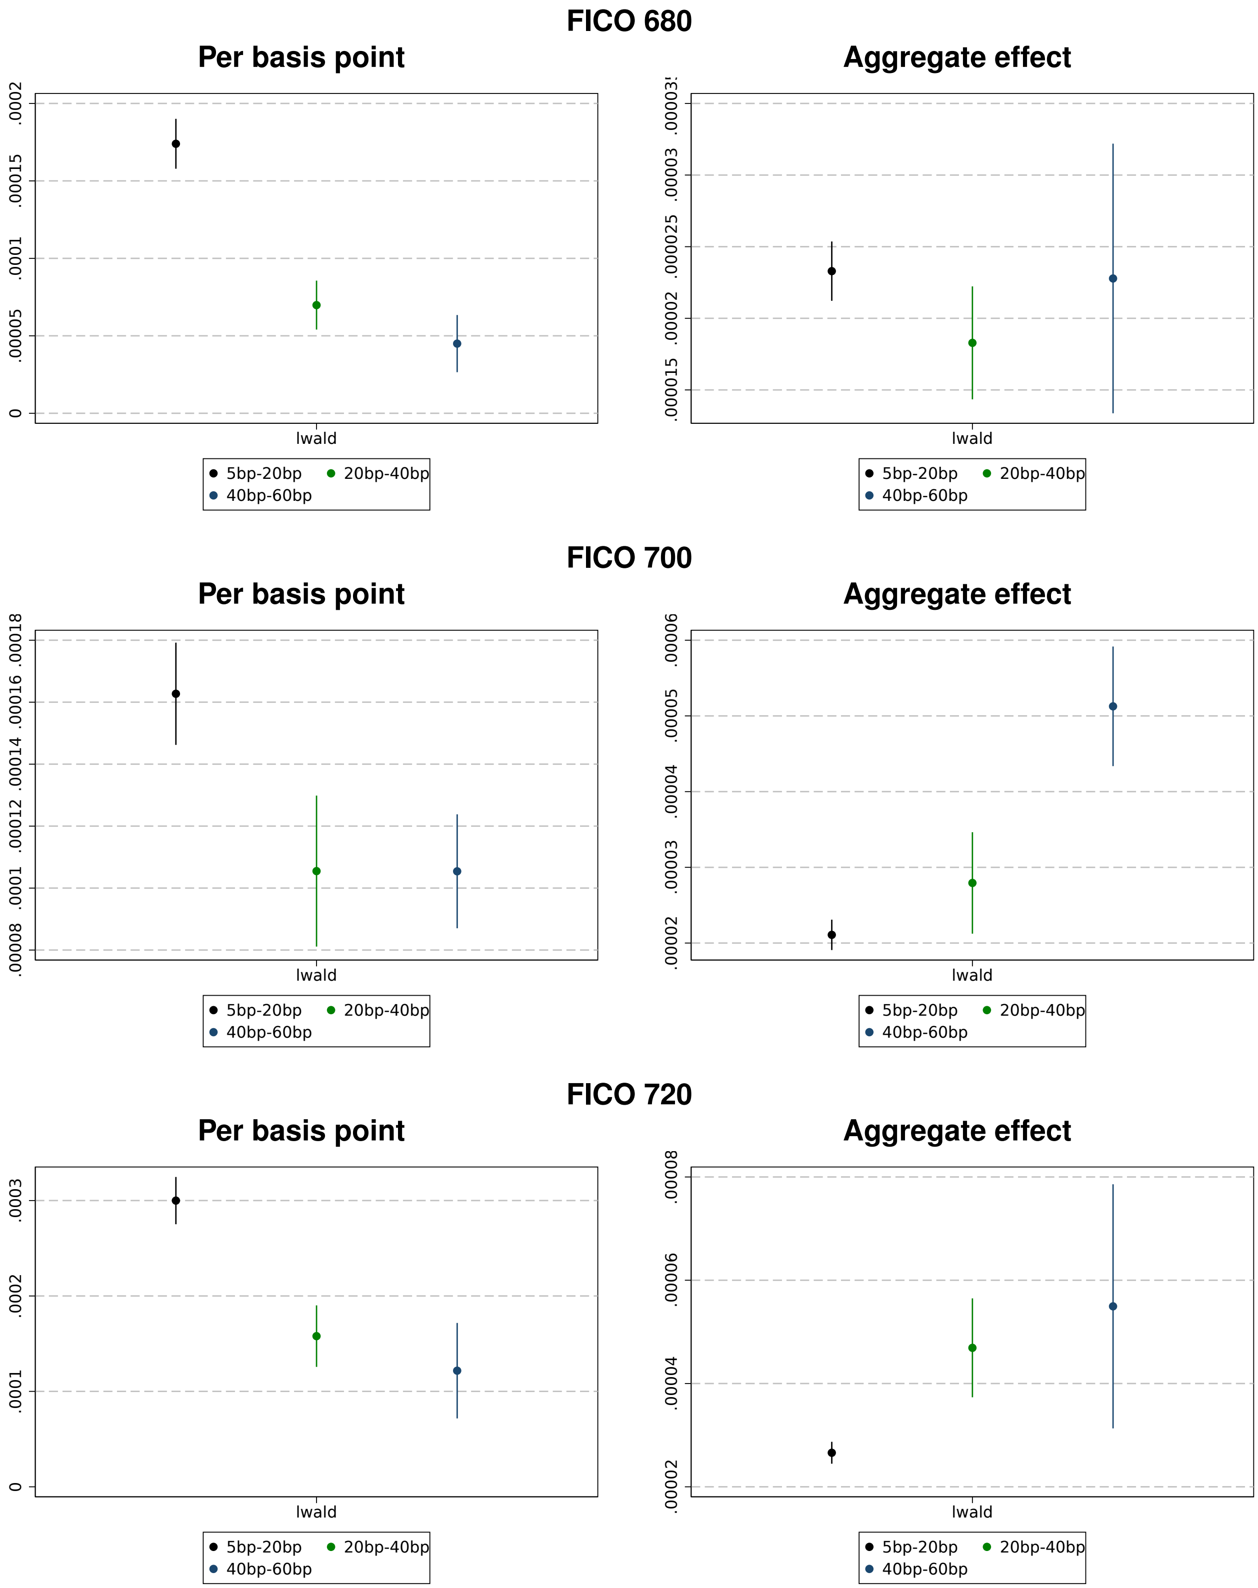

Figure 7: Extensive elasticity estimates, by size of mortgage rate spread. Markers denote point estimate and lines show 95% confidence interval from 1000 bootstraps. Source: Author's calculations, McDash LLC, Optimal Blue, New York Fed Consumer Credit Panel/Equifax.

