

Regional Redistribution Through the U.S. Mortgage Market: Online Appendix*

Erik Hurst
University of Chicago
Booth School of Business
and NBER

Benjamin J. Keys
University of Chicago
Harris School of Public Policy

Amit Seru
University of Chicago
Booth School of Business
and NBER

Joseph S. Vavra
University of Chicago
Booth School of Business
and NBER

April 2016

A-1 Empirical Appendix

In this appendix, we provide more details on our matching procedure and results from a number of robustness exercises.

A-1.A Matching Procedure

As described in Section III.C of the text, our primary set of comparisons are between loans originated in the non-GSE jumbo market and comparable loans originated in the GSE market in the 2000–2006 period since the jumbo market dries up after this period (Figure A-1). In order to make a closer comparison, in addition to our standard sample restrictions (30-year fixed rate full documentation loans made for either purchase or refinance of a single-family residence or condo), we also select our subsample of GSE loans to be from the same markets as the non-GSE loans in our sample. This restricts our sample to using GSE loans from 106 of the 374 possible MSAs.

To further make similar comparisons, we select our subsample of GSE loans to match the distribution of both FICO and LTV scores in the non-GSE market in a way that is symmetric across the loan amount distribution. To do so, we randomly select loans from ranges of FICO scores and LTV scores based on the quartiles of the non-GSE distribution in each “bin” of loan amounts used in the regression-discontinuity style estimates. For instance, there are 32,123 loans in the bin representing loans that are between the conforming loan limit and 1.2 times the conforming loan limit in the non-GSE market. To match FICO and LTV in the bin on the other side of the conforming loan limit (between 0.8 and 1.0 of the limit), we split the GSE loans in that bin into each quartile of the FICO \times LTV distribution, and select roughly 2,000 loans from each of these 16 bins.

This procedure leads to a very close match of the distribution of FICO and LTV scores in the non-GSE market. Appendix Table A-1 shows the results of the procedure. Not only are the means very similar across FICO and LTV, but the entire distributions line up closely because of this procedure.

A-1.B Robustness and Extensions

In this part of the Appendix, we discuss in detail a variety of robustness results we studied. Some of the robustness specifications looked at whether different factors could explain our results (e.g., regional variation in points or fees or regional variation in pre-payment risk). Other robustness specifications looked at different outcome variables (e.g., loan quantity adjustment). Still other robustness exercises looked at other specifications (e.g., exploiting time variation in the conforming limit within a MSA). Finally, our last set of robustness exercises looked at regional

*E-mail: Erik.Hurst@chicagobooth.edu; benkeys@uchicago.edu; Amit.Seru@chicagobooth.edu; Joseph.Vavra@chicagobooth.edu.

variation in GSE and prime jumbo loans to other policies that may affect the risk adjusted return of making a loan (e.g., local lender competition or local bankruptcy laws)

A-1.B.1 Exploring Regional Variation in Mortgage Quantities Instead of Interest Rates

In our primary analysis, we explored the adjustment of mortgage prices in response to spatial variation in regional risk. One may also expect some adjustment to occur on the quantity side, i.e., both on the extensive (loan approval) and intensive (loan amount, conditional on approval) margins.¹ Unfortunately, we are not able to explore variation on the extensive margin, because the only available data on the extensive margin (HMDA database) does not have borrower-level variables, which are crucial for differentiating borrower-level risk from location specific risk. We are, however, able to explore quantity movements on the intensive margin using our data. Appendix Figure A-2 shows the relationship between lagged default rates and LTV residuals for both the GSE sample (top panel) and the prime jumbo sample (bottom panel). These figures are similar to Figure ?? of the main text. We residualize LTV controlling for FICO score and time effects in a way similar to our residualization of interest rates. As seen from Appendix Figure A-2, there is little LTV adjustment across MSAs in response to differences in lagged default rates in either sample. If anything, borrowers in riskier places are slightly more leveraged on average. Moreover, there are no statistical differences in the response rates between the two samples. Given this finding, we focus our analysis only on the price margin and not the quantity margin.

Additionally, we focused on a sample where rejection rates are quite small. Specifically, for high quality borrowers with LTV ratios less than 0.8, the chance of the loan being rejected is very small. Yet, even within this sample, we found that the prime-jumbo interest rates varied significantly with ex-ante predictable local default risk while the GSE interest rates did not. We highlight these results more in the “Points, Fees, and Private Mortgage Insurance” subsection below.

A-1.B.2 Prepayment Risk

One potential concern with our base empirical work is that we did not account for other potential local risks that could affect local loan pricing. In particular, aside from default risk, the biggest risk lenders face is prepayment risk. If prepayment risk differs dramatically between GSE loans and prime jumbo loans in a way that is correlated with local default risk, the lack of variation in GSE mortgage rates with local default risk may not be surprising.

In our data, we can track prepayments and thus create a measure of predicted local prepayment risk. In particular, we follow the same procedure as in the main text for local default risk and create three different measures of predicted prepayment risk: the regression-based approach for both samples using lagged local GSE prepayment rates, a perfect foresight model in which the predicted prepayment rate is the actual local prepayment rate for each sample, and a random walk model in which the predicted prepayment rate is the actual lagged local prepayment rate for each sample. Our first finding is that predicted prepayment rates, conditional on loan and borrower observables, are very similar for GSE and prime jumbo loans. For example, using our RD approach, annual predicted prepayment rates were only 1 percentage points lower for prime jumbo loans above the conforming threshold relative to the GSE loans below the threshold (19% vs. 20%).

What matters is whether predicted prepayments are differentially correlated with predicted default rates across the two samples in a way that undoes the results documented above. To explore this, we added predicted prepayment rates as an additional control to all our main empirical specifications. Table 6 from the main paper shows one such specification. We focus on our base RD specification where predicted local default is based on a regression of actual default on lagged GSE default and loan-level controls. Column (1) of Table 6 redisplayes our estimates from column (5) of Table 3. We do this to facilitate comparison across our robustness specifications. Column (2) shows our RD estimates when we add the regression-based measure of predicted prepayments as an additional control. Notice that controlling for predicted prepayment risk does not change the RD estimates in any meaningful way. Again, this is not surprising given the fact that conditional prepayment probabilities barely differ between the samples. These results suggest that predicted prepayment differences are not driving the differential interest rate sensitivities to local default risk between the GSE and private samples.

A-1.B.3 MSA Fixed Effects

Another potential concern with the interpretation of our previous results is that identification could be driven by across-MSA differences in the composition of GSE versus private loans rather than from differential responses of these loans to common local conditions. To address this concern, we re-estimated all our specifications including

¹Note that most models would imply that when lenders reduce loan quantities they would also raise loan prices, so the fact that there is no regional variation in GSE loan prices strongly suggests there is no quantity variation.

MSA fixed effects. This allows us to compare GSE loans within an MSA to prime jumbo loans within the same MSA. Column (3) of Table 6 from the main text controls for MSA fixed effects in our RD specification, while column (4) controls for both MSA fixed effects and local prepayment risk. As can be seen from the table, the estimated difference in interest rate responsiveness to local default risk, $(\beta^{jumbo} - \beta^{GSE})$, is essentially unchanged in all the specifications.

