Regulating Household Leverage

Online Appendices

Anthony A. DeFusco   Stephanie Johnson   John Mondragon

A Missing Debt-to-Income Ratios

The data we use in our analysis are sourced from loan-level records that are voluntarily provided to CoreLogic by a consortium of different loan servicers. A common issue confronting researchers who use this type of servicing data is that many servicers either do not record or choose not to report DTI ratios to third-party data aggregators (Foote et al., 2010). In studies such as ours that require a non-missing DTI, this can lead to a large number of loans being dropped. To the extent that these loans differ in systematic ways from those with non-missing DTIs, dropping them could bias our results. In this appendix we provide evidence that this issue is not a first order concern in our setting.

Since much of our analysis relies on a differences-in-differences framework that compares outcomes for jumbo relative to conforming loans before and after the policy change, the key concern is that DTIs are differentially missing from one of these two market segments in a way that is correlated with the timing of the policy change. In Figure A.1 we show that this is not the case. This figure plots time series trends in the share of loans that are dropped from our analysis sample due to having a missing DTI. The series in orange cricles plots the share of jumbo loans dropped from each origination-month cohort and the series in blue triangles plots the corresponding shares for conforming loans. While the overall incidence of missing DTIs is quite high (approximately 50 percent), it does not seem to be systematically different across jumbo and conforming loans and, more importantly, does not exhibit any differential trends around the time of the policy change, which is marked by the vertically dashed grey line.

As a further check on the potential magnitude of this issue we also explore the extent to which loans with and without a reported DTI differ across other observable characteristics in our data. The first two columns in Table A.1 report means and standard deviations for five loan-level characteristics that we can observe. The sample in column 1 includes only loans with non-missing DTIs, which is the same set of loans contained in our full analysis sample and described in column 1 of Table I. Column 2 reports the analogous statistics for the sample of loans with missing DTIs. In column 3, we report the difference in means between these two samples along with its standard error. While the large sample size leads many of these differences to be statistically significant, the economic magnitude of the differences are negligible in all cases. The FICO scores, LTVs, interest rates, and property types being financed are essentially identical across the two samples. The only potentially meaningful difference is in the average loan amount, which is roughly $8,000 higher in the non-missing DTI sample. However, even this amount is minimal compared to the standard deviation in that sample, which is almost $190,000. Finally, In column 4 we also report the standardized difference in means, which scales the difference in column 3 by the average of the standard deviations in each sample to provide a measure of economic significance.[1] All of the differences are less than five percent of a standard deviation and in most cases come in well below one percent of a standard deviation. Together with Figure A.1, these results lead us to believe that missing DTIs are not a major source of bias in our analysis.

B Additional Results and Robustness Checks

B.1 Selection on Observables in the Interest Rate Regressions

In Section IV.B, we argue that the post-policy increase in interest rates for high-DTI jumbo loans is unlikely to be driven by borrower selection. This argument is based on two patterns in the data. First, when we control flexibly for borrower characteristics such as LTV and FICO, our baseline results do not change, which suggests a limited role for selection on observables. Second, when we continue to condition on the same observables but also allow our effect to vary non-parametrically in the borrower's DTI, we find that the increase in interest rates for high-DTI loans is uniform across all DTIs greater than the 43 percent threshold. While not conclusive, this fact helps to allay some concerns over selection on unobservables since such selection would presumably be most severe at DTIs just above the threshold and therefore imply a non-uniform interest rate premium. As further evidence that selection is playing a limited role in our interest rate regressions, this section directly examines the extent to which there are any differential changes in observable borrower characteristics around the DTI threshold subsequent to the policy change.

In Figure A.2 we plot average borrower characteristics by DTI separately for jumbo loans originated before (blue triangles) and after (orange circles) the implementation of ATR. We focus on two borrower characteristics that are known to be important determinants of interest rates: LTV and FICO. Panel A shows that there is no differential change in LTVs for high-DTI borrowers relative to low-DTI borrowers following the policy change. The slightly downward sloping relationship between DTI and mean LTV is nearly identical in both time periods and suggests that the policy did not meaningfully change the composition of borrowers along this dimension. This result is confirmed by columns 1 and 2 of Table A.2, which report estimates from difference-in-differences regressions that use the borrower's LTV as the dependent variable. The estimates suggest that the policy lead to at most a 0.3 to 0.4 percentage point increase in mean LTVs at high DTIs, which is very small relative to the average LTV of 75 percent.

The results for FICO, plotted in Panel B of Figure A.2, reveal a small decrease in mean credit scores at high DTIs in the post period that is particularly concentrated among DTIs just above the

[1] In the case of the condo indicator, which is a binary variable, the standardized difference is calculated as $(\overline{x}_1-\overline{x}_1)/\sqrt{\overline{x}_1(1-\overline{x}_1)+\overline{x}_2(1-\overline{x}_2)}/2$, where $\overline{x}_i$ is the mean in sample i.

threshold. This pattern is consistent with the argument we lay out in Section IV that borrower selection, if present, should be most severe just above the threshold. It is also consistent with the results in Figure II, which shows that the raw average increase in interest rates for high-DTI loans subsequent to the policy change is slightly higher at DTIs just above the threshold. Reassuringly, however, when we condition on FICO in Figure III this pattern disappears completely. This suggests that including FICO in our regressions does a good job of controlling for whatever small degree of selection may be occurring subsequent to the policy change. Moreover, the difference-in-differences regressions using FICO as the outcome reported in columns 3 and 4 of Table A.2 imply that the reduction in average FICO scores is only about 5 points, which is economically very small. While our data do not allow us to directly observe the correlation between FICO scores and borrower unobservables, evidence from existing regression discontinuity studies that exploit similarly small variation in FICO scores suggests that borrowers who differ by only 5 FICO points are likely to be very similar along many other dimensions.[2] This leads us to believe that selection on unobservables is unlikely to be a concern in our setting. Finally, to the extent that these unobservables are correlated with FICO scores, the fact that our main point estimate only falls by 2-3 basis points when we control flexibly for borrower FICO suggests that these unobservables are unlikey to be an important driver of our results.[3]

B.2 Substitution Into the Conforming Market

Our approach to estimating the counterfactual DTI distribution in the absence of the ATR/QM regulation relies on the assumption that the distribution among conforming loans was unaffected by the policy (Assumption 1). However, one way to avoid taking a non-QM loan while still maintaining a high DTI would be to substitute into the conforming market. This substitution may be optimal for borrowers with high DTIs and loan amounts that are only slightly larger than the conforming limit. If this type of substitution were prevalent, it may lead us to over-estimate the intensive margin effect of ATR/QM on the DTI distribution and under-estimate the extensive margin effect since our estimate of the counterfactual distribution would feature too many loans above the 43 percent threshold and too few below it.

To gauge the extent to which this bias may be affecting our results, we look to see if there were changes in the amount of "bunching" at the conforming loan limit among high- relative to low-DTI loans after the policy was put into effect. The easiest way for a high-DTI jumbo borrower to substitute into the conforming market would be to decrease her loan size by the minimum

[2] See, for example Figures 5 and A2 from Agarwal et al. (2018), who show that borrowers in consecutive 5-point FICO bins are similar along many dimensions including income, overall indebtedness, and default.

[3] Indeed, we are still able to reject the null of a zero interest rate effect when we use a conservative bias-adjusted treatment effect that is based on the degree of coefficient stability and the change in R-squared moving from the uncontrolled regression in column 1 of Table II to the fully controlled specification in column 4 (Oster, 2016).

amount required to qualify as conforming. Therefore, if high-DTI borrowers are shifting into the conforming market we should see an increase in the amount of bunching at the conforming limit among high-DTI loans relative to low-DTI loans subsequent to the policy change.

To measure the amount of bunching at the conforming limit we follow the approach in Kleven and Waseem (2013). For a given sample of loans, we first center each loan at the conforming limit in the year that the loan was originated. A value of zero thus represents a loan size exactly equal to the conforming limit. We then group these normalized loan amounts into \$5,000 bins with upper limits equal to m_j ($j = -J, \ldots, L, \ldots, 0, \ldots, U, \ldots, J$), and count the number of loans in each bin, n_j. To obtain estimates of bunching and the counterfactual loan size distribution, we define an excluded region around the conforming limit, $[m_L, m_U]$, such that $m_L < 0 < m_U$ and fit the following regression to the count of loans in each bin

$$n_j = \sum_{i=0}^{5} \beta_i (m_j)^i + \sum_{k=L}^{U} \gamma_k \mathbb{1}\left(m_k = m_j\right) + \epsilon_j. \tag{1}$$

The first term on the right hand side is a 5-th degree polynomial in loan size and the second term is a set of dummy variables for each bin in the excluded region. Our estimate of the counterfactual distribution is given by the predicted values of this regression omitting the effect of the dummies in the excluded region. That is, letting $\hat{n}_j$ denote the estimated counterfactual number of loans in bin j, we can write

$$\hat{n}_j = \sum_{i=0}^{p} \hat{\beta}_i (m_j)^i. \tag{2}$$

Bunching is then estimated as the difference between the observed and counterfactual bin counts in the excluded region at and to the left of the conforming loan limit,

$$\hat{B} = \sum_{j=L}^{0} \left(n_j - \hat{n}_j\right) = \sum_{j=L}^{0} \hat{\gamma}_j, \tag{3}$$

while the amount of missing mass due to bunching is $\hat{M} = \sum_{j>0}^{U} \left(n_j - \hat{n}_j\right) = \sum_{j>0}^{U} \hat{\gamma}_j$. We set the lower limit of the excluded region to $-\$10,000$, and choose the upper limit to minimize the difference between bunching and missing mass to the right of the conforming limit in the excluded region. Standard errors are calculated using a bootstrap procedure as in Chetty et al. (2011).[4]

Figure A.3 reports results from this exercise. Each panel plots the observed loan size distribu-

[4] At each iteration (k) of the bootstrap loop we draw with replacement from the estimated errors, ϵ_j, in equation (1) to generate a new set of bin counts, n_j^k. We then re-estimate bunching using these new counts. Our estimate of the standard error for a given parameter is the standard deviation of the estimates across these bootstrap replications.

tion and our estimate of the counterfactual for a given sample of loans. The top row includes all loans regardless of DTI, with the columns distinguishing between loans originated before ATR/QM (2013) and loans originated afterward (2014). The second row includes only loans with DTIs in the region just below the 43 percent cutoff. We use the same DTI bins that we used to estimate the quantity effect in Section V, so that this row includes all loans with $DTI \in (38, 43]$. Similarly, the third row reports results for loans with DTIs strictly above the 43 percent cutoff. Each panel also reports an estimate of the amount of "excess mass" at the conforming limit, which we measure as the ratio of the number of extra loans bunching at the limit relative to the predicted counterfactual number of loans in that region, scaled by the width of the bin to convert to a density. For example, the excess mass of 6.79 reported in the top left panel implies that there was roughly 6.79 times more mass at the conforming limit in 2013 than would have otherwise been expected. This reflects the underlying incentive to bunch at the limit documented by DeFusco and Paciorek (2017), which results from differences in interest rates and underwriting standards that apply to jumbo loans even in the absence of ATR/QM.