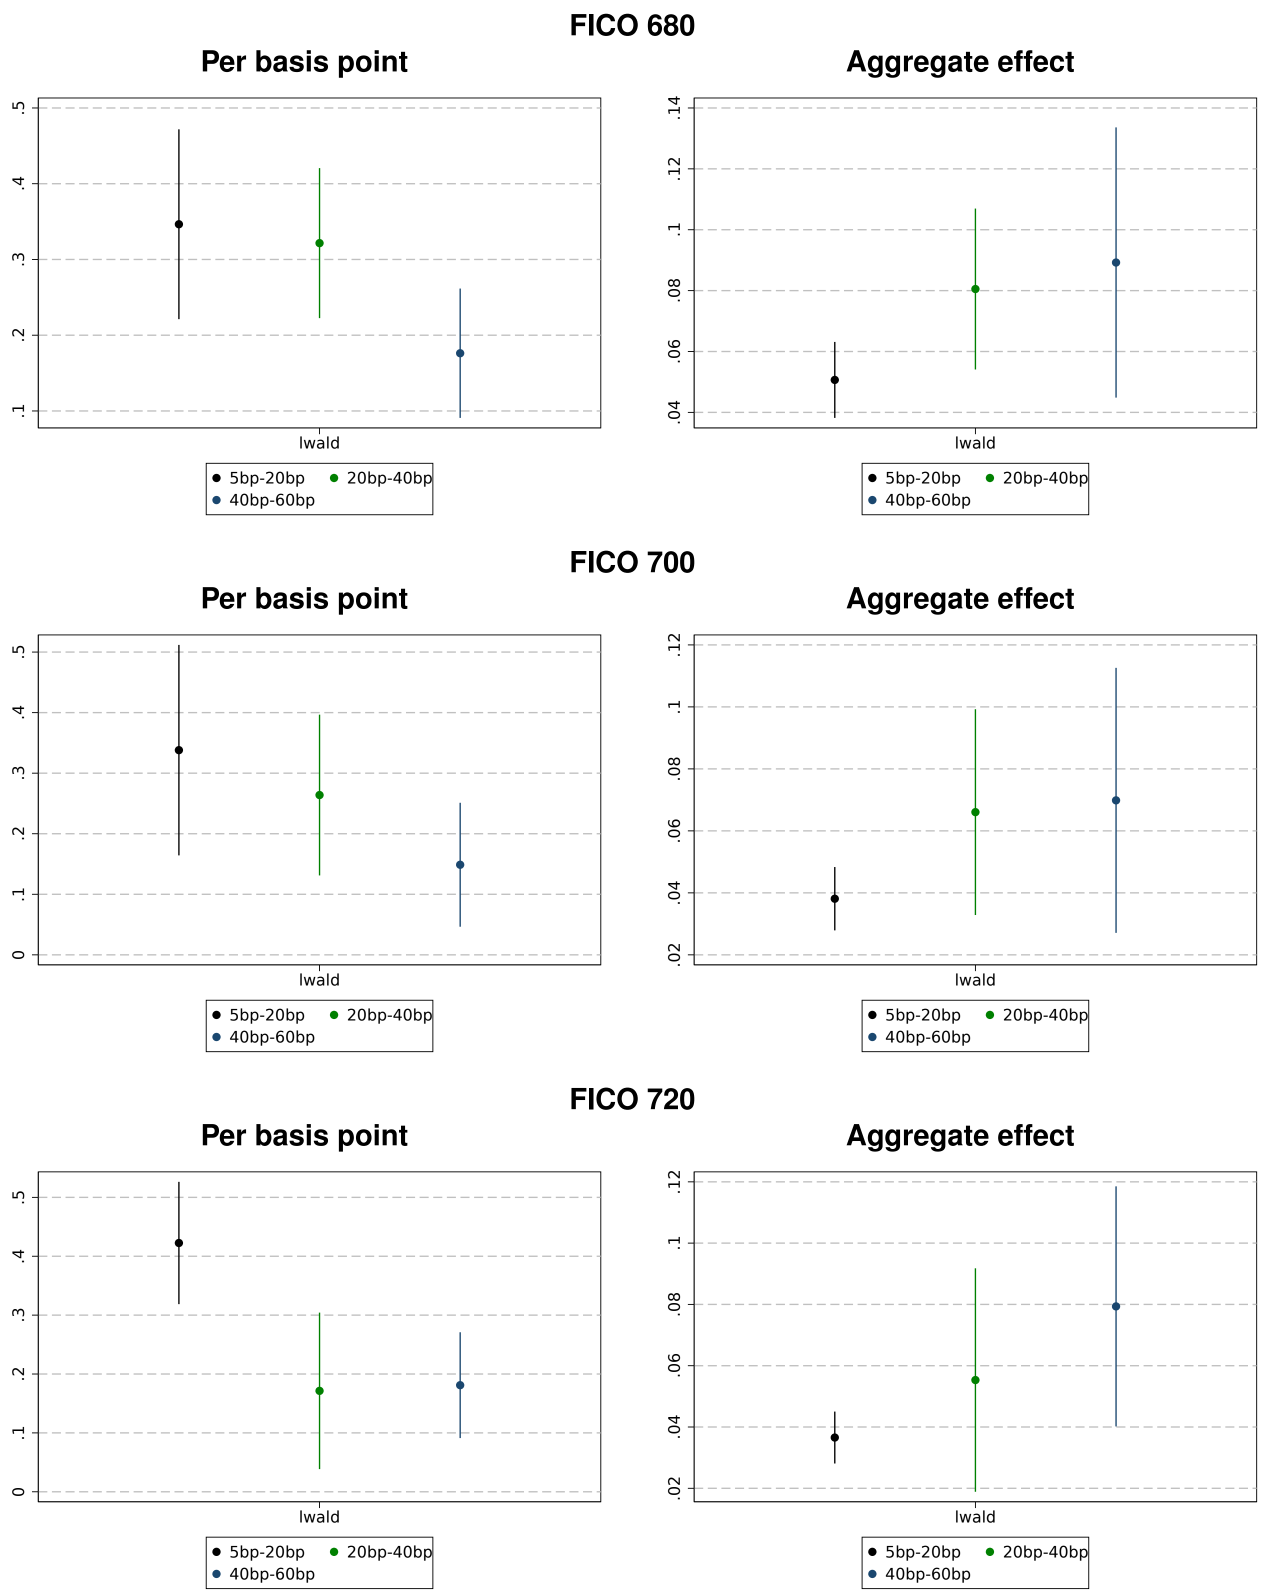

Figure 8: Intensive elasticity estimates, by size of mortgage rate spread. Markers denote point estimate and lines show 95 percent confidence interval from 1000 bootstraps. Estimates are shown in raw form, rather than being normalized per 10,000 individuals as in the text. Source: Author's calculations, McDash LLC, Optimal Blue, New York Fed Consumer Credit Panel/Equifax.

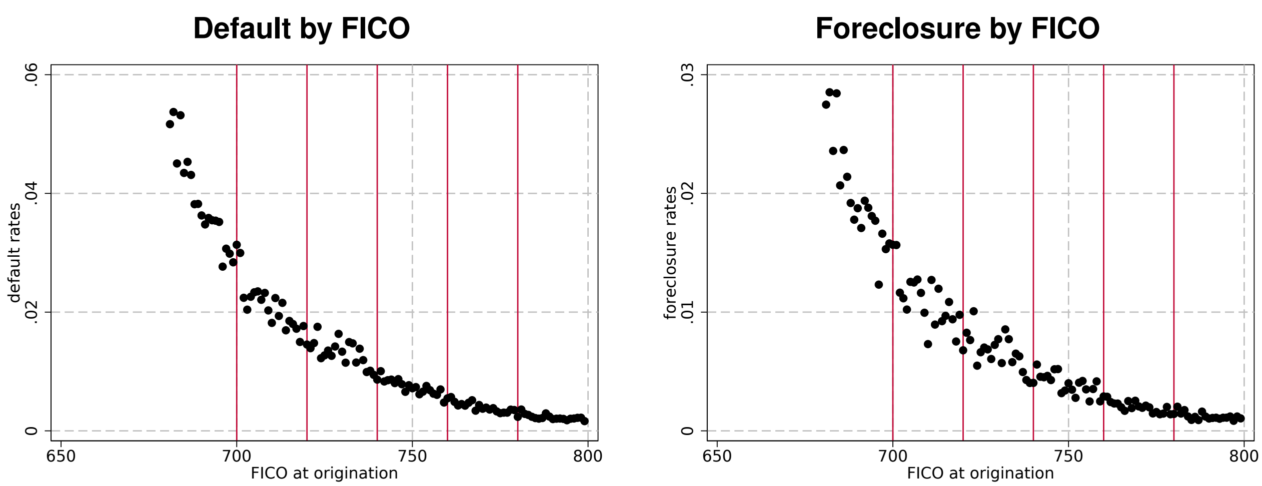

Figure 9: Default and foreclosure trends across breakpoints. Sample is all conventional 30-year fixed rate purchase mortgages between January 2008 and December 2015. A loan is considered in default if it has ever been 61 days delinquent within 36 months of origination. A loan is considered to be in foreclosure if the foreclosure process has ever been started within 36 months of origination. I test whether a significant discontinuity exists at the FICO breakpoints in Table 7. Source: McDash LLC, Author's calculations.