A-1.B.4 Points, Fees, and Private Mortgage Insurance

Up until this point, we have not examined regional variation in points paid or other loan fees because points and fees are not recorded in our data. It may be the case that mortgage rates do not vary across MSAs in the GSE sample, but that points and other fees do vary with local default risk. To address this concern, we obtained additional data from one of the GSEs to directly estimate the relationship between effective interest rates and regional risk. The measure of effective interest rates in this data nets off any points and fees (including closing costs) that are charged to the borrower. As shown in Appendix Table A-2, we find no significant relationship between effective interest rates in the *universe* of GSE loans that meet our sample criteria and regional risk, as measured by lagged GSE default. The effect of a two-standard-deviation increase in regional risk is an insignificant 5 basis points increase in effective interest rate. In results not shown, no component of the effective interest rate (either points or fees) were found to be statistically associated with regional risk for these loans.

One additional concern with our analysis is that the securitized lenders may require borrowers with higher LTVs to purchase private mortgage insurance (PMI). The high LTV borrowers (usually those with an origination LTV greater than 80%) would pay for PMI that would insure the lenders against part of the principal loss during the default. Instead of variation in local predicted default risk showing up as variation in local mortgage rates, it could show up as variation in local PMI premiums. Again, this would affect our results only if the PMI was differentially used between the GSE and prime jumbo samples. Because we do not observe whether the loan had private mortgage insurance, we cannot control for this directly in our analysis. However, we reestimate our key results on a subsample of loans that are explicitly not required by the GSEs to have PMI, namely loans with loan-to-value ratios less than or equal to 80 percent at origination. When we restrict both the GSE and prime jumbo samples to only loans with $LTV \leq 0.80$, we lose roughly 30 percent of our sample. Column (5) of Table 6 from the main text shows our RD estimates when we restrict the sample to loans with $LTV \leq 0.80$. Notice that, even in this restricted sample where PMI is not required our RD coefficients are nearly identical to our base case. If anything, the differences between GSE and non-GSE loans are slightly larger among this subsample where the GSEs are directly bearing the cost of any default.

As an additional robustness specification, we examine rate quotes from LoanSifter, a firm that collects mortgage contract quotes across a range of U.S. markets and contract types. The advantage of this data is that the prices quoted are intended to be holding points constant (in our case, at zero). We use rate quotes from September 2009 through November 2010 (collapsed to the quarterly level) across 57 metro areas, for quoted prices for a 30-year fixed-rate conforming loan with a 20 percent downpayment (that is, an 80% LTV ratio) and a 750 FICO score. The time period reflects the availability of data and follows the worst of the Great Recession, leading to substantially larger lagged default rates than in prior periods. As shown in Appendix Figure A-3, we observe no relationship between lagged GSE default rates and quoted mortgage rate prices.

A-1.B.5 Additional Robustness Tests

Appendix Table A-3 presents a few additional robustness checks on our main empirical findings. The first column shows our baseline estimates of the regression-discontinuity style effect of the difference in the relationship between interest rates and regional risk across the GSE and non-GSE markets and uses the same sample and same specification as described in the main text. Column (2) restricts our baseline sample to only those markets with at least 10 non-GSE loans in a given quarter. Column (3) includes a set of lagged FICO and lagged LTV controls, thus conditioning on both contemporaneous and prior composition of loan quality in each market. Across all of these additional specifications, the results are consistent in magnitude. These results suggest our main findings are not sensitive to our selection of MSAs or the exact specification of our FICO and LTV controls.

A-1.B.6 Exploiting Time Variation in the Conforming Limit

We also conducted a robustness specification that exploits time variation in the conforming limit, shown in Appendix Table A-4. In particular, we focus on specific set of loans of a given dollar amount that changed between being above the conforming limit to being below the conforming limit during the early 2000s. For instance, a loan amount of \$350,000 was above the limit (and thus excluded from the GSE market) in the period prior to 2005, but in 2005

the conforming loan limit was raised such that a loan of this size could be purchased and insured by Fannie Mae or Freddie Mac. We then ask if interest rates for loans of \$350,000 respond to local predicted default risk prior to 2005 but not respond to local predicted default risk after 2005. To explore this, we reestimate our main specifications using only the set of loans in the range of \$276,000 and \$417,000 that change conforming status between 2001 and 2006. For these loans, prior to the conforming limit change, they responded significantly to local predicted default risk. However, for loans in this same range, interest rates no longer responded to local default risk once they became conforming loans. The magnitudes of the differences in Appendix Table A-4 are consistent with our estimates in Table ???. We view this as further evidence that GSEs do not vary interest rates in response to local predictable default risk.

A-1.B.7 Variation of Mortgage Rates to Other Local Risk Factors

In our final analysis, we show that local mortgage rates for loans securitized by the GSEs do not vary with other dimensions that could also induce local adjustment for risk such as local mortgage recourse laws, local bankruptcy laws, or local lender concentration. In particular, in Appendix Table A-5 (Panel A), following Scharfstein and Sunderam (2013), we use two measures of MSA-level lender concentration, the market share of the four largest lenders and the Herfindahl-Hirschman Index (HHI). As seen in the table, we find no significant relationship between residualized interest rates and lender concentration among GSE loans. In the prime jumbo sample, we find a small statistically significant negative relationship, suggesting that interest rates in this market were, if anything, slightly lower in more concentrated markets.

In Panel B of Appendix Table A-5, we assess if systematic variation in different state laws that capture the ability of lenders to recover their assets in a timely manner could explain the dispersion in residualized interest rates. In particular, we estimate the relationship between rates and whether the state allows for a “deficiency judgment” against the debtor (“recourse”), whether the state requires a judicial procedure to complete a foreclosure, and an indicator for whether the state is in the top half of the distribution in the generosity of homestead exemptions in personal bankruptcy. All of the coefficients on the state laws are statistically insignificant in both the GSE and prime jumbo samples. In other words, there is no spatial variation in mortgage rates that are associated with local bankruptcy or creditor protection laws.