Between 2013 and 2014 the overall amount of bunching decreased in all samples of loans, possibly reflecting the reduction in the interest rate spread on jumbo loans relative to conforming loans during this period. Importantly, however, this decrease was equally pronounced among high-DTI loans and loans with DTIs just below the QM threshold. Excess mass decreased by roughly 16.5 percent (from 6.70 to 5.60) among low-DTI loans and by 18 percent (from 6.32 to 5.18) among high-DTI loans. If high-DTI jumbo borrowers were differentially substituting into the conforming market after the policy, then we would have expected the decrease in bunching in the high-DTI market to be substantially less than that in the low-DTI market, where there is no extra incentive to bunch due to ATR/QM. If anything, we document that the decrease in bunching for high-DTI loans was slightly larger. We take this as fairly strong evidence in favor of our assumption that the DTI distribution among conforming loans was not materially affected by the policy.

As another way to evaluate whether substitution into the conforming market should be a major source of concern for our analysis, we can also perform simple back-of-the envelope calculations to determine whether such substitution would be optimal for the typical high-DTI borrower in our sample. For example, we can consider the incentives of the average jumbo borrower with a DTI above the QM threshold in 2013. As discussed in Section V.B, this borrower had a DTI of 45 percent, a loan size of $622,000 and an interest rate of 4.08 percent, which would imply a fully amortizing monthly payment of $2,998. If we assume that this mortgage was the only debt the borrower carried, then a 45 percent DTI would imply a monthly income of $6,663. The national conforming loan limit in 2013 was $417,000. If this borrower were to reduce her loan size to $417,000 then her monthly payment (assuming the same interest rate) would fall to \$2,010. At that monthly payment, however, the borrower's DTI would be only 30 percent, which is well below the 43 percent QM threshold. Rather than substituting to the conforming market, this borrower would be much better off simply lowering her DTI to 43 percent, which would require a jumbo loan for \$584,303. The fact that the average high-DTI jumbo borrower would be better off lowering her DTI than substituting to a conforming loan may explain the lack of differential bunching at the conforming limit among high-DTI borrowers subsequent to the policy change. This fact also provides further support for our assumption that the policy had no meaningful effect on the DTI distribution in the conforming market.

B.3 The Effect of Changing $\overline{d}$ on the Estimated Counterfactual DTI Distribution

In Table III of the main text, we show how changing the lower limit of the bunching region, $\overline{d}$, affects our estimates of the intensive and extensive margin quantity effects. There are two channels through which changes in $\overline{d}$ will affect the estimates. First, holding constant the counterfactual number of loans in each bin, changes in $\overline{d}$ will alter the range of "integration" over which differences between the empirical and counterfactual distribution are computed. This type of change will have no effect on the estimated number of loans missing from above the 43 percent limit. However, depending on the relationship between the true empirical distribution and the counterfactual as estimated using our preferred value of $\overline{d} = 38$, such a change could increase or decrease the number of loans we deem to be bunching under the limit. This will affect both the intensive and extensive margin quantity effects since both depend on the number of loans bunching below the limit.

The second channel through which changes in $\overline{d}$ could affect our estimates is through their effect on the counterfactual distribution itself. Holding constant the range of integration, changes in $\overline{d}$ will lead to a different estimate of the counterfactual. This is because the counterfactual distribution is constructed from ratios of the number of loans in a given DTI bin to the total number of loans below $\overline{d}$ (see Assumption 3 and equation (4)). This effect could also increase of decrease both the extensive and intensive margin quantity estimates.

The numbers reported in Table III reflect the combined effect of these two channels. In this section, we investigate the effect of the second channel alone. Holding constant the range of integration, the effect of a change to $\overline{d}$ on our estimates will depend only on the extent to which it alters the counterfactual. In Figure A.4, we explore how the counterfactual distributions estimated using the three alternative values for $\overline{d}$ that we consider (30, 35, and 40) differ from our preferred counterfactual which sets $\overline{d} = 38$. In Panels A–C we plot the number of loans in a given DTI bin from our preferred counterfactual on the x-axis against the number of loans in the same bin for each of the three alternatives. Each dot in the figure represents a single DTI bin.

In every panel, nearly all of the points fall exactly along the 45-degree line, which is what would be expected if changing $\overline{d}$ had no effect on the counterfactual. While the distributions are not truly identical, the differences between them are very small. This can be seen in Panel D, which plots the distribution of differences between the number of loans in a given DTI bin from our preferred counterfactual and the corresponding bin in each of the three alternatives. The maximum difference in any given bin is only 11 loans and the bin counts are within only one loan of each other in more than 25 percent of cases.

Given the results in Figure A.4, it is not surprising that our estimated quantity effects do not differ much when we hold the range of integration constant but use one of these alternative counterfactuals to calculate the bunching parameters. This can be see in Table A.3. The first column of the table repeats our preferred estimates, which set the lower limit of the bunching region to $\overline{d} = 38$. Columns 2–4 report analogous estimates from alternative specifications which set this limit to 30, 35, or 40 percent when constructing the counterfactual, but hold the lower limit constant when calculating the number of loans bunching below the limit or missing from above. The estimates are all very close to one another, which suggests that the primary difference between the numbers we report in Table III is coming from the effect of changes in $\overline{d}$ on the range of integration rather than on the estimated counterfactual distribution.

B.4 Documentation Status and the Relationship between DTI and Default

One potential concern with the performance results reported in Section VI is that they rely on the implicit assumption that the relationship between DTI and default is policy invariant. However, it is possible that the implementation of ATR/QM actually led to a change the nature of the relationship between DTI and default. For example, if the policy causes lenders to put more work into verifying income and debt, then DTI may become a stronger predictor of default going forward. This would mean that the slope of the relationship we estimate between DTI and default is too flat, which would lead us to underestimate the effect on the aggregate default rate.

To address this issue, we explore whether the relationship between DTI and default changes meaningfully with loan documentation status. To do so, we re-estimate the results reported in Panel B., column 4 of Table VI for the 2008 loan cohort separately by documentation status. These results are reported in Table A.4. The first column simply repeats the results from Table VI for reference. The second column restricts the analysis to the subset of loans that CoreLogic reports as having "full documentation." This sample should be reflective of the relationship between DTI and default in a scenario in which lenders are carefully verifying the borrowers income and debts. For completeness, the third column also reports results for the sample of "low documentation" loans. Comparing across columns, it is clear that our results do not depend on documentation status. While the implied reduction in the aggregate default rate is slightly larger if we use the relationship between DTI and default from the full-doc sample, the difference is statistically insignificant and economically minimal. This leads us to believe that even if ATR/QM led to an increase in the level of documentation and verification that lenders perform, our qualitative conclusion would remain the same. The relationship between DTI and default is simply not strong enough to generate meaningful improvements in the aggregate default rate given the share of loans that we estimate were affected by the policy.

C Stylized Theoretical Model

Our empirical results show that non-QM lending volume fell by more than would be expected given the observed increase in interest rates and that this decline was particularly concentrated among non-retail and non-portfolio lenders. In this section, we present a simple model grounded in realistic features of the U.S. mortgage market that helps to rationalize these results.[5]

The model features two key mechanisms. First, we argue that the ATR/QM rule exacerbated a classical agency friction in the market for securitized loans by increasing the cost of improper and difficult to verify documentation on high-DTI mortgages. This agency friction differentially impacts mortgage originators whose business model depends on selling loans into the secondary market. Second, we assume borrowers face a search friction when deciding where to apply and that they cannot effectively distinguish between lenders more or less affected by the regulation. In the model, these two frictions mean that high-DTI loans become unprofitable for non-portfolio lenders, but that borrowers cannot easily target their applications toward the less-affected portfolio lenders. Together, this means that many borrowers arrive at lenders constrained by the policy and who are unwilling to lend at an acceptable price, while other borrowers arrive at lenders only marginally affected by the law and who therefore charge little to no premium.

We present the model in detail below, but first provide some intuition. To see why the ATR/QM rule might operate through the securitization market, recall that the legal liability associated with non-QM lending depends crucially on the level and quality of documentation collected by the mortgage originator at the time of loan approval. In particular, even non-QM loans with DTIs greater than 43 percent can be deemed compliant with the law if the lender can prove that they correctly documented the borrower's income and arrived at a reasonable, good faith determination of the borrower's "ability-to-repay."[6] Because this investigation and documentation is costly, however, mortgage originators who pass non-QM loans to investors would

[5]The basic structure of our model builds on the framework in Bubb and Kaufman (2014), who study screening and moral hazard in the secondary mortgage market. We extend their framework to allow for a more realistic pricing mechanism and to include two types of lenders. We abstract, however, from the credit score threshold problem emphasized in their paper as the credit quality threshold in our setting (DTI = 43) was set exogenously by regulation.

[6]This arises as a result of the "General ATR Option" discussed in Section II.

prefer not to investigate the loans. In contrast, the investors in these loans, who bear both the default risk and a portion of the legal liability, would prefer that the loans be properly documented. This creates a classical agency problem. Investors cannot verify whether the originator properly documented the loan, and originators cannot credibly commit to doing so.[7] As a result of this agency conflict, investors will assume that all loans delivered by such originators come with prohibitively high default costs. Critically, lenders who keep all or some portion of originated loans on their balance sheet will be able to overcome this agency conflict. For these lenders, who internalize both the costs of default and the benefits of proper investigation, the ATR/QM rule will cause only a slight increase in the cost of origination.

This agency friction implies that the regulation will increase the cost of high-DTI loans substantially at some lenders, while having relatively minor effects at others. In a perfectly competitive market, all borrowers would move to less-affected lenders and there would be a very small quantity decline, consistent with the small increase in price. To reconcile the empirical puzzle of the small increase in price and large fall in quantities, we assume that applying to lenders is costly and that borrowers cannot perfectly direct their search toward lenders offering the best terms or highest chance of approval. As a result, some high-DTI borrowers will arrive at lenders severely affected by the ATR/QM rule while others will arrive at less affected lenders. Those arriving at affected lenders will face arbitrarily high interest rates and no loan will be originated, which we interpret as a denial. Those arriving at less-affected lenders will instead be able to borrow at only a slight premium. In the data, interest rates will not be observed on loans that are not originated. This means that the estimated interest rate premium will be based only on the relatively small costs faced by portfolio lenders. The quantity response, however, will include loans rejected by non-portfolio lenders and will therefore appear large relative to the observed change in accepted interest rates.

Our assumptions about borrower search are consistent with several first-order empirical facts about the U.S. mortgage market. In particular, the assumption that mortgage applications are costly is consistent with the fact that nearly all borrowers only apply to a single lender when searching for a mortgage.[8] Similarly, the assumption that borrowers cannot perfectly direct their search accords with the fact that up to 18 percent of mortgage applications result in rejection despite these apparently large application costs.[9] The most commonly cited reason for these rejections is DTI, which further suggests that borrowers have a particularly difficult time deter-

[7] DTI is a notoriously difficult underwriting criterion for third-party investors to verify. For example, without performing an additional investigation it is difficult for an investor to check if a lender ignored an obligation like alimony, or overstated sources of income.

[8] The CFPB reports that 77 percent of borrowers taking out a mortgage in 2013 applied to only one lender (`http://files.consumerfinance.gov/f/201501_cfpb_consumers-mortgage-shopping-experience.pdf`).