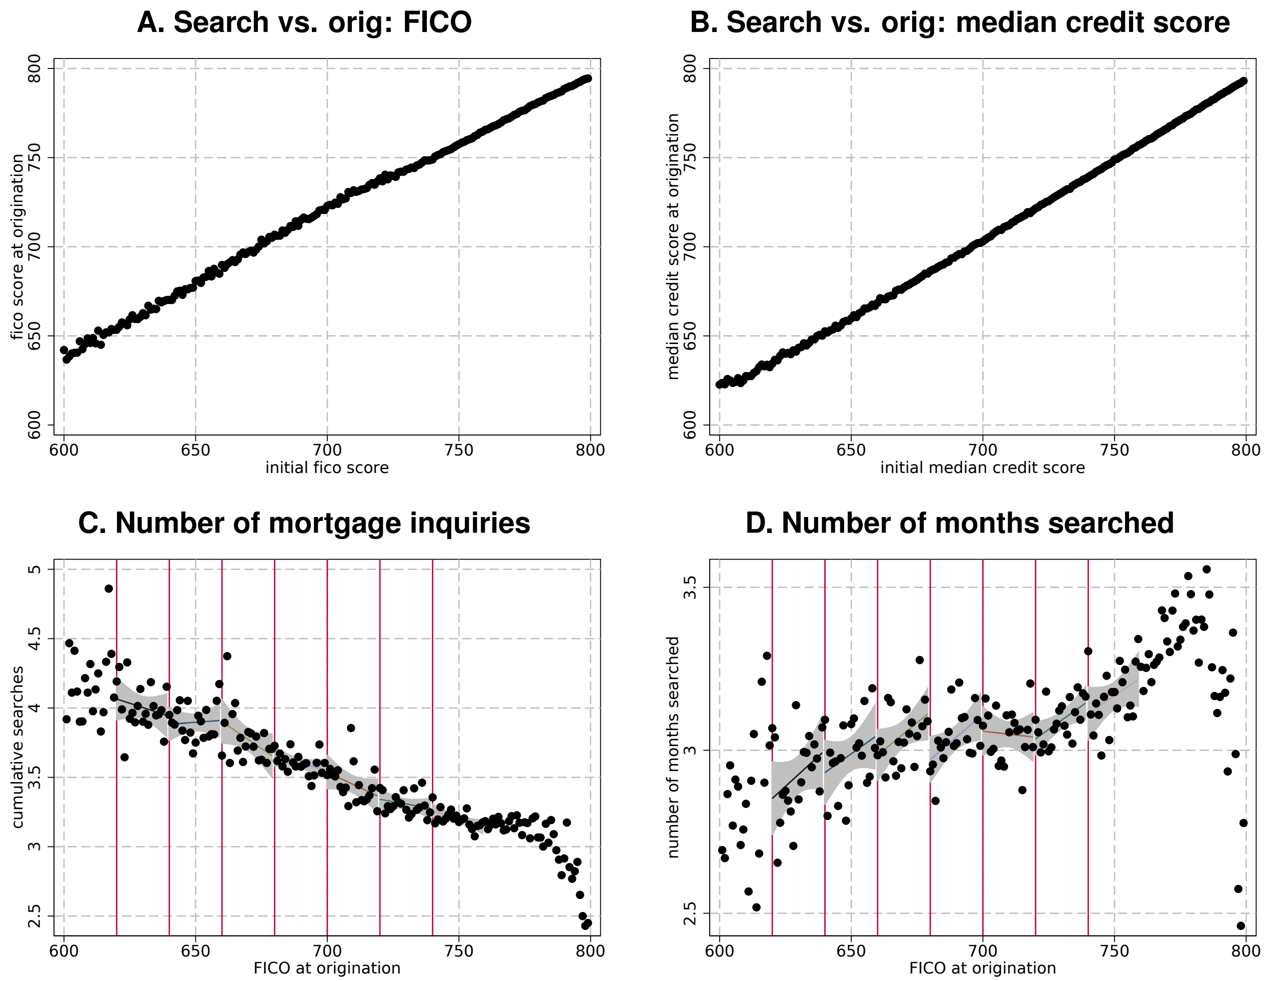

Figure 10: Mortgage shopping. All panels aggregate the baseline dataset, i.e. individual conforming purchase loan data from October 2008 to June 2015. Panel A shows the FICO score at origination (taken as a mean across loans) compared to the FICO score when mortgage shopping started, and Panel B shows the median credit score at origination (again the mean across loans) compared to the median credit score at the time of initial mortgage search. Panel C shows the number of mortgage inquiries in the 6 months previous to origination versus the credit score at origination. The number of cumulative searches shows no discontinuity across mortgage cutoffs. Panel D shows the number of months searched on average before the mortgage was originated, with the 95% confidence intervals shaded in grey. Source: McDash LLC, New York Fed Consumer Credit Panel/Equifax, Author's calculations.

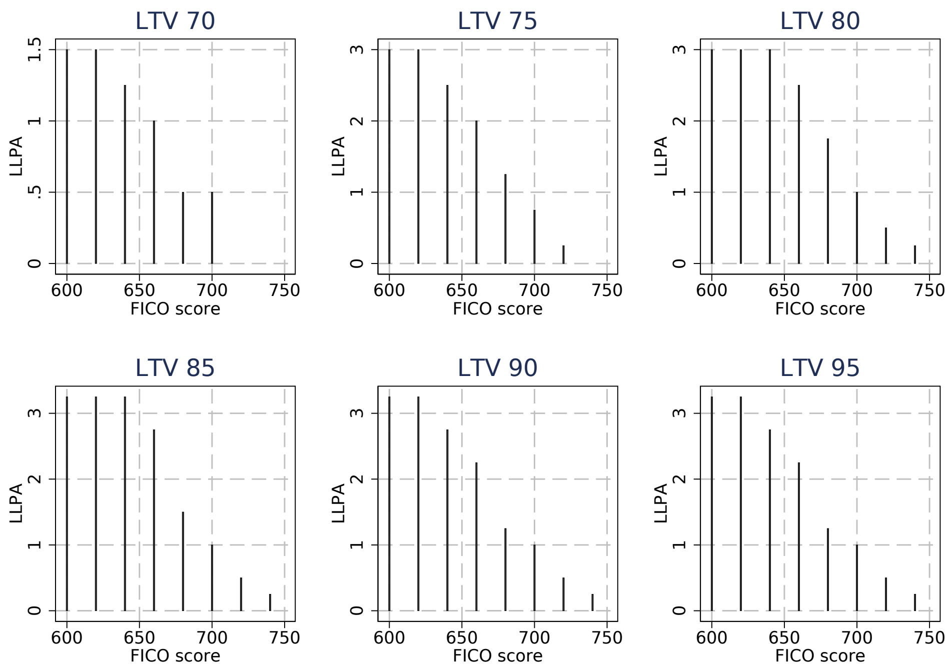

Figure 11: LLPAs for LLPA matrix effective April 1, 2011, which dominates the sample (effective until August 31, 2015). LLPAs are virtually constant across LTVs but vary much more across FICO scores. For the analysis, I assume the LLPAs are always those for a maximum LTV of 80, since borrowers tend to bunch at 80 LTV. Source: Fannie Mae.