A-2 Computational Appendix

In this appendix, we describe the solution to the model described in the body of the text. Reviewing the baseline model setup, the household state vector is defined as $\mathbf{s}_{jk} = (b_j, m_j, h_j, z_j, r_j^{m, fixed}; \gamma_{jk})$, and the model is solved by backward induction from the final period of life. When working, households solve:

$$V_j(\mathbf{s}_{jk}) = \max \left\{ V_j^{adjust}(\mathbf{s}_{jk}), V_j^{noadjust}(\mathbf{s}_{jk}), V_j^{refi}(\mathbf{s}_{jk}), V_j^{rent}(\mathbf{s}_{jk}) \right\}$$

with

$$\begin{aligned} V_j^{adjust}(\mathbf{s}_j) &= \max_{c_j, b_{j+1}, m_{j+1}, h_{j+1}} U_{jk}(c_j, h_{j+1}) + \beta E_j(V_{j+1}(\mathbf{s}_{j+1, k})) \\ & \text{s.t.} \\ c_j &= b_j(1+r) - b_{j+1} + (\chi_j + z_j)(\gamma_{k,j})^{\phi_y + \phi_r^y \phi_r} - (1 + r_{k,j}^{m, market})m_j + m_{j+1} \\ & \quad + (\gamma_{k,t})^{\phi_h + \phi_r^h \phi_r} h_j (1 - \delta^h) (1 - F) - (\gamma_{k,t})^{\phi_h + \phi_r^h \phi_r} h_{j+1} \\ b_{j+1} &\geq 0, \quad m_{j+1} \geq 0 \\ \log z_{j+1} &= \rho_z \log z_j + \eta_{j+1} \\ \log \gamma_{k,j+1} &= \rho_\gamma \log \gamma_{k,j} + \varepsilon_{k,j+1} \\ m_{j+1} &\leq (1 - \theta) (\gamma_{k,t})^{\phi_h + \phi_r^h \phi_r} h_{j+1} \\ r_{k,j}^{m, market} &= r + \bar{\Psi} \gamma_{k,j}^{-\phi_r} \\ r_{j+1}^{m, fixed} &= r_{k,j}^{m, market}, \end{aligned}$$

when households choose to adjust the size of their owner-occupied house. The value function for non-adjusters is given by:

$$\begin{aligned}
V_j^{noadjust}(\mathbf{s}_j) &= \max_{c_j, b_{j+1}, m_{j+1}} U_{jk}(c_j, h_j) + \beta E_j(V_{j+1}(\mathbf{s}_{j+1, k})) \\
&\text{s.t.} \\
c_j &= b_j(1+r) - b_{j+1} + (\chi_j + z_j)(\gamma_{k,j})^{\phi_y + \phi_r^y \phi_r} - (1 + r_j^{m, fixed})m_j + m_{j+1} - \delta^h (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j \\
b_{j+1} &\geq 0, \quad m_{j+1} \geq 0 \\
\log z_{j+1} &= \rho_z \log z_j + \eta_{j+1} \\
\log \gamma_{k,j+1} &= \rho_\gamma \log \gamma_{k,j} + \varepsilon_{k,j+1} \\
m_{j+1} &\leq (1 - \theta) (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j \\
h_{j+1} &= h_j \\
r_{j+1}^{m, fixed} &= r_j^{m, fixed},
\end{aligned}$$

The value function for households who refinance but do not move is given by:

$$\begin{aligned}
V_j^{refi}(\mathbf{s}_j) &= \max_{c_j, b_{j+1}, m_{j+1}} U_{jk}(c_j, h_j) + \beta E_j(V_{j+1}(\mathbf{s}_{j+1, k})) \\
&\text{s.t.} \\
c_j &= b_j(1+r) - b_{j+1} + (\chi_j + z_j)(\gamma_{k,j})^{\phi_y + \phi_r^y \phi_r} - (1 + r_{k,j}^{m, market})m_j + m_{j+1} \\
&\quad - \delta^h (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j - F^{refi} (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j (1 - \delta^h) \\
b_{j+1} &\geq 0, \quad m_{j+1} \geq 0 \\
\log z_{j+1} &= \rho_z \log z_j + \eta_{j+1} \\
\log \gamma_{k,j+1} &= \rho_\gamma \log \gamma_{k,j} + \varepsilon_{k,j+1} \\
m_{j+1} &\leq (1 - \theta) (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j \\
r_{k,j}^{m, market} &= r + \bar{\Psi} \gamma_{k,j}^{-\phi_r} \\
r_{j+1}^{m, fixed} &= r_{k,j}^{m, market} \\
h_{j+1} &= h_j,
\end{aligned}$$

and a household that chooses to sell its current house² and rent has value function

$$\begin{aligned}
V_j^{rent}(\mathbf{s}_j) &= \max_{c_j, b_{j+1}, m_{j+1}, h_{j+1}^f} U_{ijk}(c_j, h_{j+1}^f) + \beta E_j(V_{j+1}(\mathbf{s}_{j+1, k})) \\
&\text{s.t.} \\
c_j &= b_j(1+r) - b_{j+1} + (\chi_j + z_j)(\gamma_{k,j})^{\phi_y + \phi_r^y \phi_r} - (1 + r_{k,j}^{m, market})m_j \\
&\quad + (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_j (1 - \delta^h) (1 - F) - r^f (\gamma_{k,t})^{\phi^h + \phi_r^h \phi_r} h_{j+1}^f \\
b_{j+1} &\geq 0, \quad m_{j+1} = 0 \\
\log z_{j+1} &= \rho_z \log z_j + \eta_{j+1} \\
\log \gamma_{k,j+1} &= \rho_\gamma \log \gamma_{k,j} + \varepsilon_{k,j+1} \\
r_{k,j}^{m, market} &= r + \bar{\Psi} \gamma_{k,j}^{-\phi_r} \\
h_{j+1} &= 0.
\end{aligned}$$

The problem for a retired household is identical except that social security benefits replace labor earnings, and future payoffs are discounted at rate $\beta(1 - d_j)$ where d_j is an age-specific probability of death.

In order to implement the solution to this model numerically, we proceed as follows. First, we note that since $r_{k,j}^m > r \forall k, j$ it is never optimal to simultaneously hold both positive m and positive b . Thus, we can replace m and

²If previously a renter, the household will start the period with $h_j = 0$ so will have nothing to sell when it chooses to rent again.

b with a single variable a . This financial asset variable is positive if $b > 0$ and negative if $m > 0$, and households face $r_{k,j}^m$ when $a < 0$ and r when $a > 0$. Thus, the household state reduces to $\mathbf{s}_{jk} = (a_j, h_j, z_j, r_j^{m, fixed}, \gamma_{jk})$. Second, notice that if $F^{refi} = 0$ as in our baseline model, then $V_j^{refi}(\mathbf{s}_{jk}) \geq V_j^{noadjust}(\mathbf{s}_{jk})$, so we can eliminate $V_j^{noadjust}(\mathbf{s}_{jk})$ from the problem and simply set $r_{j+1}^{m, fixed} = r_{k,j}^{m, market}$. Similarly, if $F^{refi} = F$, as in our robustness check with high refinancing costs, then $V_j^{adjust}(\mathbf{s}_{jk}) \geq V_j^{refi}(\mathbf{s}_{jk})$ and we can eliminate the refinancing choice from the problem.