[9] Rejection rates are from (Bhutta et al., 2017).

mining what their DTI is prior to applying or how it may affect the outcome of their application at a particular lender.[10]

In our model, we will make the simplifying assumption that rejected borrowers exit the market entirely and are unable to re-apply at other lenders. In reality, a borrower's decision to re-apply at another lender will reflect both the pecuniary and non-pecuniary costs of re-applying in addition to borrower beliefs about the benefits of doing so. To the best of our knowledge there are no credible studies of how borrowers respond when a credit application is denied. However, there is substantial evidence that search costs in general may be non-trivial in the U.S. mortgage market. For example, the wide price dispersion documented in both Bhutta et al. (2018) and Alexandrov and Koulayev (2018) for GSE-eligible loans implies that borrowers do not effectively search across even the most commoditized space of mortgage products. [11] Bhutta et al. (2018) find that the difference in price paid between the 90th and 10th percentiles of the price distribution is equivalent to $7,500 in upfront costs, with even larger costs for lower FICO borrowers. Alexandrov and Koulayev (2018) find similar dispersion and argue it is at least in part facilitated by borrower's mistaken beliefs that there is no price dispersion in the mortgage market. While searching across prices and searching across underwriting criteria are clearly distinct, borrowers that mistakenly believe there is no price dispersion may also mistakenly believe there are no differences across lenders in underwriting criteria, and so may re-apply at lower rates than they should.

While direct estimates of borrowers' beliefs about the likelihood of approval are hard to come by, survey evidence indicates that this type of discouragement effect may be large. The Survey of Consumer Finances asks households if there was a time in the last five years when they thought about applying for credit but did not because they thought they might be denied. Using the pooled surveys from 1995-2013, we calculate that 45 percent of borrowers who were denied a loan sometime within the last 5 years also indicate having been discouraged from applying for new credit due to low perceived likelihood of approval. This is three to nine times larger than the same share among households who had either successfully applied for a loan in the last 5 years or not applied for one at all. Admittedly, the 45 percent number does not exactly match our setting as some of these borrowers may be accurately predicting their likely denial. However, given the opacity of the underwriting process, especially with respect to the DTI calculation and how it depends on the QM regulation, we think it is plausible that many of the borrowers who are rejected at one lender due to the regulation are unaware that they may find another lender that would approve their loan.

[10] While HMDA does not require lenders to report a reason for denial, DTI is the most frequently cited reason among those that do (Bhutta et al., 2017).

[11] Argyle et al. (2017) document similar price dispersion in the auto lending market.

C.1 Setup

We model a mortgage market in which potential borrowers interact with two types of lenders: portfolio lenders and mortgage companies. Portfolio lenders can originate loans and either sell them to investors or retain them in their own portfolios. Mortgage companies can also sell their loans to investors, but earn zero return on any loans they choose to retain. Loan pools originated by either type of lender are indexed by the borrower's DTI θ. Investors purchasing loans from lenders specify both the share of each originated pool they want to purchase $\sigma(\theta)$ and the interest rate they require on these loans $R(\theta)$. In exchange for these loans, investors offer a price $T(\theta)$ to the lender. Lenders take these prices as given and screen applicants on behalf of the investor. This relationship reflects the actual use of "rate sheets" in the mortgage market, whereby investors specify the rates they require as a function of the borrower characteristics on each loan they purchase. For simplicity, we assume that DTI does not affect the probability of default δ, and we ignore other risk factors like FICO as we condition on these in our empirical work.[12]

Borrowers

There is a unit mass of potential borrowers of each type θ that arrive to apply at each lender randomly every period. Thus, each "borrower" represents a pool of potential loan originations with a given DTI. Once a borrower has arrived at a lender, the lender will offer the borrower the investor's price $R(\theta)$. The borrower has a smooth demand curve $D(R)$ for loan size up until a reservation price $\overline{R}$, above which the borrower walks away without taking a loan. We denote the borrower's total demand function as:

$$\overline{D}(R) = \begin{cases} D(R) & \text{if } R \leq \overline{R} \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

This demand function is known to the investor, but neither the investor nor the borrower know the borrower's DTI until the borrower arrives at a lender to apply for a loan. If the borrower walks away from the lender they cannot search for a loan again.[13]

[12]Our assumption that DTI does not affect the probability of default could easily be relaxed and would not affect the results. We make this assumption for simplicity and because it is consistent with our empirical results on the relationship between DTI and default near the ATR cutoff.

[13]Our model is essentially static, but can be easily extended to multiple periods so that the borrower walking away can reapply next period. In this version of the model, the reservation price $\overline{R}$ could also be pinned down endogenously as a function of the cost of walking away and the expected value of being re-matched to a lender randomly again in the next period. This would not change the central results in any meaningful way.

Lenders

Lenders originate and screen loan applications on behalf of investors. If the investor's required price for a given DTI is below the borrower's reservation price, then lenders can choose to either accept a pool of loans outright and pay an origination cost of c_A, or they can further investigate and fully document the borrower's application, which costs $c_I > c_A$. We denote the lender's choice between accepting or investigating loans as $a \in \{A, I\}$. To capture how the ATR/QM rule affects the market, we will allow the lender's choice between accepting and investigating to potentially affect the loss given default in a way that depends on the borrower's DTI. Specifically, the return in default for loans with DTI θ is given by $R(\theta)(1-\rho(\theta,a))$, where $0 < \rho(\cdot,\cdot) < 1$. Crucially, we assume that investors cannot verify whether or not the lender investigated the loan and can therefore not condition either their payments to the lender or the rate offered to the borrower on this decision.

There are two types of lenders in the market: mortgage companies and portfolio lenders. Mortgage companies cannot hold loans on portfolio due to their small balance sheets and therefore earn zero return on any loan not sold to an investor. Suppressing the fact that R, σ, and T all depend on θ, the mortgage companies' payoffs are given by:

$$\Pi^{MC}(a;\theta) = \overline{D}(R)(\sigma T - c_a). \tag{5}$$

As long as $c_I > c_A$ and the profit from originating and securitizing a loan is positive, the mortgage company will always decide to accept a loan without investigation ($\Pi^{MC}(A;\theta) > \Pi^{MC}(I;\theta)$). This follows from the fact that neither the securitization rate σ nor the price of securitized loans T can be conditioned on whether the lender investigates the loans. This inability to verify whether or not the mortgage company has investigated a loan implies that the investor must always assume that mortgage companies will shirk by not paying the investigation cost.

Portfolio lenders can also sell some fraction of their loans to investors. Unlike mortgage companies, however, portfolio lenders are able to collect the expected return on any loans they choose not to sell.[14] The total payoffs to a portfolio lenders from both the loans they sell and those they retain are given by:

$$\Pi^{P}(a;\theta) = \overline{D}(R)(\sigma T + (1-\sigma)R(1-\delta\rho(\theta,a)) - c_a). \tag{6}$$

Because portfolio lenders are exposed to potential losses on a fraction of the loans they originate, the net returns to accepted and investigated loans will differ not only due to the cost of investiga-

[14] A simplified, but somewhat trivial version of this model would only allow portfolio lenders to retain loans on balance sheet and would yield similar results. We present this model because it highlights the role of the agency friction.

tion but also due to the loss given default. As before, the securitization rate, interest rate, and the price paid by the investors cannot vary with the investigation decision since it is not observable to the investor.[15]

Investors

Investors purchase loans from both mortgage companies and portfolio lenders while also setting the price at which they are willing to lend to borrowers. To capture the idea that investors generally have more diversified portfolios or access to better servicing technology, we assume that loans held by investors earn a higher expected return than loans held in portfolio. In particular, the expected return on loans held by investors is given by $\beta(R(1-\delta\rho(\theta,a)))$, where we assume that $\beta > 1$. This assumption implies that there are potential gains to trade between portfolio lenders and investors and will incentivize portfolio lenders to securitize a portion of the loans they originate.[16]

The investor's problem is to choose the securitization rate σ and the interest rate R that maximize the total surplus split between them and the lender given the lender's actions.[17] The surplus is split according to the transfer price T. In general, this transfer price is not uniquely determined and would be pinned down by the relative bargaining power of the lender and the investor, but we will verify its existence below.

In solving this problem, investors know which type of lender they are dealing with and can therefore set separate policies for mortgage companies and portfolio lenders. When dealing with portfolio lenders, investors are aware of the fact that the total surplus will depend on both the return on loans they purchase as well as those the lender chooses to retain. Thus, investors dealing with portfolio lenders maximize the following surplus:

$$S^{P}(\sigma,R;a,\theta)=\overline{D}(R)((\sigma\beta+1-\sigma)R(1-\delta\rho(\theta,a))-c_{a}). \tag{7}$$

The surplus that is split between investors and mortgage companies is similar, except for the fact that loans not sold to investors will earn zero return. Investors dealing with mortgage companies will therefore maximize the following surplus:

$$S^{MC}(\sigma,R;a,\theta)=\overline{D}(R)(\sigma\beta R(1-\delta\rho(\theta,a))-c_{a}). \tag{8}$$

[15] We assume also that the rate on loans retained by the lender cannot vary with the investigation decision. This is equivalent to assuming that the lender commits to selling a random share of a pool of identical loans to the investor. Cream-skimming would induce an additional agency friction distracting from the central question of how the ATR policy affected the market.

[16] If the return on securitized and portfolio loans were the same ($\beta = 1$), then there would be no reason for portfolio lenders to sell loans to investors. This version of the model yields similar results, but provides less useful insight into the nature of the agency problem.

[17] This is the same solution concept used by Bubb and Kaufman (2014).

C.2 *Equilibrium*

An equilibrium in this model is composed of a set of securitization policies $\sigma^{MC}(\theta), \sigma^{P}(\theta)$; interest rates $R^{MC}(\theta), R^{P}(\theta)$; and transfer prices $T^{MC}(\theta), T^{P}(\theta)$; such that each agent is maximizing their respective objective functions (5)–(8) given borrower demand (4). To characterize this equilibrium, we first solve for the investor's optimal securitization policies taking the interest rate as given. Given these securitization policies, we then solve for the optimal interest rate as a function of borrower demand. Finally, we show that there exist transfer prices that support this equilibrium.

Securitization Policies

For investors facing mortgage companies, the securitization rule is simple. Investors know that if investigation is costly ($c_I > c_A$), then mortgage companies will always choose to accept rather than investigate loans. As a result, the investor will securitize all loans so long as they have positive expected value conditional on the mortgage companies' lack of investigation:

$$\sigma^{MC}(\theta) = \begin{cases} 1 & \text{if } \beta R(1-\delta\rho(\theta,A)) \geq c_A \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

For investors dealing with portfolio lenders, the securitization policy is more complicated. Because portfolio lenders are exposed to losses on all loans not purchased by the investor, the investor can use the securitization rate to incentivize portfolio lenders to investigate loans when it is beneficial for them to do so. In particular, if there is a benefit to investigation ($\rho(\theta,I) < \rho(\theta,A)$) and if the cost to investigation is such that the surplus from investigating is higher than the surplus from not investigating ($S^P(\sigma,R;I,\theta) \geq S^P(\sigma,R;A,\theta)$), then the investor would like to lower the securitization rate to the point where the lender is just indifferent between investigating and accepting.[18] The securitization rate at which the portfolio lender is indifferent satisfies the following equality:

$$\sigma^P T + (1-\sigma^P)R(1-\delta\rho(\theta,A)) - c_A = \sigma^P T + (1-\sigma^P)R(1-\delta\rho(\theta,I)) - c_I.$$

Solving for σ^P gives the optimal securitization rate

$$\sigma^P(\theta) = 1 - \frac{c_I - c_A}{R\delta(\rho(\theta,A) - \rho(\theta,I))}. \tag{10}$$

[18] If there is no benefit to investigation ($\rho(\theta,I) \geq \rho(\theta,A)$), then the securitization rate for portfolio lenders is identical to the mortgage company's.