Mortgage Rates									
	mean	sd	min	max	p10	p25	p50	p75	p90
660	4.77	0.99	3.44	13.71	3.70	4.17	4.58	5.23	6.16
680	4.50	0.69	3.32	6.57	3.56	4.04	4.40	5.03	5.49
700	4.35	0.66	3.18	6.55	3.45	3.88	4.28	4.85	5.23
720	4.24	0.63	3.09	6.37	3.37	3.80	4.19	4.77	4.99
740	4.20	0.62	3.05	6.36	3.34	3.77	4.15	4.73	4.91
Mortgage Rate Differences									
	mean	sd	min	max	p10	p25	p50	p75	p90
660	0.27	0.86	0.06	13.33	0.08	0.10	0.14	0.22	0.36
680	0.27	0.50	0.09	7.95	0.11	0.13	0.16	0.25	0.47
700	0.15	0.07	0.00	0.57	0.10	0.11	0.13	0.17	0.22
720	0.11	0.06	0.04	0.47	0.07	0.07	0.09	0.11	0.15
740	0.05	0.03	0.00	0.30	0.03	0.03	0.04	0.05	0.07

Table 1: Summary statistics, mortgage rates. Rates are for conforming 30-year FRM. The numbers shown reflect the mean across the entire baseline sample for the exact FICO score shown, on the weekly level, from October 2008 to December 2014. Higher FICO scores tend to benefit from lower mortgage rates due to lower upfront payments induced by LLPAs. Source: Optimal Blue and Fannie Mae; Author's calculations.

Date announced	Date effective	Overview of changes
November 6, 2007	March 1, 2008	First LLPA announcement
March 31, 2008	June 1, 2008	LLPAs increase for low-FICO borrowers
August 11, 2008	November 1, 2008	LLPAs decrease for high-LTV loans
December 29, 2008	April 2, 2009	LLPAs generally increased
September 22, 2009	January 1, 2010	Increased mortgage insurance
December 23, 2010	April 1, 2011	LLPAs changed for most loans with LTV at or above 70%
April 17, 2015	September 1, 2015	Increases in certain LLPAs

Table 2: LLPA history. For each change in LLPAs, the change was first announced via a press release and later became effective. Each LLPA change corresponded to a new matrix of LLPAs across credit scores and LTV ranges. Source: Fannie Mae.

	FICO Beacon 5.0	Vantage Score Solutions	Equifax Risk Score
min	300	501	280
max	850	990	850
mean	758	746	763
sd	49	62	52
p1	608	568	606
p5	659	621	661
p10	685	654	690
p25	730	710	736
p50	774	760	775
p75	795	795	801
p90	806	813	820
p95	811	820	824
p99	816	830	829

Source: New York Fed Consumer Credit Panel/Equifax

Table 3: Summary statistics about credit scores, sample. Subsample taken for all borrowers in dataset, up to 6 months before mortgage origination, as available in the New York Fed CCP / Equifax data. Source: New York Fed Consumer Credit Panel/Equifax.

	N mortgages	Total population	Mortgage propensity	Orig. amt	Appraisal amt	LTV	DTI
620	5		1.4	167377	246468	0.69	32.2
640	7	465,131	1.8	171328	258060	0.69	33.2
660	11	490,996	2.8	174968	254756	0.71	32.9
680	21	509,516	5.2	194279	279461	0.72	31.9
700	34	534,861	8.3	202698	289301	0.72	31.3
720	43	558,789	10.4	212630	303407	0.72	31.0
740	48	558,600	12.1	210294	298745	0.72	32.4
760	71		18.9	219076	314683	0.72	30.6

Table 4: Summary statistics, RD. The numbers shown reflect the mean across the entire baseline sample for the exact FICO score shown, on the weekly level (i.e. there are 21 purchase mortgages per week originated for FICO 680). The total borrowers is the count of all individuals with a credit history in the data, scaled to account for the fact that the data is a random 5% of the U.S. population. Source: McDash LLC and New York Fed Consumer Credit Panel/Equifax.

	(1) 660	(2) 680	(3) 700	(4) 720
Origination amount	53198.0***	52384.9***	78886.0***	29012.9
	[39846.5,66549.4]	[35044.0,69725.8]	[60119.6,97652.4]	[-8994.4,67020.3]
Base origination amount	174044.9	190357.9	200912.6	207352.7
Appraisal amount	46611.3***	46340.8***	63180.6***	-5841.7
	[26945.0,66277.7]	[15933.3,76748.3]	[36835.4,89525.7]	[-71197.7,59514.4]
Base appraisal amount	217942.7	242577.5	257840.3	273459.8
Loan-to-value ratio	0.0135	0.00148	0.0352***	0.0334
	[-0.00451,0.0315]	[-0.0176,0.0205]	[0.0140,0.0564]	[-0.0246,0.0913]
Base loan-to-value ratio	0.80	0.78	0.78	0.76
Mortgage propensity	21.58***	21.78***	37.70***	29.25***
	[19.48,23.69]	[19.53,24.03]	[34.13,41.26]	[19.63,38.88]
Base mortgage propensity	14.04	18.46	22.62	28.32

95% confidence intervals in brackets

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Base RD results. Intensive and extensive response of borrowers across interest rate discontinuities. All numbers are per 100 basis points of rate spread. Restricted to rate gaps of 5 basis points or more. Mortgage propensity is defined as the number of individuals who get a mortgage per 100,000 individuals in the population per week. Base bandwidth is 10 FICO points, but the results are robust to choosing bandwidths of 5 or 19 FICO points. Source: New York Fed Consumer Credit Panel/Equifax, Optimal Blue, McDash LLC, Author's calculations.