In order to rectangularize the choice set and simplify the computational problems imposed by the endogenous liquidity constraint, we follow Diaz and Luengo-Prado (2010) in reformulating our problem in terms of voluntary equity. In particular, define $q_j \equiv a_j + (1 - \theta)p_{k,j}h_j$. After substituting the budget constraint into the utility function to eliminate non-durable consumption as a choice variable, the value function can then be rewritten in terms of two endogenous variables q_j and h_j , the choice of which is restricted to be strictly positive. Note that $q_{j+1} \equiv a_{j+1} + (1 - \theta)\gamma p_{k,j+1}h_{j+1}$ but that a_{j+1} and h_{j+1} are chosen in period j . Thus, shocks to house prices mean that voluntary equity realized at the start of period $j + 1$ may differ from that chosen at the end of period j . Define $\widetilde{q}_{j+1} \geq 0$ to be the choice of voluntary equity for period $j + 1$ made in period j . The state of realized voluntary equity relevant for $j + 1$ then evolves as $q_{j+1} = \widetilde{q}_{j+1} + (1 - \theta)h_{j+1}(p_{k,j+1} - p_{k,j})$. This implies that although households are constrained to always choose $\widetilde{q}_{j+1} \geq 0$, actual voluntary equity can become negative if house prices fall by a large enough amount. To account for this, we solve the model for states that include negative voluntary equity even though households are constrained to choose non-negative values for this variable.³

We discretize the problem so it can be solved on the computer by first discretizing γ and z using the algorithm of Tauchen (1986). We use 13 grid points for z and 5 grid points for γ . We then approximate $V_j^{adjust}(q_j, h_j, z_j, \gamma_j, r_j^{m, fixed})$, $V_j^{refi}(q_j, h_j, z_j, \gamma_j, r_j^{m, fixed})$, and $V_j^{rent}(q_j, h_j, z_j, \gamma_j, r_j^{m, fixed})$ as multilinear functions in the endogenous states. In our benchmark calculation, we use 50 knot points for q_j (we space these points more closely together near the constraint) and 36 knot points for h_j . The presence of fixed adjustment costs on housing together with the borrowing constraint make the household policy function highly non-linear. For this reason, we follow Berger and Vavra (2015) and compute optimal policies for a given state-vector using a Nelder-Meade algorithm initialized from 3 different starting values, to reduce the problem of finding local maxima. The value of adjusting, not adjusting and renting are then compared to generate the overall policy function. We proceed via backward induction from the final period of life.

In order to compare the constant interest rate model to that with risk-based interest rates, we solve the model for both of these scenarios. To compute consumption equivalents, we then simulate a panel of 100,000 households over their life-times to find the endogenous joint-density of household states over regional economic activity. We record average welfare by region in the variable interest rate world and then compute the percentage change in household consumption required to make average welfare in that region equal to average welfare in that region in the constant interest rate simulation.

We consider two extensions of the model that require modifications of the computation. The first extension changes the fixed-rate, variable-balance mortgages in our baseline model to fixed-rate, fixed-balance mortgages. This adds the additional restriction that when not adjusting $m_{j+1} = m_j$. Importantly, in this environment it is no longer true that it is never optimal to simultaneously hold both positive m and positive b . Thus, we must solve the model with an additional state variable rather than the simplified version after substituting a for b and m . In addition, this specification requires an alternative method to rectangularize the state-space. Since $0 \leq m_{j+1} \leq (1 - \theta)p_{k,j}h_j$, applying the analogous previous change of variables $q_j \equiv -m_j + (1 - \theta)p_{k,j}h_j$ delivers the non-rectangular constraint that $0 \leq q_j \leq (1 - \theta)p_{k,j}h_j$. Instead, we transform the state-variable into a leverage ratio: $q = \frac{m}{(1 - \theta)ph}$. The problem can be restated in terms of this leverage ratio with a restriction that $-1 \leq q \leq 0$. Leverage is not defined for renters, but this is not problematic since by definition they will always have $m = 0$, which is enough information to fully characterize the rental problem. This version of the model is substantially more computationally challenging, as it adds an additional continuous state-variable as well as an additional continuous choice variable. Now instead of choosing net-debt, households must separately choose mortgage debt and savings. When adjusting, this means that the choice is now three dimensional: households choose b , q , and h . We modify the Nelder-Meade algorithm accordingly, but in order to compute the model in a reasonable amount of time, grid sizes are reduced from our baseline: we solve the model with 20 grid points in b and q and 22 grid points in h . We also reduce the grid size for z to 11 points. Despite this reduction in grid sizes, this version of the model takes roughly 50 times longer to solve than the baseline. While versions of the model with smaller grid sizes can be solved reasonably quickly, we found that results were sensitive to reducing grid sizes further. Increasing grid sizes modestly did not appear to affect the

³Shocks to house prices in the model are not large enough to ever reach a situation where q is so negative that households would be unable to choose $\widetilde{q}_{j+1} \geq 0$ without having negative consumption.

conclusions.

The second extension introduces the option to default, as in Campbell and Cocco (2015). We assume that any point in time, households can default on fraction Ω of existing mortgage debt (net of the value of housing), but that they must then sell their house and are excluded from owning for the rest of their lives. Since $\theta = 0.2$ and we abstract from heterogeneity in initial LTV, the majority of households in the model have positive housing equity and do not default. Thus, a high value of Ω is required to match observed default rates. We define a two-year lagged default variable as in the data, and setting $\Omega = 0.95$ delivers an average value for this variable of around 2.5%, in line with results for our sample period.⁴ The introduction of the option to default introduces one new value function, which is analagous to those in our baseline problem. It also introduces one additional binary state-variable which tracks whether or not a household has previously defaulted and is thus excluded from the ownership market.

A-3 Model Robustness

A-3.A Alternative Mortgage Contracts

In this section, we show that our results are robust to a variety of alternative specifications for mortgage contracts. In our baseline model, we assume that mortgages are fixed rate, since this characterizes the majority of mortgage contracts in the U.S. Allowing for the presence of fixed rates is potentially important since it can affect the extent to which households are subject to interest rate fluctuations which might arise from variation in local risk. For example, if local shocks were iid and refinancing was costless, then households with access to fixed rate mortgages might actually prefer a variable interest rate policy, since they could lock in low rates when conditions are good and keep them if conditions decline tomorrow.

In row 2 of Table A-6, we assess the importance of fixed rate mortgages for our conclusions by computing a version of the model where all mortgages are instead adjustable rate. In particular, we assume that the mortgage rate paid by a household in period t is always equal to $r_{k,t}^{m,market}$. This row shows that whether rates are fixed or variable has little quantitative effect on our conclusions. When conditions are bad, the subsidy from the constant interest rate policy is slightly larger under adjustable rate mortgages (2.35%) than under fixed rate mortgages (2.26%) since abandoning the constant interest rate policy makes all rates rise under adjustable rate mortgages. Conversely, the tax from the constant interest rate policy in good regions is slightly larger under fixed rate mortgages (-2.12%) than under adjustable rate mortgages (-1.88%). This is because when conditions are good and interest rates are not constant, fixed rate mortgages allow households to lock in a relatively low rate that remains permanently, while this is not possible with adjustable rate mortgages.

While it may seem surprising that fixed rates make little difference for our conclusions, this is true for a number of reasons. 1) Anyone who is buying a house or refinancing will be subject to the current rate, including all first time home buyers. These agents tend to be more constrained and as a result are more affected by interest rate variation than those who are not adjusting and maintain the old rate. For renting households contemplating buying, the distinction between the FRM and the ARM makes much less difference in terms of how they are affected by current rates; 2) Regional shocks are reasonably persistent; 3) There is an important life-cycle component to housing and mortgage adjustment decisions; and 4) Idiosyncratic shocks are large compared to the regional variation in interest rates.⁵

The last two facts imply that households are largely making their decisions based on life-cycle considerations and idiosyncratic factors rather than timing purchases or delaying purchases to take advantage of low rates. When combined with persistent regional shocks, this means that most households end up having a fixed rate mortgage which is equal to or very close to the prevailing interest rate so that the distinction between fixed and adjustable rates is not particularly important.