Because we assume $\rho(\theta, I) < \rho(\theta, A)$ and $c_I > c_A$, the second term in this expression is positive and implies that the investor must leave the portfolio lender with some skin in the game ($\sigma^P < 1$) to make them indifferent between investigating and accepting.

This expression provides useful intuition for how the agency conflict between the investor and the lender affects incentives. As the cost to investigation c_I increases, portfolio lenders will have less of an incentive to investigate loans and more of an incentive to shirk. To offset this incentive, investors can lower the securitization rate, which forces lenders to internalize the costs of not investigating by increasing the number of loans on their books. Conversely, because the lender is fully exposed to both gains and losses on any loans it retains, investors can increase the securitization rate when either the promised return on the loan R or the expected cost of not investigating $\delta(\rho(\theta, A) - \rho(\theta, I))$ increase. Intuitively, the investor does not need to force the lender to retain as many loans when the outcome of each retained loan becomes more consequential for the lender.

Demand and Loan Pricing

The optimal securitization policies we solved for in the previous section take the interest rate as given. In this section, we solve for the investor's optimal interest rates given these securitization policies and borrower demand.

Let $S^i(\sigma^i, R; a, \theta)$ denote the surplus from a relationship with lender type $i \in \{P, MC\}$ at generic interest rate R given the investor's optimal securitization rate for that lender type. The investor's objective is to choose the interest rate schedule $R^i(\theta)$ that maximizes this surplus separately for each lender type.

In general, the investor's optimal price for lender type i can fall into only one of three regions. If the investor offers a price below the borrower's reservation price, she will choose this price optimally to maximize total surplus given the smooth demand function $D(R)$. We denote this price as $R_i^*(\theta)$. Alternatively, if total surplus is higher at the borrower's reservation price, the investor will raise the price to $\overline{R}$. Finally, if costs are so high that lending is unprofitable even at the reservation price, then the investor will offer some indeterminate price above the reservation price knowing that borrowers will walk away. We interpret this third scenario as a denial, since there is no price the borrower would accept that the investor is willing to offer. Given this, the investor's pricing function can be expressed as follows:

$$R^i(\theta) = \begin{cases} R_i^*(\theta) & \text{if } S^i(\sigma^i, R_i^*(\theta); a, \theta) \geq S^i(\sigma^i, \overline{R}; a, \theta) \\ \overline{R} & \text{if } S^i(\sigma^i, R_i^*(\theta); a, \theta) < S^i(\sigma^i, \overline{R}; a, \theta) \text{ and } S^i(\sigma^i, \overline{R}; a, \theta) \geq 0 \\ R \in \left(\overline{R}, \infty\right) & \text{otherwise.} \end{cases}$$

Since we have solved for the optimal securitization rates already, the only thing remaining that needs to be determined to fully characterize investor pricing are the optimal interior prices $R_P^*(\theta)$ and $R_{MC}^*(\theta)$. We solve for these prices separately taking the optimal securitization policies as given. In doing so, we continue to restrict attention to the most interesting region of the parameter space where investigating a loan generates more surplus than simply accepting it outright.

The price offered through mortgage companies is straightforward since all loans offered by mortgage companies are securitized so long as the surplus is positive. Assuming positive surplus and taking into account the fact that mortgage companies will always choose to accept rather than investigate loans, the investor chooses the price to maximize the following

$$\max_R D(R)(\beta R(1-\delta\rho(\theta,A))-c_A),$$

which yields the solution

$$R_{MC}^*(\theta)=\frac{c_A}{\beta(1-\delta\rho(\theta,A))}-\frac{D}{D'}, \tag{11}$$

where D and D' are evaluated at $R_{MC}^*(\theta)$. Intuitively, the first term shows that investors in loans originated by mortgage companies will increase prices as either the cost of origination increases or as the cost/likelihood of default increase. The second term reflects the fact that there are limits to borrower search, which allows the investor to charge a standard monopolist's markup exploiting the shape of the demand curve.[19]

The expression for pricing at portfolio lenders is slightly more complicated. When there are benefits to investigation, the investor sets the securitization rate for portfolio lenders to ensure that they investigate all loans.[20] Because this optimal securitization rate is less than one, changes in the interest rate at portfolio lenders will end up affecting the total surplus through both the return on loans purchased by the investor and those retained by the lender. Plugging the investor's optimal securitization rate (10) into the objective function (7) and taking into account the fact that portfolio lenders will always choose to investigate given this securitization rate, the investor chooses interest rates to maximize

$$\max_R D(R)((\sigma^P\beta+1-\sigma^P)R(1-\delta\rho(\theta,I))-c_I).$$

[19] We assume that the derivative of the ratio D/D' with respect to R is negative for simplicity (this would be true, for example, if the demand function were linear). This ensures that demand effects do not swamp the direct effect of cost changes.

[20] If there are no benefits to investigation then the investor would choose the same securitization rate at both types of lenders and the expression for pricing at the portfolio lender would be identical to the mortgage company's.

This yields the following solution for the optimal interest rate at portfolio lenders:

$$R_P^*(\theta) = \frac{c_I}{\beta(1-\delta\rho(\theta,I))} - \frac{D}{D'} + \frac{\beta - 1}{\beta\delta(\rho(\theta,A)-\rho(\theta,I))}. \tag{12}$$

The price charged on loans from portfolio lenders is similar to that charged on loans at mortgage companies, but for the final term. This last term captures the premium due to agency costs. Since $\beta > 1$ (investors value the loan returns more), this additional term will increase the price on loans. The size of this premium will shrink, however, as the expected costs of not investigating $\delta(\rho(\theta,A)-\rho(\theta,I))$ grow. Intuitively, increasing the cost of not investigating will reduce the agency friction since portfolio lenders are exposed to losses on the loans they retain.

Transfer Prices

To close the model, we need to show that there exist transfer prices $T^{MC}(\theta)$ and $T^P(\theta)$ that support the securitization policies and interest rates described above. These transfer prices simply split the surplus between lenders and investors and will not, in general, be unique. They could be pinned down if we were to explicitly model the bargaining process between investors and lenders. However, because the transfer prices do not affect the real outcomes in the model, we do not take a stand on this bargaining process and instead simply prove existence.

In order for the investor to be willing to purchase any quantity of loans from a mortgage company the following inequality must hold

$$T^{MC}(\theta) \leq \beta R^{MC}(\theta)(1-\delta\rho(\theta,A)).$$

Similarly, the transfer price must also incentivize the mortgage company to originate and sell these loans to the investor, which means

$$T^{MC}(\theta) \geq c_A.$$

Putting these inequalities together, a valid transfer price between investors and mortgage companies exists so long as the following inequality holds:

$$\beta R^{MC}(\theta)(1-\delta\rho(\theta,A)) - c_A \geq 0.$$

Similarly, the transfer price between investors and portfolio lenders must satisfy

$$\beta R^P(\theta)(1-\delta\rho(\theta,I)) \geq T^I(\theta) \geq c_I.$$

In either case, if these inequalities are violated then not only will lenders and investors not trade loans, but neither type of lender would decide to originate any loans at all. This is obvious in the case of mortgage companies, since we assumed they cannot profitably hold loans on their balance sheet. With portfolio lenders, it follows from the fact that the portfolio lender's return on holding loans is always less than the investor's since $\beta > 1$. Thus, if the investor cannot hold the loans at a surplus then neither can the portfolio lenders

$$R^P(\theta)(1-\delta\rho(\theta,I)) < \beta R^P(\theta)(1-\delta\rho(\theta,I)) < c_I.$$

Therefore, as long as there is a positive surplus from originating loans there will also exist transfer prices supporting trade between the lenders and investors.

C.3 Effects of ATR Policy

The model outlined above can be used to understand how frictions in borrower search and agency conflicts in the secondary mortgage market may contribute to the empirical results we document. To demonstrate this, we start from a pre-ATR equilibrium in which the two types of lenders charge interest rates that are arbitrarily close to each other and portfolio lenders choose to investigate all loans.[21] We then model the introduction of the ATR rule as an increase in lender costs along two dimensions. First, the law required more stringent documentation on all loans. We model this as an increase in the cost of investigation from c_I to $c_I + \xi_c$. Second, the law increased the expected cost of default on non-QM loans for which the lender cannot prove that they arrived at a reasonable, good faith determination of the borrower's ability-to-repay. We model this as an increase in the loss given default conditional on not investigating for all loans with DTIs greater than some threshold θ^{ATR}. Specifically, we assume that the loss in default conditional on not investigating increases from $\rho(\theta,A)$ to $\rho(\theta,A)+\xi_\rho$ for all loans with DTI $\theta > \theta^{ATR}$.[22] We are interested in how these changes affect prices and quantities in the potentially affected segment of the market $\theta > \theta^{ATR}$.

We summarize the effects of the policy change in Proposition 1. In this proposition, variables subscripted by a 0 refer to pre-ATR equilibrium values and variables subscripted by "ATR" refer to equilibrium values after the policy change.

[21]The parameter restrictions we impose preclude the existence of an equilibrium in which both lenders charge identical prices. However, the difference in prices can be made arbitrarily small with an appropriate choice of parameters as we outline in the proof of Proposition 1 below.

[22]For simplicity, we assume that default costs on investigated loans, $\rho(\theta,I)$, remain unchanged for all DTIs. This assumption could be relaxed as long as the change in default costs for investigated loans was smaller than that for non-investigated loans.

Proposition 1. *Consider an intitial equilibrium where* $|R_0^P(\theta) - R_0^{MC}(\theta)| \leq \epsilon_R$ *for any* $\epsilon_R > 0$. *There exist cutoffs* ξ_ρ^* *and* ξ_c^* *such that* $\forall \xi_\rho > \xi_\rho^*$ *and* $\forall \xi_c < \xi_c^*$

1. *Mortgage companies exit the non-QM market:*

$$S_{ATR}^{MC}(\sigma, R; a, \theta) \leq 0 \quad \forall \theta > \theta^{ATR}.$$

2. *Portfolio lenders investigate all loans and do not exit the market:*

$$S_{ATR}^{P}(\sigma, R; I, \theta) \geq S^P(\sigma, R; A, \theta) \geq 0 \quad \forall \theta.$$

3. *The change in the average accepted interest rate,* $\Delta R(\theta)$, *for non-QM loans is approximately equal to the change the portfolio lender's price:*

$$\lim_{\epsilon_R \to 0} \Delta R(\theta) = R_{ATR}^P(\theta) - R_0^P(\theta) \quad \forall \theta > \theta^{ATR}.$$

4. *The change in the aggregate quantity of non-QM lending is given by*

$$\Delta Q = -Q_0^{MC} + Q_0^P \epsilon_D \Delta R(\theta),$$

where Q_0^i *denotes the initial quantity of non-QM lending by lender type* i *and* ϵ_D *is the elasticity of demand implied by the smooth portion of the borrower's demand function* $D(R)$.

Intuitively, this proposition states that as long as the increase in the cost of default conditional on not investigating is sufficiently large and the increase in the overall cost of investigating is not too large, then mortgage companies will be driven out of the market, the observed price response will be driven entirely by the change in prices at portfolio lenders, and the quantity response will primarily reflect the extensive margin loss of loans due to exiting mortgage company lenders. We describe the intuition for these results below before providing the proof.