A. Mortgage propensity

	2008-2010	2011-2012	2013-2014
680	11.37***	19.58***	35.43***
	[9.279,13.47]	[16.64,22.53]	[31.21,39.65]
base mortgage propensity	13.27	14.29	14.66
700	18.92***	21.28***	25.29***
	[15.00,22.83]	[17.97,24.59]	[20.99,29.60]
base mortgage propensity	17.82	18.39	19.23
720	28.63***	40.23***	45.46***
	[24.22,33.04]	[33.58,46.88]	[38.15,52.76]
base mortgage propensity	22.48	22.41	23.00

B. Origination amount

	2008-2010	2011-2012	2013-2014
680	33388***	70809***	58534***
	[22359,44417]	[51207,90411]	[25625,91443]
base origination amount	167790	167501	188002
700	52156***	54337***	48993***
	[31564,72748]	[34410,74265]	[5529,92457]
base origination amount	185195	184218	202330
720	93437***	74414***	64210***
	[63696,123178]	[42999,105828]	[27730,100690]
base origination amount	198089	195685	209894

Table 6: Regression discontinuity estimates over time. The base RD is run across three time periods: October 2008-2010, 2011-2012, and 2013-2014. Source: McDash LLC and New York Fed Consumer Credit Panel/Equifax; Author's calculations.

	(1) 620	(2) 640	(3) 660	(4) 680
Discontinuity in P(FHA)	0.634***	1.020***	0.473***	0.441***
	[0.438,0.830]	[0.837,1.200]	[0.224,0.723]	[0.197,0.684]
Discontinuity in P(securitization)	0.0234***	-0.00696**	0.00324**	0.00251**
	[0.0112,0.0357]	[-0.0114,-0.00250]	[0.000417,0.00606]	[0.000524,0.00450]
Discontinuity in P(default)	0.0304***	0.0156***	-0.00940***	-0.00269
	[0.0113,0.0495]	[0.00453,0.0267]	[-0.0142,-0.00463]	[-0.00659,0.00122]
	(5) 700	(6) 720	(7) 740	(8) 760
Discontinuity in P(FHA)	-0.275***	0.417***	0.137*	-0.139*
	[-0.462,-0.0879]	[0.256,0.578]	[0.004,0.277]	[-0.293,0.0153]
Discontinuity in P(securitization)	-0.0000809	0.000591	0.000624	-0.000504
	[-0.00244,0.00228]	-0.00119,0.00237]	[-0.000728,0.00198]	[-0.00163,0.000627]
Discontinuity in P(default)	-0.000688	-0.000579	-0.000283	
	[-0.00187,0.000492]	[-0.00149,0.000329]	[-0.001102,0.000454]	

95% confidence intervals in brackets
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Discontinuities in supply-side propensities. Each number listed is the regression discontinuity for the FICO listed for each variable separately. The propensity to get an FHA loan is defined as the number of individuals getting an FHA loan per 100,000 population per week. A loan is considered securitized if it is ever bought in the first 36 months after origination. Default is defined as ever having been 61 days delinquent or more at any point 36 months after origination. Source: McDash LLC and New York Fed Consumer Credit Panel/Equifax; Author's calculations.

	(1) 680	(2) 700	(3) 720	(4) 740
Bankcard balance, current	17.09	-24.78	-46.18*	-26.32
	[-27.38,61.55]	[-73.38,23.82]	[-100.2,7.858]	[-73.61,20.98]
Base amount, bankcard balance	7542.8	8110.6	8505.3	8275.9
Car debt, conditional on having car debt	-15.84	-50.66	-74.92	22.17
	[-144.7,113.0]	[-147.2,45.91]	[-174.4,24.57]	[-69.48,113.8]
Base amount, car debt	15716.2	15274.9	15099.3	14947.3
Credit utilization	-0.442	0.0444	0.157	0.106
	[-1.172,0.289]	[-0.165,0.254]	[-0.122,0.437]	[-0.0987,0.311]
Base amount, credit utilization	0.689	0.628	0.584	0.551

95% confidence intervals in brackets
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: RD on full population. $+/-10$ FICO points is shown, although the results are qualitatively similar $+/-20$ and $+/-5$ FICO points. The Equifax credit score is adjusted using historical relationships to FICO scores as described in the main text. Credit utilization measured as the total balance held by an individual divided by the total credit available to that individual. Source: New York Fed Consumer Credit Panel/Equifax, Author's calculations.