In our baseline model, we assume that mortgages have fixed rates but calibrate the model such that these mortgages can be refinanced with no cost. This specification has two advantages: First, it simplifies the solution of the model, since households always refinance when $r_{k,t-1}^{m,fixed} < r_{k,t}^{m,market}$ and never refinance when $r_{k,t-1}^{m,fixed} > r_{k,t}^{m,market}$. Second, this specification eliminates the possibility of lock-in effects which complicate the interpretation of transfers. In the presence of fixed rate mortgages with costless refinancing, households always take advantage of the opportunity to move to lower rates when local conditions improve. If refinancing is instead costly, then households can potentially get locked into the "wrong" rate. For example, if a household happens to refinance for idiosyncratic reasons when conditions are bad, it will continue to pay a high rate until refinancing again, even if local conditions improve.

⁴Note that nothing necessarily restricts Ω to be less than 1. $\Omega > 1$ could perhaps reflect the value of living in a house for free during foreclosure. Increasing Ω increases default rates but does not affect our conclusions for the size of implicit transfers.

⁵These conclusions hold under a wide range of empirically realistic calibrations of the local shock process.

Row 3 of Table A-6 shows that such lock-in effects are nevertheless quantitatively unimportant. In this row, we impose a very large refinancing cost $F^{refi} = F$, so that mortgages will only be refinanced by moving. Despite this large cost, results are nearly identical to our baseline. This again occurs because life-cycle conditions are the primary determinant of housing decisions: it is not worth staying in a house of the wrong size for several years in order to save 25 basis points or less on the mortgage rate. Since moving decisions are largely made independently of local interest rates, most households end up endogenously facing market rates which are near their current fixed rate. This in turn means that making it more difficult to refinance and take advantage of local rate variation does not have huge effects. Row 4 of Table A-6 similarly shows that increasing the fixed cost of moving also has little effect on our results.

Finally, for tractability, our baseline model assumes that mortgage *rates* are fixed but that mortgage *balances* can be adjusted without triggering a rate reset unless the household actually moves. In reality, the ability to adjust mortgage balances is asymmetric: most fixed rate mortgage contracts allow balances to be paid down more quickly than required without triggering a rate reset, but contracts do not allow mortgage debt to be increased while maintaining the same rate. The symmetry in our baseline model is thus undesirable, but it dramatically simplifies the computation of the model. Modeling mortgages with fixed balances requires separately tracking mortgage debt and liquid assets rather than just net debt, and it also introduces an additional continuous choice variable. In order to assess the importance of this simplification, we have solved a version of the more complicated model with fixed-rate, fixed-balance mortgages.⁶ This model is substantially more complicated to solve and requires less precise model solutions, but the fifth row of Table A-6 shows that it delivers nearly identical results. This is not particularly surprising because it turns out that in our baseline model only 1.5% of households actually choose to increase mortgage debt without refinancing or moving. Thus, even though we allow for this unrealistic behavior in our baseline model in order to substantially simplify computations, it is very unusual and so has little quantitative effect on our conclusions.

A-3.B Modeling default

In our result thus far, we do not explicitly model household default decisions and instead capture the effects of local conditions on credit risk and interest rates through their effects on the risk adjustment factor (Ψ). While it would be desirable to build a model with endogenous credit risk and interest rates rather than exogenously linking these variables to local conditions, this would substantially complicate the model. To what extent would such a model change any of our conclusions? Since we are interested in evaluating the effects of a change in interest rate policy, the main way in which endogenizing default could matter would be if default responds strongly to interest rates. These responses might then alter the size of implied transfers.

For example, the first half of the paper shows that under the constant interest rate status quo, default risk rises when local conditions deteriorate. If defaults also rise with interest rates, then abandoning the constant interest rate policy would further increase the response of default risk and interest rates to local conditions, which would amplify our conclusions. Conversely, endogenous default might instead dampen our results if when rates rise, households are able to use default to avoid paying higher interest rates. However, the presence of fixed rate mortgages in our baseline model reduces the importance of this second mechanism, since interest rates do not automatically rise unless households refinance.

What are the quantitative effects of endogenous default? A full-fledged model of strategic default and endogenous interest rates is beyond the scope of this paper, but it is straightforward to provide some sense of the sensitivity of our results to explicitly modeling default using a simple version of Campbell and Cocco (2015).⁷ Reassuringly, this model suggests that our results are unlikely to be altered by explicitly modeling local credit risk. In particular, we find that default responds only very mildly to local interest rate variation: In particular, a one-standard deviation increase in interest rates (12.5 bp) in the model increases default from 2.58% to 2.66%. Dividing this .08% change by the 2.9% cross-MSA standard deviation in the last row of Table 1 implies that this is only 0.027 (cross-MSA) standard deviation change in default. Furthermore, applying our estimates from the first half of the paper that a 1% increase in predicted default should lead interest rates to rise by 13.48 basis points means that a 0.08 percentage point increase in default should only increase interest rates by 1.07 basis points. This would increase the size of implied transfers but in a fairly negligible way that is well within the range of robustness we considered in Table ???. Thus, we conclude that formally modeling default and endogenous interest rates will greatly complicate our model without substantively altering its quantitative implications.

Table A-6 also shows that allowing for the option to default does not mitigate the effects of given interest rate variation. This does not imply that the option to default does not matter for welfare, but rather that any such

⁶See the computational appendix for details.

⁷See the computational appendix for additional description of this extended model.

welfare consequences do not interact strongly with variable interest rates. Put differently, regional risk matters for welfare, and the ability to default interacts with that, but this interaction is not much altered by interest rates which also vary with regional risk.

A-3.C Additional Robustness

Table A-7 explores how the overall level of regional risk interacts with the GSE constant interest rate policy in addition to alternative calibration choices for the model. The benchmark results are redisplayed at the top of the table, and the next rows show the sensitivity of our results to different values of ϕ^y . In our base specification, we set $\phi^y = 1$. As a robustness exercise, we show results for $\phi^y = 0.5$, and $\phi^y = 0.0$. Basically, the different estimates of ϕ^y determine the extent of regional income risk. In the extreme case, $\phi^y = 0$, there is no regional income risk at all. The next rows of the table show the sensitivity of our results to reducing ϕ^h , where ϕ^h is the responsiveness of local house prices to changes in local economic activity. If housing supply is perfectly elastic in all periods, $\phi^h = 0$. To explore the sensitivity of our results, we recompute results setting $\phi^h = 0.25$ and $\phi^h = 0$. The main take-away from Table A-7 is that varying sensitivity of regional income and house prices has essentially no effect on how households evaluate interest rate variation. In other words, although regional income and house price variation matter for the level of risk faced by households, they have little effect on the way that households are affected by changes in the interest rate across the two policies.