Prices

The effect of the regulation on lender pricing is demonstrated most clearly by considering the smooth pricing equations (11) and (12). Focusing first on portfolio lenders, if we plug the new origination and default costs into the pricing equation (12), we can see that the lender's new price for any borrower with DTI $\theta > \theta^{ATR}$ will be given by[23]

$$R_{ATR}^P(\theta) = \frac{c_I + \xi_c}{\beta(1 - \delta\rho(\theta, I))} - \frac{D}{D'} + \frac{\beta - 1}{\beta\delta(\rho(\theta, A) - \rho(\theta, I) + \xi_\rho)}.$$

[23]Using equation (12) implicitly assumes that it is never surplus maximizing for the portfolio lender to switch from investigating to not investigating loans. This will be true as long as the increase in the cost of investigating ξ_c is sufficiently small.

This expression highlights the two primary effects of the regulation on pricing at portfolio lenders. First, because the cost of investigation has gone up, prices will rise. This is reflected in the first term, which is increasing in ξ_c. Second, because the default cost conditional on not investigating has gone up, and because portfolio lenders are exposed to losses on loans they retain, the agency conflict between the portfolio lender and investor will weaken subsequent to the policy change. This effect will cause the price to fall and is reflected in the third term, which is decreasing in ξ_ρ.

Both of these effects are illustrated in Panel A of Figure A.6 which plots the price response as a function of the size of the default cost shock under reasonable assumptions about the model's functional forms and assuming a fixed increase in the cost of investigation. In this figure, the lender's pre-ATR price is marked by the horizontal line at R_0^P. When there is no change in the default cost ($\xi_\rho = 0$), prices at portfolio lenders will unambiguously increase relative to the pre-ATR price. The size of this price increase is governed by the increase in the cost of investigation ξ_c and is reflected in the upward shift from the dotted to the solid blue line. However, as the size of the default cost grows, prices at portfolio lenders will decrease due to the reduced agency frictions. This effect is reflected in the negative slopes of both blue lines. For sufficiently large increases in the default cost, the latter effect will dominate, which would lead prices to actually fall below their pre-ATR level. We view this case as relatively unlikely, which is why prices are always higher than their pre-ATR level for the range of parameter values we plot. Nonetheless, these results imply that the increase in prices at portfolio lenders should be bounded above by the size of increase in the cost of documentation and may even be smaller than this in some cases.[24]

In contrast, the price response at mortgage companies is unambiguously positive and has no such upper bound. This can be seen by plugging the new cost parameters into equation (11), which yields the following expression for mortgage company pricing:

$$R_{ATR}^{MC}(\theta) = \frac{c_A}{\beta(1-\delta(\rho(\theta,A)+\xi_\rho))} - \frac{D}{D'}.$$

Since mortgage companies cannot commit to investigating loans, increases in the cost of investigation have no effect on prices. Moreover, because they never investigate loans, increases in default costs conditional on not investigating lead to strictly higher mortgage company pricing. This effect is reflected in the first term of the expression above, which is increasing in ξ_ρ. For sufficiently large increases in default costs, R_{ATR}^{MC} will exceed the borrower's reservation price and lending by mortgage companies will collapse entirely.

The orange line in Panel A of Figure A.6 demonstrates this possibility. The mortgage company's price will increase with the cost of default until it hits the reservation price $\overline{R}$. Eventually,

[24]For the extreme case in which the portfolio lender is unable to securitize any loans at all, the price change would be governed exclusively by the change in investigation costs.

lending even at this price becomes unprofitable because the price cannot increase to compensate for further increases in default costs without pushing the borrower out of the loan entirely. When this occurs, mortgage companies will post indeterminate and arbitrarily high interest rates above the borrowers reservation price and no lending will occur. This scenario is indicated in the figure by the wavy orange line above the borrower's reservation price. The vertically dashed line at ξ_ρ^{ATR} denotes a shock that we think reflects the actual ATR policy. The shock is large enough that the price of loans at portfolio lenders increases modestly, while mortgage companies drop out of the market entirely.

Quantities

Panel B of Figure A.6 turns to loan quantities and illustrates the effect of the policy given the shock ξ_ρ^{ATR} marked in Panel A. The y-axis plots prices while the x-axis plots loan quantities. The light blue line is the borrower's demand curve, which discretely jumps to zero at the borrower's reservation price. Prior to the policy change, both lenders charge the same price ($R_0^P \approx R_0^{MC}$) and therefore originate the same quantity of loans. The effect of the ATR policy is reflected in the two upward shifts in lender prices. The portfolio lender price increases moderately to the horizontal solid blue line. This induces a moderate decline in lending quantities that is governed by the slope of the demand curve. At the same time, the price at mortgage companies (the wavy orange line) has increased beyond the reservation price, which leads to a complete collapse in loan quantities for these lenders. Since we assume that borrowers cannot substitute across lender types after arriving to apply for a loan, the aggregate quantity response in this case would be given by the change in quantities at each type of lender weighted by their appropriate pre-ATR market shares. The observed price response, however, would simply be the change in the portfolio lender's price since the new rates at mortgage companies are never observed.

Together, these results are able to rationalize our central empirical findings. The large price increase driving out most of the lending volume is not observed because it is never accepted by the borrower. We interpret these missing loans as outright rejections, which are responsible for the large reduction in lending on the extensive margin. The lenders who remain in the market, however, are only subject to a moderate increase in costs and therefore do not raise prices as much. This smaller price increase is empirically observed and leads to a much smaller intensive margin reduction in lending. While substitution between lenders due to denials would undo some of the extensive margin response, the large number of missing loans we document and the heterogeneity of this response across lender types suggest that it was not sufficient to undo the differential impact of the regulatory shock on non-retail and non-portfolio lenders. In this way, we think the model parsimoniously captures how differences in lender business models and frictions in borrower search can interact to affect the impact of a seemingly straightforward regulatory change leading it to have unexpectedly large real effects.

Proof of Proposition 1

First we show that prices at portfolio lenders and mortgage companies can be made arbitrarily close to each other in the initial equilibrium by an appropriate choice of parameters. This is trivially true if both lenders are charging the reservation price. If the lenders are not charging the reservation price we can push the difference to be as small as needed. Consider the following absolute difference, which is simply a re-arrangement of the absolute difference of lender prices in equations (11) and (12)

$$| R^*_{MC} - R^*_P | + \left| \frac{D(R^*_{MC})}{D'(R^*_{MC})} - \frac{D(R^*_P)}{D'(R^*_P)} \right| \leq \left| \frac{c_A}{\beta(1-\delta\rho(\theta,A))} - \frac{c_A}{\beta(1-\delta\rho(\theta,A))} + \frac{\beta-1}{\beta\delta(\rho(\theta,A)-\rho(\theta,I))} \right|.$$

Clearly, the right-hand side of this expression can be made as small as possible by increasing β or decreasing the difference between the costs of default across the two actions $(\rho(\theta,A)-\rho(\theta,I))$, and this will still preserve the parameter assumptions we require for the particular equilibrium actions we examine above. This implies that the lender prices can be made arbitrarily close to each other in the initial equilibrium

Given an equilibrium where prices are arbitrarily close, we can solve for the increase in ATR default penalties at which the mortgage company exits the market (i.e. the point at which the mortgage company surplus becomes negative). Under our assumption that the derivative of D/D' with respect to R is positive, it follows that $\lim_{\xi_\rho \to \overline{\xi}_\rho} R^*_{MC} = \infty$ where $\overline{\xi}_\rho \triangleq \frac{1}{\delta} - \rho$. Because R^*_{MC} is unbounded, it follows that there is a point at which the lender will have to charge the reservation price $\overline{R}$. We next examine the mortgage company's surplus function at the reservation price and show that it falls below zero for sufficiently large ξ_ρ. The mortgage company's surplus at the reservation price is zero when[25]

$$\sigma\beta\overline{R}(1-\delta\rho(\theta,A)+\xi_\rho)-c_A = 0 \Rightarrow$$
$$\xi_\rho = \xi^*_\rho \triangleq \frac{\left(1-\delta\rho(\theta,A)-\frac{c_A}{\beta\overline{R}}\right)}{\delta}.$$

Since surplus is clearly decreasing in ξ_ρ, it must be true that for any $\xi_\rho > \xi^*_\rho$, the surplus of

[25]This expression assumes that all mortgage company loans issued at the reservation price would be securitized in equilibrium. The only other possible equilibrium according to equation (9) is for no mortgage company loans to be securitized. In that equilibrium mortgage company surplus is already equal to zero.

originating loans through the mortgage companies is negative for borrowers at the affected DTIs.

We next show there is a cutoff for the shock to the cost of investigating loans ξ_c below which the portfolio lender will always continue to operate. Again, given the assumptions about the demand curve, it is clear that the portfolio lender's smooth price solution (12) is increasing in ξ_c so long as it continues to be optimal to investigate, which will be true in the space we are examining where operating without investigating has a negative surplus. Allowing for the cost shock to have driven up the portfolio lender's price to the reservation price, the surplus for the portfolio lender will be zero at the following point (see (7))

$$\xi_c^* \triangleq (\sigma\beta + 1 - \sigma)\bar{R}(1 - (1 - \delta\rho(\theta, I))) - c_I.$$

Thus, the portfolio lender will continue to operate while investigating loans so long as $\xi_c \leq \xi_c^*$.

Therefore, assuming that $\xi_c \leq \xi_c^*$ and $\xi_\rho > \xi_\rho^*$, the change in the average accepted price will simply be the change between the post-ATR price at the portfolio lender and the market-share weighted average of the mortgage company and portfolio lender's initial prices. Let ϵ_R be the initial difference in prices between portfolio and mortgage company lenders. Denoting pre-ATR values with a 0 subscript and post-ATR values with the subscript "ATR" we have

$$\begin{aligned}
\Delta R &= \frac{Q_0^P}{Q_0}(R_{ATR}^P - R_0^P) + \frac{Q_0^{MC}}{Q_0}(R_{ATR}^P - R_0^{MC}) \\
&\leq \frac{Q_0^P}{Q_0}(R_{ATR}^P - R_0^P) + \frac{Q_0^{MC}}{Q_0}(R_{ATR}^P - R_0^P + \epsilon_R) \\
&= (R_{ATR}^P - R_0^P) + \frac{Q_0^{MC}}{Q_0}\epsilon_R,
\end{aligned}$$

where Q_0^i is the total amount of debt originated by lender type i before the policy change and $Q_0 = Q_0^P + Q_0^{MC}$ is the aggregate amount of debt issued by both lender types. Thus, the change in average accepted prices converges to the change in the portfolio lender's price as ϵ_R approaches zero.

The quantity change when $\xi_c \leq \xi_c^*$ and $\xi_\rho > \xi_\rho^*$ is straightforward to compute. The mortgage company exits the market entirely while the change in borrowing from portfolio lenders simply follows the demand curve:

$$\Delta Q = -Q_0^{MC} + Q_0^P \epsilon_D \Delta R.$$

This decline in quantities is largely driven by the exit of mortgage companies from the market, with minimal contribution from the demand elasticity of borrowers and the increase in prices at portfolio lenders.