	Conforming	FHA
2008	1910642	1099531
2009	1473077	1509258
2010	1231015	1259203
2011	1222996	878942
2012	1289119	669041
2013	1300823	498998
2014	1076761	308979

Table A.1: Conforming and FHA mortgage counts by year. Purchase mortgages only. Data source covers approximately 70 percent of the mortgage market. Source: McDash LLC.

	mean	sd	p1	p10	p25	p50	p75	p90	p95	p99
No mortgage	673.32	111.17	400.00	522.00	588.00	685.00	775.00	808.00	820.00	826.00
Has mortgage	722.22	96.31	423.00	582.00	671.00	751.00	797.00	820.00	825.00	830.00
Total population	689.39	108.96	406.00	534.00	609.00	712.00	784.00	815.00	822.00	828.00

Table A.2: Summary statistics about Equifax Risk Score, general population. Data is pulled for all of 2007 and the first half of 2008, so both time-series and cross-sectional data are included. The population is 5 percent of all individuals with a credit history in the United States. The mean credit score shown here is lower than in Table 3, reflecting the fact that purchase mortgages after the crisis disproportionately went to higher-credit individuals than historical norm. Unsurprisingly, credit scores for individuals without a mortgage are substantially lower than those with a mortgage. Source: McDash LLC and New York Fed Consumer Credit Panel / Equifax.

Appendix

A Additional Summary Statistics

The two included tables, Table A.1 and Table A.2 give further detail on the mortgage market and risk scores.

B Back of the Envelope: Mortgage Payments

Note the standard annuity formula which applies to fixed rate mortgages is:

$$A = P\frac{i(1+i)^n}{(1+i)^n - 1}$$

where A is the periodic payment amount, P is the principal amount on the loan net of down-payment, i is the periodic interest rate, and n is the total number of payments.

For a 30-year fixed rate mortgage, $n = 360$ months. Our standard back-of-the-envelope calculation assumes an origination amount of $200k and a mortgage rate of 5 percent, so $i = 5/12$[31]. For our baseline, the above formula gives an annual payment of $1073.64.

To think about what happens when the mortgage payment stays fixed and we lower the interest rate to 4 percent (=0.33 percent monthly), we solve for P_2 such that:

[31] This is the standard in the mortgage industry, so think of i as an annual percentage rate (APR) rather than an annual interest rate

	5 percent				4 percent			
Month	Payment	Principal	Interest	Balance	Payment	Principal	Interest	Balance
1	1073.64	240.31	833.33	199759.69	954.83	288.16	666.67	199711.84
60	1073.64	307.12	766.52	183964.59	954.83	350.68	604.15	180895.03
180	1073.64	505.84	567.80	135767.82	954.83	522.80	432.03	100955.11
360	1073.64	1069.19	4.45	186511.57	954.83	951.66	3.17	143739.01

Table B.3: Amortization Details. For $200k 30-year fixed rate mortgage at 5 percent versus 4 percent.

interest rates		holding orig. amount constant		holding mortgage payment constant
mortgage rate	monthly rate	monthly payment	total cost	orig. amount
3	0.25 percent	843	303,555	254,656
3.25	0.27 percent	870	313,349	246,697
3.5	0.29 percent	898	323,312	239,094
3.75	0.31 percent	926	333,443	231,830
4	0.33 percent	955	343,739	224,886
4.25	0.35 percent	984	354,197	218,246
4.5	0.38 percent	1,013	364,813	211,895
4.75	0.40 percent	1,043	375,586	205,817
5	0.42 percent	1,074	386,512	199,999
5.25	0.44 percent	1,104	397,587	194,428
5.5	0.46 percent	1,136	408,808	189,091
5.75	0.48 percent	1,167	420,172	183,977
6	0.50 percent	1,199	431,676	179,074

Table B.4: Hypothetical monthly payments or alternate origination amount, as interest rates change.

$$\bar{A} = 1073.64 = P_2 \frac{0.0033 \times (1.0033)^{360}}{(1.0033)^{360} - 1}$$

which gives us $P_2 = 225,976$. **Further intuition**. It may come as a surprise to some readers that reducing the interest rate from 5 percent to 4 percent – or 20 percent – only decreases the mortgage payment by approximately 10 percent, from 1073 to 954. This section gives further intuition.

Part of the monthly payment goes toward paying off the principal of the loan, while the other part goes toward interest. At the beginning of a 30-year FRM, the vast majority of the payment goes toward interest payment, with very little principal paid down. The interest due for any given payment is simply the debt owed at the beginning of the period multiplied by the interest rate. So for a $200k loan, the first interest payment is 5/12*1/100*200,000 = $833.33. The total payment is fixed at $1073.64, so the amount that goes toward paying back the principal is simply $240.31.

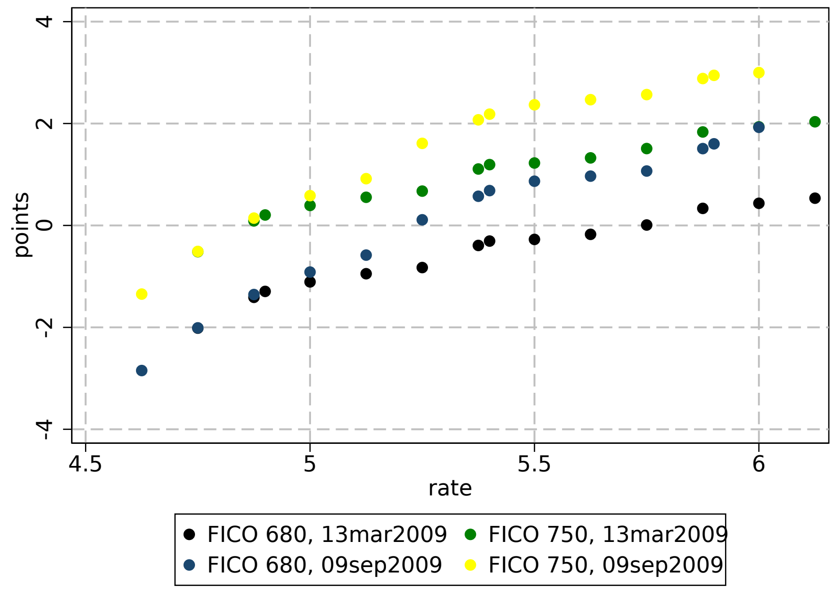

Figure C.1: Points and rates examples. Source: Optimal Blue.