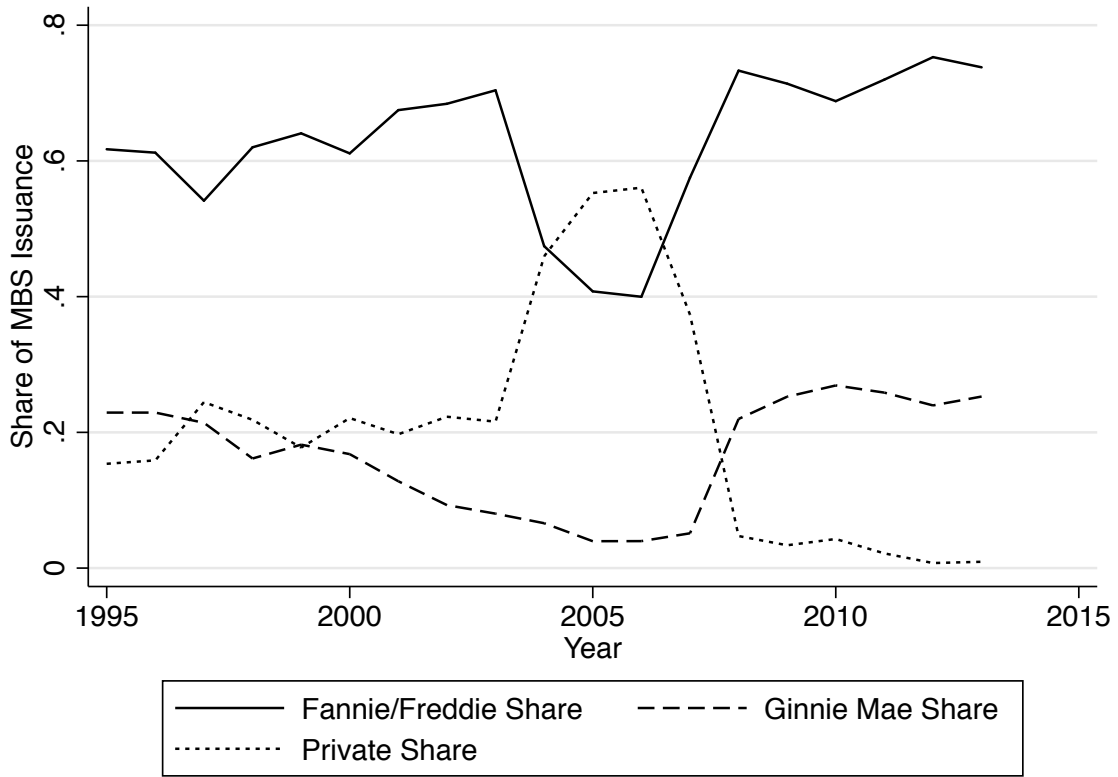
More specifically, the counterfactual experiment asks how much a household would pay to avoid having the interest rate rise today. What the non-interaction with ϕ^y and ϕ^h means is that the answer to this question depends little on whether a household is currently in a high-house-price area, a low-house-price area, a high-income area, or a low-income area. This can be seen most clearly in the second to last row, where both $\phi^h = \phi^y = 0$ so there is no regional variation in income or in house prices. This means that under the constant interest rate policy, regions are identical. The only way that regions vary when $\phi^h = \phi^y = 0$ is that in the variable interest rate counterfactual, some regions will have high interest rates and some regions will have low interest rates. Comparing that row with the benchmark shows that the welfare effects of interest rate variation in a world where house prices and income vary spatially is very similar to one where they do not. This occurs because there are offsetting interactions between regional risk and interest rate variation. For example, higher income makes households more likely to want to buy a big house, so that they dislike high interest rates more than low-income households do. But at the same time, higher income also makes households less likely to borrow, so they care less about high interest rates. The net effect is that these two effects roughly cancel each other out so the level of regional risk has little effect on implicit transfers from interest rate variation.⁸

In our baseline results, we follow Berger et al (2015) in calibrating our parameters to match the life-cycle profile of *liquid* wealth. Kaplan and Violante (2014) argue that matching liquid wealth is essential for generating realistic consumption dynamics. Nevertheless, Table A-7 shows results when we instead target the median *total* wealth-to-income of 1.52 from SCF data (see Kaplan and Violante (2010)). With higher wealth, households are better able to self-insure without borrowing and so the size of transfers is very mildly attenuated, but this effect is extremely minimal. In our baseline results, we match the initial distribution of assets for young households in the SCF, but Table A-7 shows that instead starting all households with zero housing and wealth does not importantly affect our conclusions. Finally, our baseline calibration for the housing share in total consumption is based on matching consumption and residential investment from BEA data and implies a housing expenditure share of around 15%. Davis and Ortalo-Magne (2011) focus on renters using census data and find somewhat larger expenditure shares of 24%.⁹ Table A-7 shows that if we target an expenditure share of 24%, then all of our results are moderately amplified. This is not surprising, since a larger expenditure share on housing implies a larger role for mortgage payments and sensitivity to interest rates.

⁸It is important to note that this non-interaction does not imply that ϕ^y and ϕ^h do not matter for welfare: households are much worse off in the world with regional income variation. But the relevant question is to what extent households' willingness to tolerate interest rate variation interacts with these other regional shocks. Table A-7 shows that there is no substantive interaction.

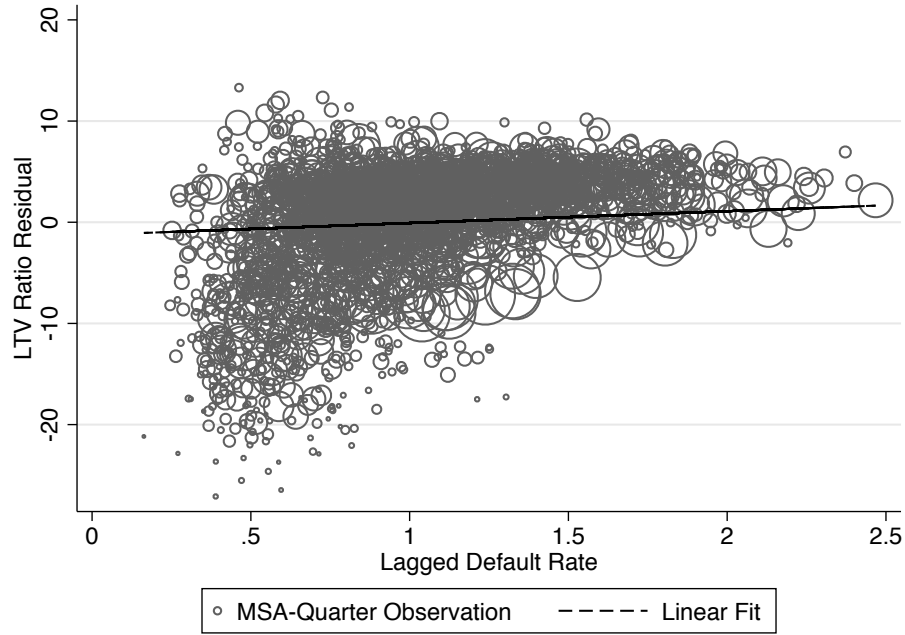
⁹See Section 2 of Davis and Ortalo-Magne (2011) for a discussion of the comparison to BEA data.