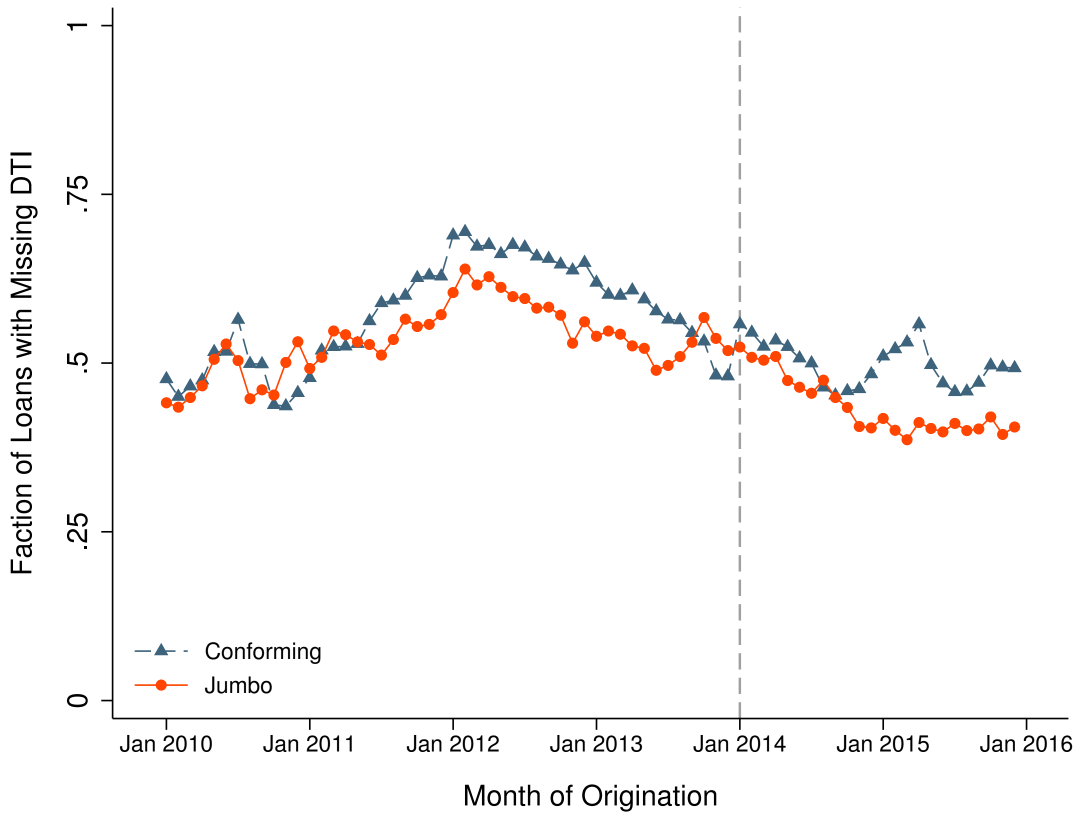

FIGURE A.1

Fraction of Loans with Missing DTI by Month of Origination

NOTE.—This figure plots the share of loans that are dropped from our analysis sample due to having a missing DTI. Shares are reported separately for jumbo and conforming loans and by month of origination. The vertically dashed grey lines marks the month that the Ability-to-Repay Rule and Qualified Mortgage Standards went into effect (January 2014).

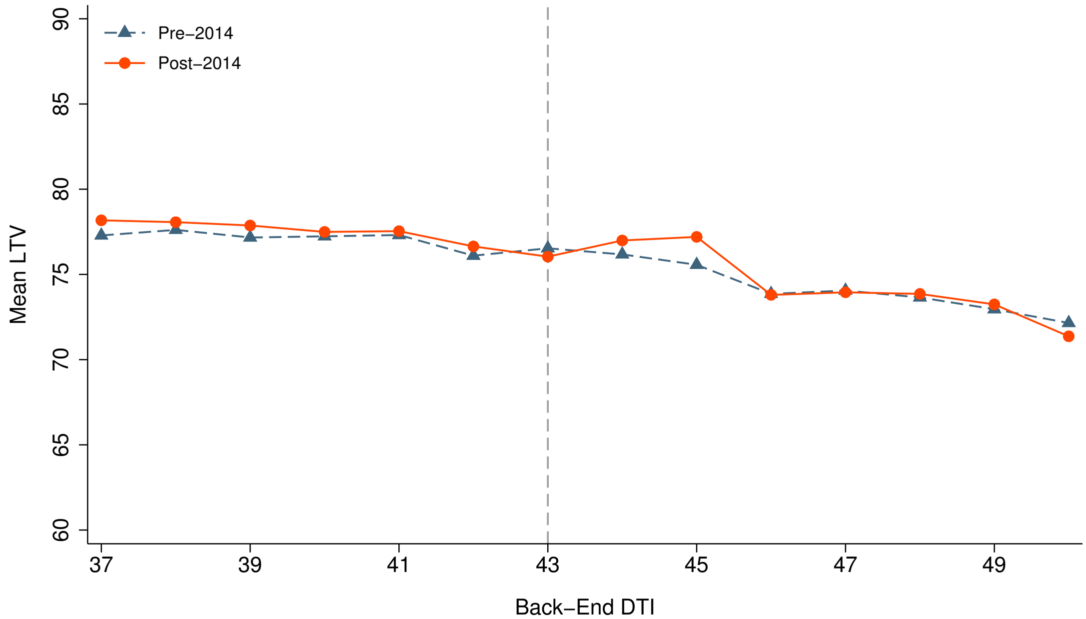

Panel A. Mean LTV by DTI

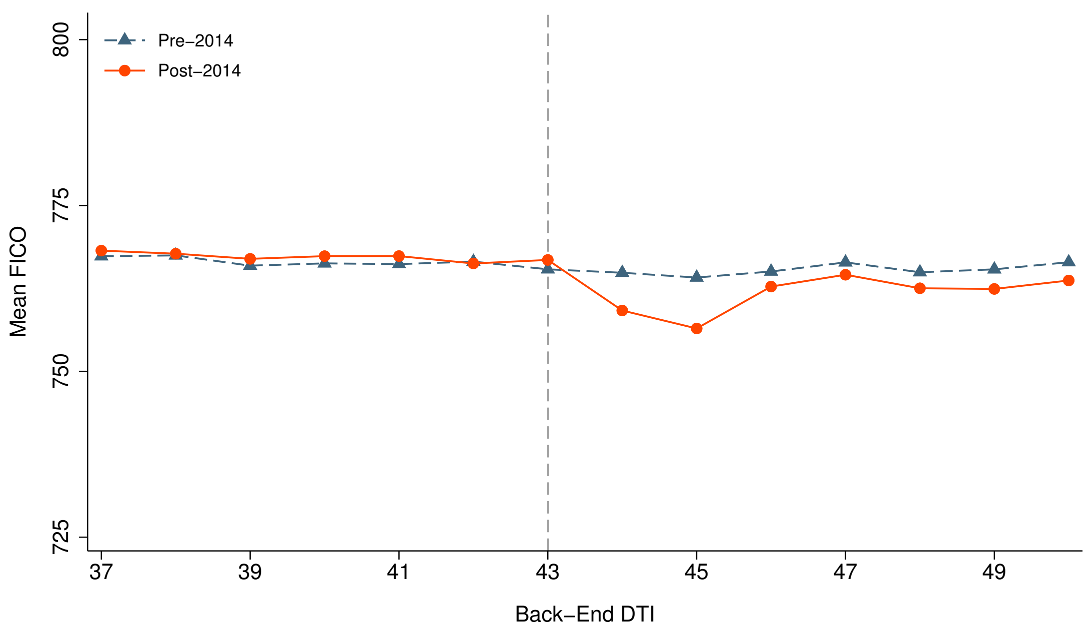

Panel B. Mean FICO by DTI

FIGURE A.2

Changes in Borrower Characteristics following the Implementation of ATR/QM

NOTE.—This figure plots mean borrower LTV ratios and FICO scores by DTI for loans originated before (blue triangles) and after (orange circles) the implementation of ATR/QM. Each dot represents the raw average LTV or FICO score for borrowers in the indicated one-percent DTI bin and time period. The vertically dashed grey line marks the QM threshold of 43 percent. DTI bins are created by rounding up to the nearest integer so that the 43 percent bin inclused all DTIs greater than 42 percent and less than or equal to 43 percent. Means are calculated using the sample of all jumbo loans with DTIs between 36 and 50 percent described in Section III.