C Points and Rates

This section gives further detail on the point-rate tradeoff for mortgages.

Figure C.1 show an example of the mortgage ratesheet points and rates for two FICO scores, FICO 680 and FICO 750, across two different (arbitrarily chosen) days.

First, it is obvious that the points offered per rate is higher for FICO 750 than it is for FICO 680. Or, put another way, on September 9, 2009, if both FICO borrowers wanted a yield spread premium of 100 (no points), then the FICO 750 borrower would have obtained a mortgage at approximately 4.75 percent, whereas the FICO 680 borrower would have obtained a mortgage at approximately 5.25 percent. Alternatively, if the higher FICO borrower were willing to borrow at the higher rate of 5.25 percent, she would have obtained about 1.8 points at closing. For a $200,000 loan, this is worth an upfront payment of $3600 in exchange for a mortgage payment that is $70 per month higher.

Solving for the implicit discount factor, we know that

$$PV = A\frac{1 - \frac{1}{(1+r)^n}}{r}$$

Recalling the average loan origination amount of $200,000, this is approximately $3600. (With an LTV of 80 percent, the appraisal amount would be $250,000 and borrowers would need to put down approximately $50k to begin with). The monthly payment on the 4.75 percent loan is 1043 (total cost of $375,586), and for the 5.25 percent loan is 1104 (total cost of $397,587).

D FHA vs Conventional Loans

There are two major types of mortgages: "conventional" (Fannie and Freddie) and Federal Housing Administration (FHA) loans. In this study, I solely measure the margins of adjustment with respect to conventional loans. This is justified by a number of reasons. First, with regard to the RD strategy, only conventional mortgages are subject to the pricing breakpoints across FICO scores, and given the smoothness of FHA mortgage acquisition, we can safely abstract from FHA switching and still maintain valid estimates. Second, FHA loans are generally more expensive, and conventional wisdom suggests that if a borrower can afford the larger down payment, she should obtain a conventional loan.

Conventional loans require a down payment of at least 20 percent or private mortgage insurance (PMI) if the LTV is higher than 80 percent. To work around this limit, some borrowers take out "piggyback" loans, or a conventional first loan with up to 80 percent LTV and a second loan (in place of PMI), often under a higher rate as it cannot be securitized via Fannie/Freddie.

FHA loans are often thought to be targeted at lower-income borrowers, with 580 FICO allowing eligibility for maximum financing. The standards tend to be looser but the fees higher. FHA loans only require 3.5 percent down payment, and this can be paid using gifted funds (whereas conventional loans often have standards in terms of personal income versus gifts), making the down payment more affordable for low-income individuals.

The higher cost of FHA loans comes from two targeted mortgage insurance premiums. First, the upfront mortgage insurance premium (UFMIP) charges the borrowers a premium of 1.75 percent [32], which can be paid upfront at closing or rolled into the mortgage. In order to compare across loans, I roll these costs into the rates. Because the FHA UFMIP is a percent of the mortgage size, the FHA rate for lenders who are otherwise identical, including identical LTVs, but with different absolute mortgage sizes will have different implied FHA rates.

Second, FHA loans come with a annual MIP (typically paid monthly with the mortgage payments) that differ depending on the borrower LTV, loan size, and loan length.

FHA loans also are more lax with the debt-to-income ratio, with conventional mortgages requiring borrowers to have DTI of 45 percent or less, while FHA allows borrowers to spend up to 56 percent of their income on monthly obligations like credit card payments...

The number of loans originated by FHA relative to conventional mortgages has increased substantially. During the subprime boom from 2003 to 2007, less than 10 percent of the purchase loans originated each year were backed by the FHA. By the end of 2009, that number ballooned to about 40 percent of all purchase originations. As the FHA increased its mortgage premia, that number has fallen to approximately 25 percent.

Is this increase in demand for FHA loans due to the easy credit standards or due to cheaper mortgage rates? Industry experts argue that conventional loans are generally less expensive for borrowers in almost all cases [33].

An additional level of complexity is that FHA loans and conforming loans have varying loan limits to qualify for the best rates. FHA standard loans are for amounts up to \$271050 and FHA jumbo loans are for amounts up to \$625500 although the maximums vary by county. On conventional loans, the conforming standard loan limit is for amounts up to \$417000; the conforming jumbo loans are up to \$625500 with maximum amounts varying by county. Mortgages exceeding this amount are not eligible for purchase by Fannie or Freddie.

[32] This has changed over time.

[33] See quote here: http://www.bankrate.com/finance/mortgages/homebuyers-get-fha-loans.aspx

Conventional PMI. One realistic simplifying assumption that allows me to run a robustness test is that PMI premia hardly change.[34] Hence, any week-to-week variation in the number of mortgages obtained can be attributed to changes in base rates rather than changes in the relative attractiveness of an under-80 LTV loan and an over-80 LTV loan.

[34] A simple archive.org search on Genworth's rate sheets indicate that many mortgage insurance premia are in effect for several months: for instance, as of March 21, 2012, the mortgage insurance premium ratesheet for August 1, 2011 was still in effect, and on April 20, 2015, the mortgage ratesheet for November 18, 2013 was still in effect. Individuals in the industry recognize that default risk changes, but note that the mortgage insurance premia are fixed for the life of the loan. These companies assume that default risk for a longer-term view (5-10 years) doesn't change that quickly, hence the mortgage insurance premia also change slowly.