Figure A-1: MBS Issuance by Issuer, 1996-2013

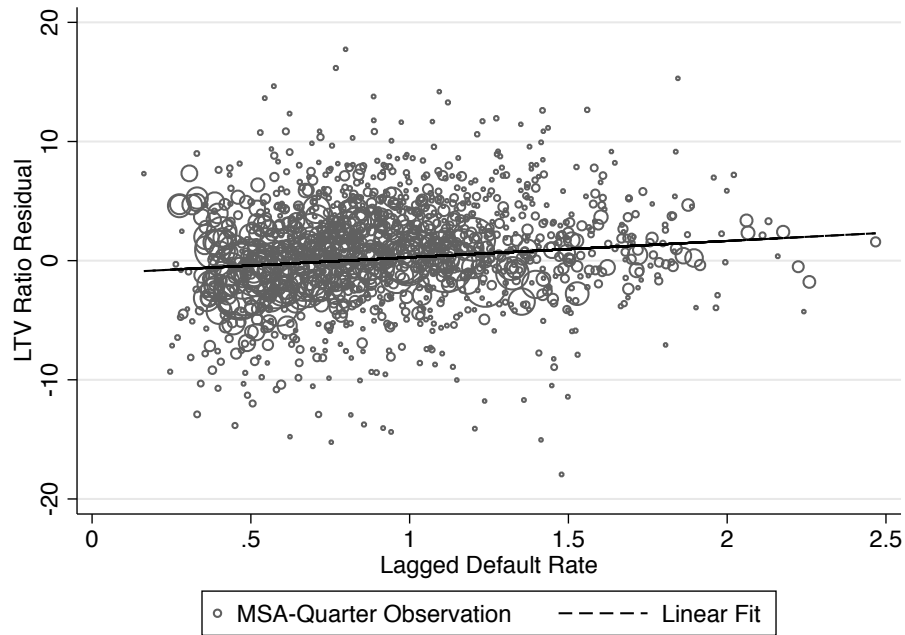


Note: This figure shows the share of mortgage-backed security issuance by the GSEs (Fannie Mae, Freddie Mac, and Ginnie Mae) and non-GSE issuers over the period 1996-2013. The focus of this paper is on the period 2001-2006, when the non-GSE issuance market was especially active, and 2007-2009, during the recession when the non-GSE market collapsed. Source: SIFMA).

Figure A-2: Relationship between LTV Ratio and Regional Risk



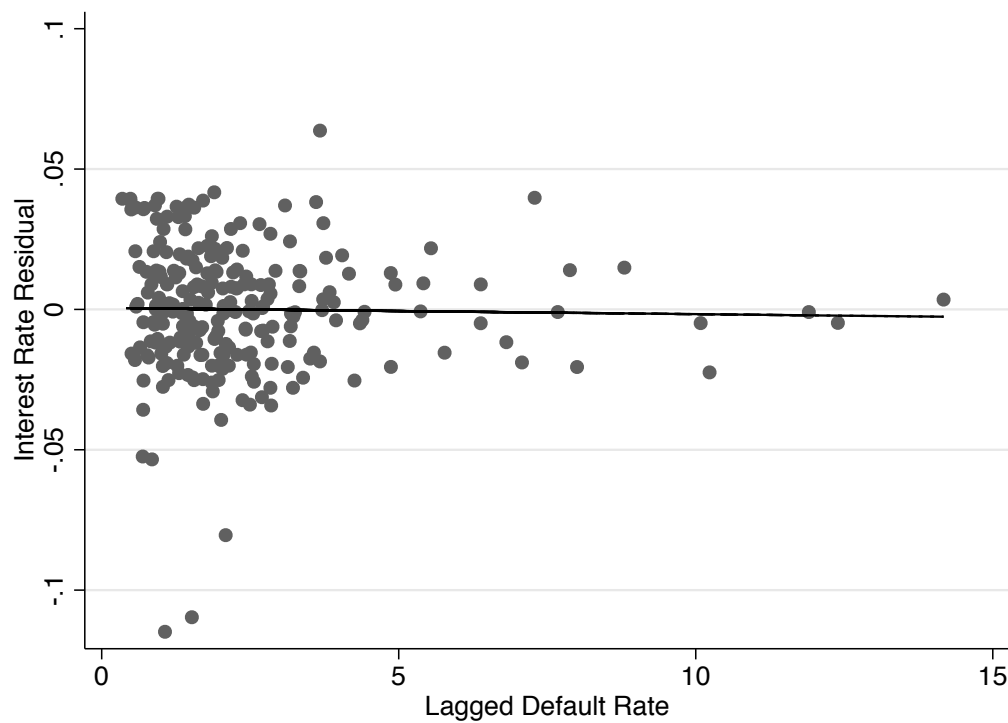
(a) GSE Loans



(b) Non-GSE Loans

Note: The figure presents the relationship between average loan-to-value (LTV) ratios and regional risk for both the GSE (panel a) and non-GSE (panel b) markets during 2001 to 2006. Regional risk is measured by the lagged default rate in the MSA. The estimated relationships are statistically and economically insignificant, and if anything show a slightly positive relationship between leverage and predicted risk during this time period. The LTV measure is residualized to remove time fixed effects and FICO scores. Each dot represents an MSA-quarter average. See text for details.

Figure A-3: Relationship between Interest Rates and Regional Risk, LoanSifter data



Note: The figure uses MSA-level quarterly data from 2009:Q4 through 2010:Q4 from LoanSifter to explore quoted conforming mortgage prices with points and fees held constant at a level of zero. The interest rate measure is residualized to remove time fixed effects. The estimated relationship between interest rates and lagged GSE default are statistically and economically insignificant. See text for details.

Table A-1: Descriptive Statistics of Matched Samples

Percentile	Interest Rate		FICO Score		LTV Ratio	
	Non-GSE	Matched GSE	Non-GSE	Matched GSE	Non-GSE	Matched GSE
10th	5.625	5.5	627	628	60	57
25th	5.99	5.875	637	638	70	70
50th	6.5	6.25	656	658	80	79
75th	7.125	6.875	698	700	85	85
90th	7.95	7.25	745	759	90	95
Mean	6.66	6.32	672	677	76.7	76.1
Observations	70,327	70,327	70,327	70,327	70,327	70,327

Note: The table provides summary statistics for the samples of matched GSE and non-GSE (“prime jumbo”) loans originated between 2001 and 2006. The matched GSE sample uses only those MSAs with prime jumbo loans present (during 2001 to 2006) and matches the distribution of FICO scores and LTV ratios in the non-GSE sample. See text for details.