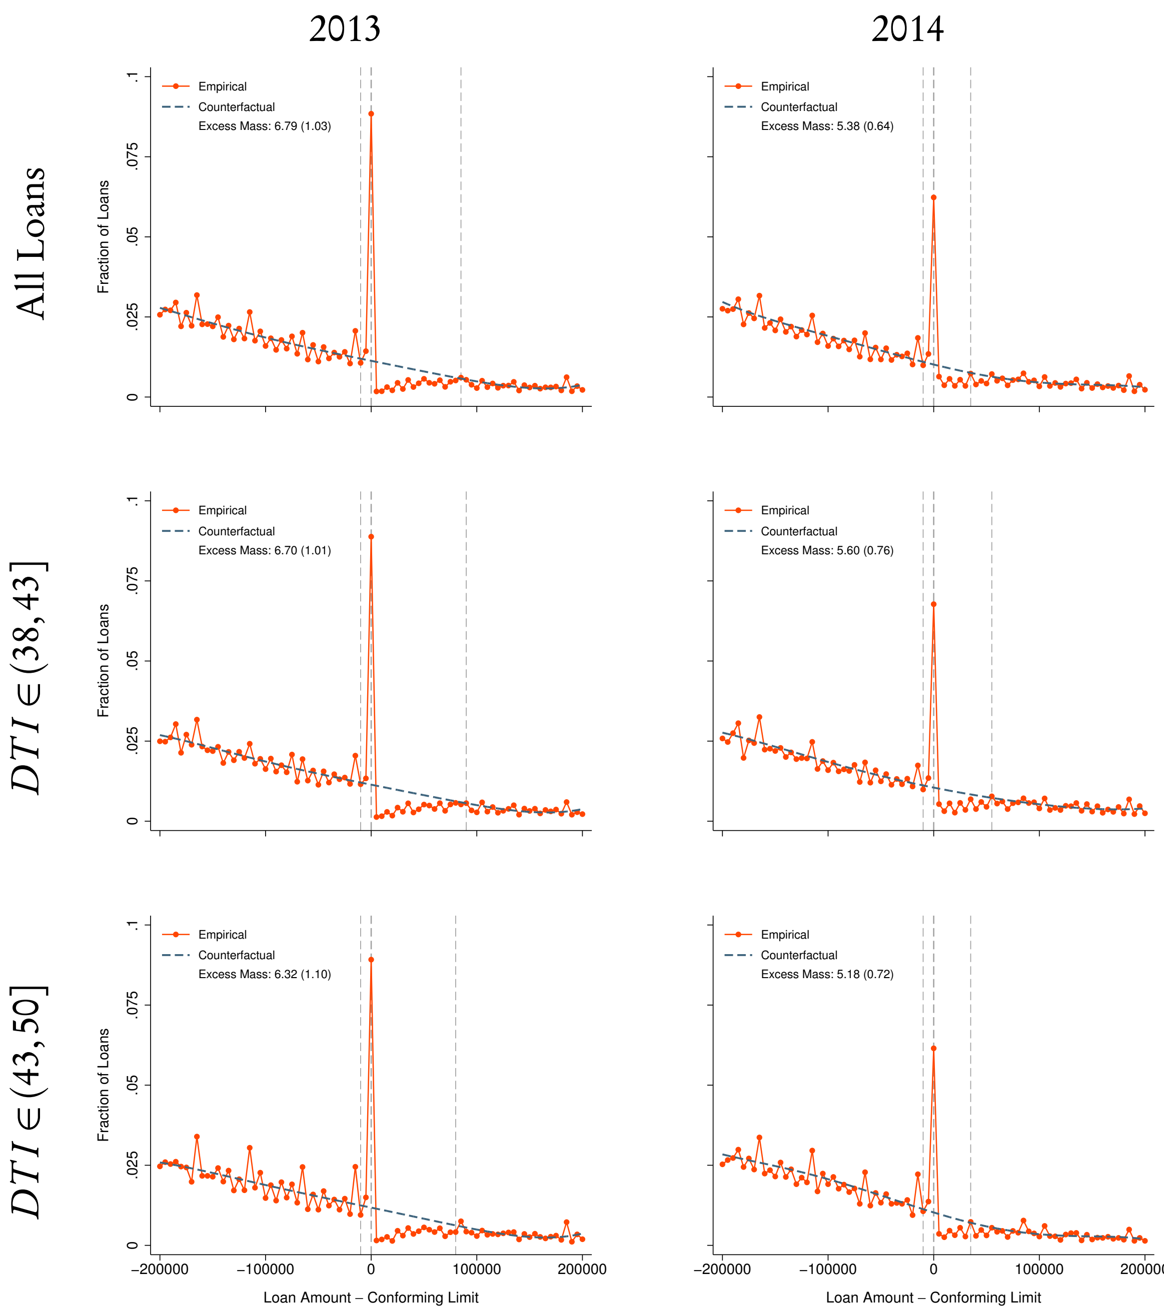

FIGURE A.3
Bunching at the Conforming Limit before and After ATR/QM

NOTE.—This figure plots the empirical and counterfactual density of loan size relative to the conforming limit by origination year and borrower DTI. In each panel, the connected line plot represents the fraction of loans in a given $5,000 bin relative to the conforming limit in effect at the time of origination. The heavy dashed line is the estimated counterfactual density obtained by fitting a 5th degree polynomial to the bin counts, omitting the contribution of the bins in the region marked by the vertical dashed gray lines. The figure also reports the estimated excess mass at the conforming limit and its standard error, calculated as described in Appendix B.2.

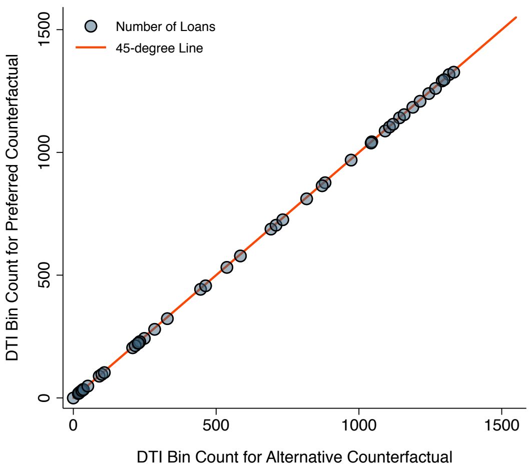

Panel A. Lower Limit: $\overline{d} = 30$

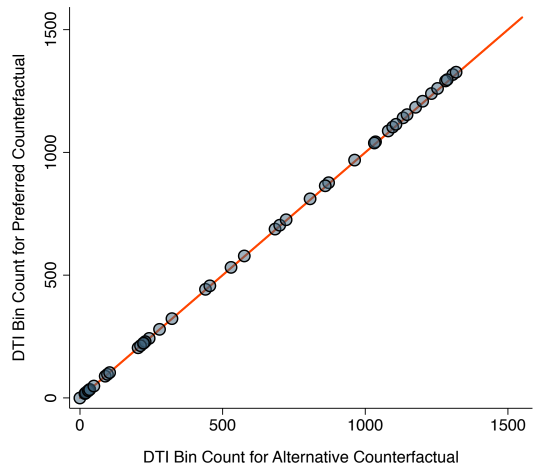

Panel B. Lower Limit: $\overline{d} = 35$

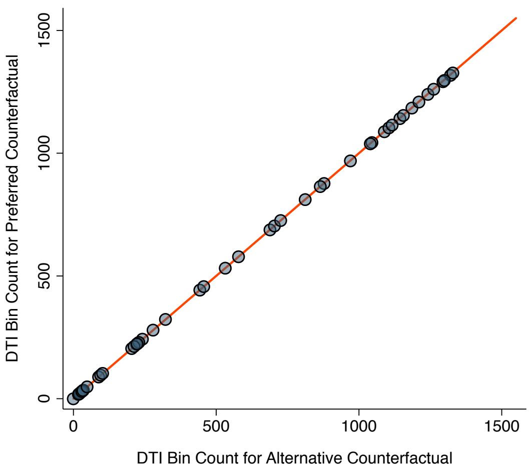

Panel C. Lower Limit: $\overline{d} = 40$

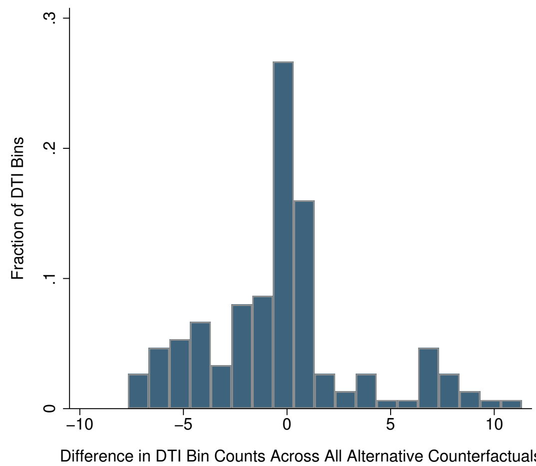

Panel D. All Pairwise Differences from $\overline{d} = 38$

FIGURE A.4

Differences in Counterfactual DTI bin Counts

Across Alternative Choices for the Lower Limit of the Bunching Region

NOTE.—This figure shows how changing the lower limit of the bunching region, $\overline{d}$, affects the estimated counterfactual DTI distribution. Panel A plots the number of loans in a given DTI bin from our preferred counterfactual, which sets $\overline{d} = 38$, on the x-axis against the number of loans in the same bin using an alternative counterfactual estimated by setting $\overline{d} = 30$. The solid orange line is the 45-degree line. Panels B and C plot analogous results using alternative counterfactuals that set $\overline{d}$ to 35 and 40 percent respectively. Panel D pools across all three alternative counterfactuals and plots the distribution of pairwise differences between the number of loans in a given DTI bin from the preferred counterfactual and the number of loans in that same bin across each of the alternatives.

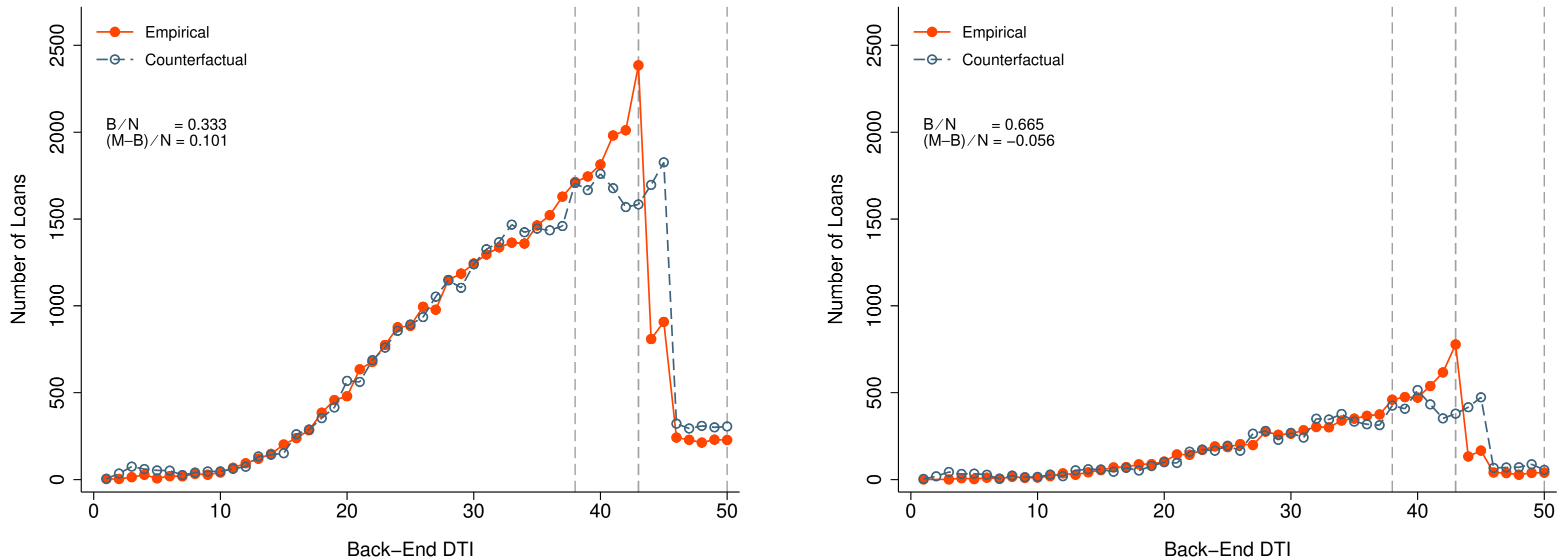

Panel A. FRMs and ARMs Combined

Panel B. ARMs Only

FIGURE A.5

Bunching, Missing Mass, and the Effect of ATR/QM on the Quantity of Credit by Product Type

NOTE.—This figure plots the empirical and counterfactual DTI distribution for jumbo mortgages by product type for loans originated in 2014, the first year that ATR/QM was in effect. Panel A. includes both fixed- and adjustable-rate mortgages whereas Panel B. is restricted to the sample of adjustable-rate mortgages only. The solid orange connected line is the empirical distribution. Each dot represents the number of loans of the indicated type originated in 2014 for which the borrower's DTI fell into the one-percent bin indicated on the x-axis. DTI bins are created by rounding up to the nearest integer so that the 43 percent bin includes all DTIs greater than 42 percent and less than or equal to 43 percent. The dashed blue connected line plots the counterfactual, which was estimated as described in Section V.A using 2013 as the pre-period. The vertically dashed grey lines mark the lower limit of the bunching region ($\overline{d} = 38$), the QM-threshold, and the maximum DTI.

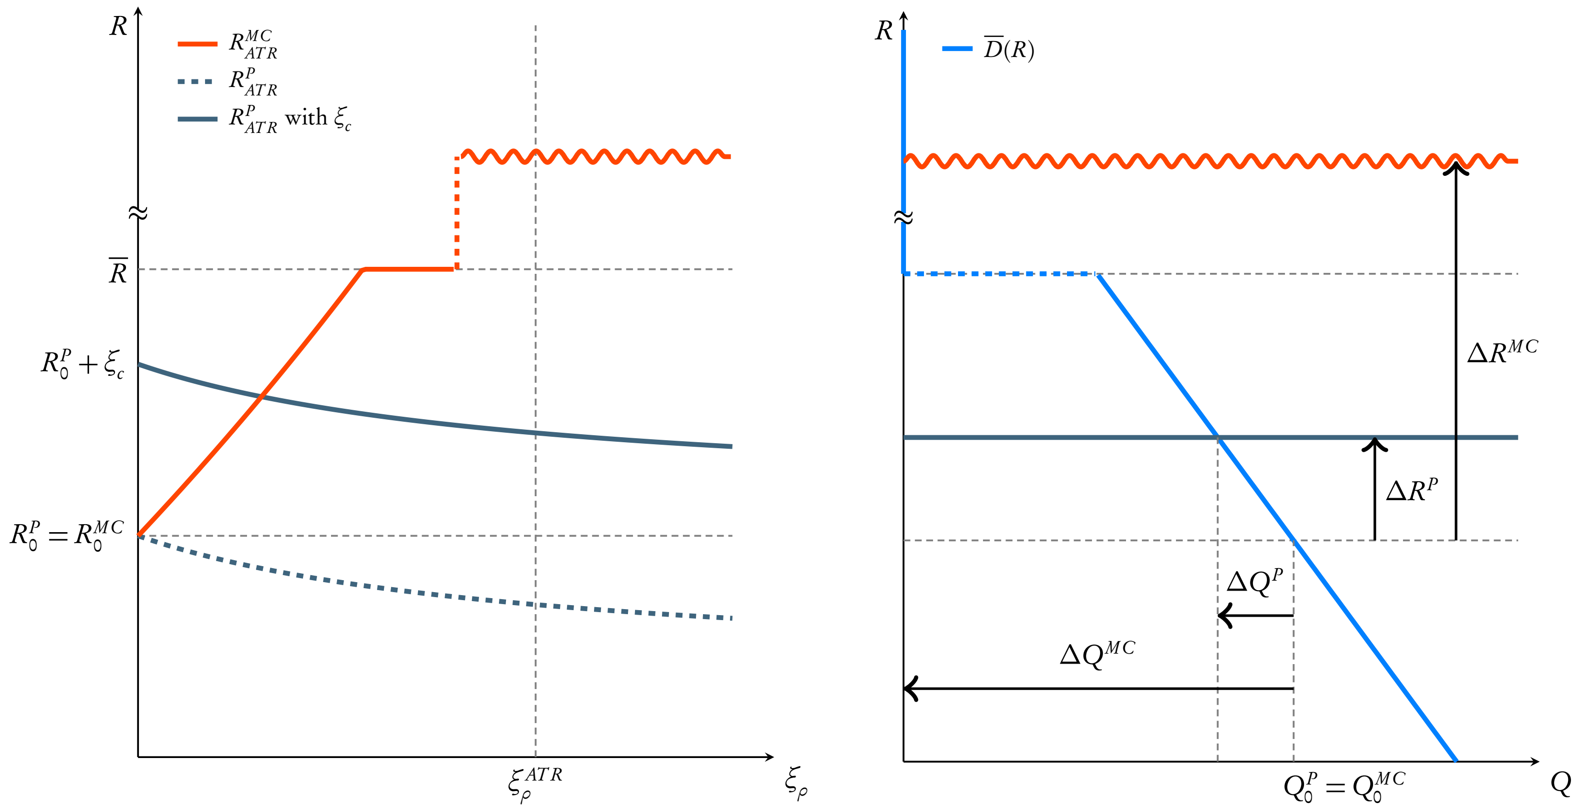

Panel A. Optimal Pricing as a Function of ATR Default Cost ξ_ρ

Panel B. Quantity Response for Fixed ATR Default Cost ξ^{ATR}_ρ

FIGURE A.6

Effects of ATR of the Price and Quantity of Credit

NOTE.—This figure illustrates how the ATR policy affects the price and quantity of credit originated by portfolio lenders and mortgage companies. Panel A plots the equilibrium interest rate for each type of lender as a function of the size of the increase in default costs conditional on not investigating loans (ξ_ρ) and the increase in the cost of origination conditional on investigating (ξ_c). Panel B plots the change in quantities associated with these equilibrium price changes for a given default cost ξ^{ATR}_ρ, marked with a vertically dashed line in Panel A.

TABLE A.1
COVARIATE BALANCE ACROSS LOANS WITH MISSING AND NON-MISSING DTIs

	DTI Non-Missing	DTI Missing	Difference in Means	Standardized Difference
	(1)	(2)	(3)	(4)
FICO Score	756.119	756.477	-0.357***	-0.008
	(43.275)	(42.432)	(0.053)	
Loan Amount ($1000's)	264.577	256.808	7.770***	0.041
	(189.747)	(185.890)	(0.233)	
Loan-to-Value	80.341	80.860	-0.518***	-0.037
	(13.891)	(13.873)	(0.017)	
Interest Rate	4.292	4.294	-0.002***	-0.004
	(0.561)	(0.527)	(0.001)	
Percent Condo	11.393	11.412	-0.019	-0.001
	—	—	(0.039)	
Number of Observations	1,195,895	1,428,985	2,624,880	2,624,880

NOTE.—This table reports loan-level descriptive statistics for both our full analysis sample (column 1) and the sample of loans that are excluded from our analysis for having a missing DTI (column 2). The first two columns report sample means along with their standard deviations in parentheses. Column 3 reports the difference in means between columns 1 and 2 along with its standard error. Column 4 reports the standardized difference in means. For continuous variables, this is calculated as the difference in means divided by the average of the standard deviations in each sample. In the case of the condo indicator, which is a binary variable, it is calculated as $(\overline{x}_1-\overline{x}_1)/\sqrt{\overline{x}_1(1-\overline{x}_1)+\overline{x}_2(1-\overline{x}_2)/2}$, where $\overline{x}_i$ is the mean in sample i. Significance levels for the difference in means reported in column 3 of 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE A.2

CHANGES IN BORROWER CHARACTERISTICS FOLLOWING THE IMPLEMENTATION OF ATR/QM

	LTV		FICO	
	(1)	(2)	(3)	(4)
DTI > 43	-2.025***	-1.992***	-1.521***	-1.533***
	(0.113)	(0.113)	(0.369)	(0.356)
DTI > 43 × Post	0.407**	0.334	-4.908***	-4.905***
	(0.207)	(0.208)	(0.832)	(0.807)
Month FEs	X	X	X	X
County FEs	X	X	X	X
FICO × Post FEs		X		
LTV × Post FEs				X
Pre-Period Mean	74.7	74.7	765.0	765.0
Number of Observations	62,748	62,748	62,748	62,748

NOTE.—This table reports estimates from difference-in-differences regressions examining changes in borrower characteristics among high-DTI jumbo loans subsequent to the implementation of ATR/QM. Each column reports a separate regression estimated at the loan level using the indicated borrower characteristic (LTV or FICO) as the dependent variable. The sample includes all jumbo loans with DTIs between 36 and 50 percent. Coefficient estimates are reported for the non-QM "treatment" dummy ($DTI > 43$) as well as its interaction with an indicator for whether the loan was originated in a month following the implementation of ATR/QM ($Post$). Column 2 includes a full set of fixed effects for the borrower's FICO score (20-point bins) interacted with the $Post$ dummy. Similarly, column 3 includes a full set of fixed effects for the borrower's LTV (5 point bins) interacted with the $Post$ dummy. The first row of the bottom panel reports the pre-period mean of the dependent variable among high-DTI jumbo loans. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

TABLE A.3

HOW ALTERING THE COUNTERFACTUAL BY CHANGING $\overline{d}$ AFFECTS THE INTENSIVE AND EXTENSIVE MARGIN ESTIMATES OF ATR/QM ON THE QUANTITY OF CREDIT

	Preferred Estimates	Effect of Changing $\overline{d}$ Holding the Bunching Region Constant		
	$\overline{d} = 38$	$\overline{d} = 30$	$\overline{d} = 35$	$\overline{d} = 40$
$\hat{B}/\hat{N}^{post}_{44+}$	0.208	0.201	0.222	0.204
$(\hat{M} - \hat{B})/\hat{N}^{post}_{44+}$	0.154	0.166	0.136	0.158

NOTE.—This table reports estimates of the intensive and extensive margin effects of the Ability-to-Repay Rule and Qualified Mortgage standards on the quantity of credit in the jumbo mortgage market. The top row reports the estimated intensive margin effect of the regulation on the allocation of credit across the DTI distribution. Each estimate represents the fraction of jumbo loans in the counterfactual no-policy distribution that were shifted from a DTI above the QM-threshold of 43 percent to below the threshold. The second row reports the estimated extensive margin effect of the policy on the total number of jumbo mortgages originated. Each estimate represents the fraction of the counterfactual number of jumbo loans that were eliminated as a result of the policy. Column one reports our preferred estimates, which set the lower limit of the bunching region to $\overline{d} = 38$. Columns 2–4 report analogous estimates from alternative specifications which set this limit to 30, 35, and 40 percent when constructing the counterfactul, but hold the lower limit constant at 38 when calculating the number of loans bunching below the limit or missing from above. All specifications use 2013 as the pre-period and 2014 as the post period. The sample therefore includes all jumbo loans that were originated in either 2013 or 2014.

TABLE A.4
ESTIMATES OF THE EFFECT OF DTI ON THE 2008 FIVE-YEAR PROBABILITY OF DEFAULT

	All Loans	Full Documentation	Low Documentation
	(1)	(2)	(3)
DTI $\leq$ 38	-0.0706***	-0.0709***	-0.0695***
	(0.0038)	(0.0052)	(0.0061)
DTI > 43	0.0320***	0.0384***	0.0227***
	(0.0044)	(0.0056)	(0.0069)
Implied Aggregate Effect	-0.0040***	-0.0045***	-0.0032***
	(0.0006)	(0.0006)	(0.0006)
High-DTI bin share	0.23	0.23	0.23
Medium-DTI bin share	0.19	0.19	0.19
Number of Observations	91,493	58,748	30,415

NOTE.— This table reports estimates of five-year default probabilities for high-DTI and low-DTI loans relative to loans in the omitted category $DTI \in (38, 43]$ for loans originated in 2008. The relationship between DTI and default probability is estimated separately by loan documentation status and includes both jumbo and conforming loans. Column 1 reports results pooling across all loans. Column 2 restricts to a sample of full documentation loans and column 3 reports results for low documentation loans only. The third row of each column also reports the implied counterfactual effect of the ATR/QM rule on the aggregate default rate estimated as described in Section VI. Estimates of the relative default probabilities are derived from a regression of whether a loan defaulted on DTI-bin dummies, fixed effects for the month of origination, the county the property was located in, the type of property as well as 20-point FICO score bins, 5-point LTV bins and the pairwise interaction between the two. We define a loan as having defaulted if the borrower was ever more than 90 days delinquent or if the property was repossessed by the lender (foreclosure or REO) within five years of the origination date. Standard errors are reported in parentheses and are clustered at the county level. Significance levels 10%, 5%, and 1% are denoted by *, **, and ***, respectively.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?," *The Quarterly Journal of Economics*, 2018, *133* (1), 129–190.

Alexandrov, Alexei and Sergei Koulayev, "No Shopping in the U.S. Mortgage Market: Direct and Strategic Effects of Providing Information," 2018. CFPB Working Paper Series 2017-01.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer, "Real Effects of Search Frictions in Consumer Credit Markets," 2017. MIT Sloan Research Paper No. 5242-17.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo, "Paying Too Much? Price Dispersion in the US Mortgage Market," 2018. Working Paper.

_ , Steven Laufer, and Daniel R Ringo, "Residential mortgage lending in 2016: Evidence from the Home Mortgage Disclosure Act data," *Fed. Res. Bull.*, 2017, *103*, 1.

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

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *The Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

DeFusco, Anthony A. and Andrew Paciorek, "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit," *American Economic Journal: Economic Policy*, 2017, *9* (1), 210–240.

Foote, Christopher, Kristopher Gerardi, Lorenz Goette, and Paul Willen, "Reducing foreclosures: No easy answers," in "NBER Macroeconomics Annual 2009, Volume 24," University of Chicago Press, 2010, pp. 89–138.

Kleven, Henrik J. and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *The Quarterly Journal of Economics*, 2013, *128* (2), 669–723.

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," 2016. Forthcoming, *Journal of Business & Economic Statistics*.