Table A-2: Relationship between Effective Interest Rates and Regional Risk

	Interest Rate (%)	Effective Interest Rate net of points and fees (%)
Coefficient on Lagged GSE Default Rate	0.372 (0.337)	1.881 (1.414)
Observations (N> 6M)	N	N
Implied Interest Rate Change in Y to a Two Standard Deviation Change in Lagged GSE Default	0.01	0.05

Note: The table provides coefficients from regressions of interest rates and effective interest rates on regional risk as measured by lagged GSE default. Standard errors in parentheses are clustered at the MSA level. The data comes from one of the GSEs (anonymously), and we are not allowed to provide precise sample sizes except to note that there are more than 6 million loans in each specification. The GSE loans included use the same sample restrictions we have made above, and the specifications include quadratics in FICO and LTV interacted with time dummies as above. Column 1 reports the relationship between interest rate and our measure of regional risk. Column 2 uses a measure of “effective” interest rate that nets off any points and fees (including closing costs) paid by borrowers. Neither regression shows a significant association between our measure of regional risk and the cost of borrowing in the GSE mortgage market.

Table A-3: Additional Robustness Results

	Baseline	N>=10	Conditional Lags
Predicted	13.48	13.93	14.48
Default	(4.56)	(3.89)	(4.04)
Lagged GSE	32.54	34.59	29.61
Default	(4.00)	(3.95)	(6.33)
Lagged Own	13.04	13.25	11.88
Default	(4.57)	(4.60)	(5.29)
Actual	2.06	2.72	--
Default	(0.44)	(0.53)	--

Note: The table provides coefficients from a number of specifications of the regression-discontinuity style estimates using different samples or different controls. Column 1 provides the baseline estimates from the main text. Column 2 restricts the sample to only those loans that are not required to purchase private mortgage insurance. Column 3 restricts the sample to use only those markets with at least 10 non-GSE loans in an MSA-quarter of origination cell. Column 4 uses the baseline sample but includes controls for both contemporaneous and lagged measures of loan quality (FICO and LTV). Standard errors in parentheses clustered at the MSA level. Standard errors for results relying on predicted default are bootstrapped (500 repetitions, clustered at MSA level) to account for the generated regressor. See text for details.

Table A-4: Focus on Loans Transitioning Around the Conforming Loan Limit

	Estimated Coefficient			
	GSE Matched Sample	Prime Jumbo Sample	Difference	p-value of difference
Predicted Default Using	-3.38	10.76	14.14	<0.001
Lagged GSE Default	(3.69)	(1.07)		

Note: The table provides coefficients from separate regressions for loans backed by the GSE and non-GSE markets during 2001 through 2006 that had loan amounts between \$276,000 and \$417,000, the range of loans that switched status based on changes in the conforming loan limit. The results show strong and statistically significant correlations between interest rates and regional risk in the non-GSE market, but no meaningful relationship in the GSE market. As in the main text, the regressions control for loan quality using measures of FICO and LTV interacted with quarter of origination. Standard errors in parentheses clustered at the MSA level. Standard errors for results relying on predicted default are bootstrapped (500 repetitions, clustered at MSA level) to account for the generated regressor. See text for details.

Table A-5: Relationship between Interest Rates, Lender Concentration, and State Laws

	GSE Sample	GSE Sample	Prime Jumbo Sample	Prime Jumbo Sample
Lender HHI	-1.15 (0.78)		-5.92 (1.34)	
Top 4 Concentration		-0.171 (0.117)		-0.909 (0.193)
Years	2001-2006	2001-2006	2001-2006	2001-2006
Observations	624	630	461	465
MSAs	105	106	105	106
R-squared	0.03	0.02	0.13	0.13

(a) MSA-level Regressions

	GSE Sample	GSE matched Sample	Prime Jumbo Sample
Recourse Law	0.034 (0.019)	0.036 (0.024)	0.063 (0.045)
Judicial Foreclosure	0.038 (0.040)	-0.007 (0.067)	0.013 (0.039)
Debtor- Friendly	0.023 (0.039)	-0.021 (0.065)	0.02 (0.038)
Years	2001-2006	2001-2006	2001-2006
Observations	1224	809	809
States	51	46	46
R-squared	0.08	0.06	0.013

(b) State-level Regressions

Note: Panel A provides coefficients from regressions of interest rate residuals on lender concentration as measured by the Herfindahl-Hirschman Index (HHI) and the market share of the top 4 lenders in the MSA. The dependent variable is residualized interest rate, following the specifications in the main text, removing FICO, LTV, and higher-order polynomials, all interacted with quarter of origination fixed effects. Panel B provides coefficients from regressions of interest rate residuals on three relevant state laws. The laws measure the ability of lenders to recover their assets in a timely manner. The regression includes measures of whether the state allows for a “deficiency judgment” against the debtor (a.k.a. “recourse”), whether the state requires a judicial procedure to complete a foreclosure, and an indicator for whether the state is in the top half in terms of homestead exemptions in personal bankruptcy.

Table A-6: Sensitivity to Mortgage Contracts and Housing Transaction Costs

	Regional Employment				
	-2 Standard Deviation	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviation
Consumption gain: Bnchmk	2.26%	1.18%	-0.02%	-1.04%	-2.12%
Adjustable Rate Mortgages	2.35%	1.30%	0.12%	-0.86%	-1.88%
FRM-Large Refinancing Cost	2.42%	1.27%	0.09%	-0.92%	-2.03%
Double Fixed Cost of Moving	2.19%	1.02%	-0.16%	-1.30%	-2.27 %
FRM-Fixed Balance	2.20%	1.24%	-0.14%	-1.57%	-1.94%
Allow for Default	2.26%	1.18%	-0.04%	-1.05%	-2.11%

Note: This table shows the robustness of our estimated consumption gains to different interest rate sensitivities to local economic conditions. See text for a description of baseline parameters and the policy experiment as well as the description of each particular mortgage environment. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate.

Table A-7: Sensitivity to Regional Economic Variation and Alternative Calibrations

	Regional Employment				
	-2 Standard Deviation	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviation
Consumption gain: Bnchmrk	2.26%	1.18%	-0.02%	-1.04%	-2.12%
Reduce Regional Income Variation in Half	2.25%	1.16%	-0.05%	-1.16%	-2.25%
Reduce Regional Income Variation to Zero	2.20%	1.11%	-0.07%	-1.20%	-2.25%
Reduce Regional House Price Variation in Half	2.18%	1.16%	0.07%	-1.04%	-2.11%
Reduce Regional House Price Variation to Zero	2.11%	1.02%	-0.11%	-1.30%	-2.25%
Reduce Regional House Price and Income Variation to Zero	2.16%	1.07%	-0.07%	-1.21%	-2.22%
Calibrate to Total Wealth	2.09%	1.05%	-0.02%	-1.14%	-2.13%
No Initial Assets Calibration	2.32%	1.11%	-0.04%	-1.13%	-2.22%
Higher Housing Share	2.64%	1.50%	-0.02%	-1.58%	-2.98%

Note: This table shows the robustness of our implied consumption response to alternative parameterizations. See text for a description of baseline parameters and the policy experiment. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. In rows 2 and 3, we reduce the value of ϕ^h , in rows 4 and 5 we reduce the value of ϕ^y , and in row 6 we reduce the value of both elasticities. Row 7 presents results for the baseline model recalibrated to match median total net wealth rather than liquid wealth. Row 8 presents results for model where agents are born with no assets or housing. Row 9 shows results for a calibration matching higher rent share data rather than BEA statistics.