Regional Risk Sharing through the U.S. Mortgage Market*

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

January 2014

*E-mail: Erik.Hurst@chicagobooth.edu; benkeys@uchicago.edu; Amit.Seru@chicagobooth.edu; Joseph.Vavra@chicagobooth.edu. We thank seminar participants at the Bank of Italy and Wharton for helpful comments and suggestions. Any remaining errors are our own.

1 Introduction

Government sponsored mortgage enterprises (GSEs) securitize the bulk of all mortgages originated within the United States. In 2008, for example, the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac) owned or guaranteed about half of the U.S.'s $12 trillion mortgage market. In 2008, Fannie Mae and Freddie Mac were placed into conservatorship by the U.S. government. While there have always been discussions to dismantle Fannie Mae and Freddie Mac, those discussions have intensified in recent years. In order to inform the public debate, it is necessary to better understand the costs and benefits of the GSEs on economic activity.

In this paper, we start to formally explore one of the aspects of the GSE pricing decisions. In particular, we are interested in extent that the policies of Fannie Mae and Freddie Mac result in the redistribution of resources across U.S. regions in a state contingent way. To shed light on this issue, the paper unfolds in three parts.

The paper begins by empirically establishing that there is essentially no spatial variation in GSE mortgage rates, conditional on borrower observables, in any recent period. As a corollary, we show that local economic conditions are unrelated to the mortgage rates on loans securitized by Fannie and Freddie. Specifically, using loan level data from Fannie Mae and Freddie Mac, we show that average interest rates in an MSA are uncorrelated with recent local house price growth, recent local unemployment rates or changes in local unemployment rates, or recent local mortgage default rates. This zero result is precisely estimated.

The lack of spatial variation in mortgage pricing by the GSEs is surprising given the extent of spatial variation in predictable default risk. Both ex-post default and ex-ante predicted default rates vary substantially, and especially so in the wake of the housing bust. Although the GSEs charge different interest rates to borrowers who take on greater leverage or who are less creditworthy, they do not charge different rates to borrowers in regions with declining house prices, weakening labor markets, or greater lagged local mortgage default rates.

To what extent should agency mortgage rates vary spatially over time, given that the spatial variation in risk is quantitatively large? To address this question, we show that mortgage rates in the private "prime jumbo" market, where loans are larger than the conforming limit but comparable on many dimensions to loans backed by the GSEs, were strongly correlated with ex-ante predicted default probabilities across geography. Thus we show that in the government-backed market there is no regional risk-based pricing, whereas the private market does set interest rates based in part on regional risk factors.

In the second part of the paper, we explore a number of explanations for why GSE mortgage rates do not vary spatially over time, and conclude that political pressure is the most reasonable explanation for the patterns we observe. The government-sponsored enterprises face a great deal of political scrutiny. Evidence from prior efforts of the GSEs to regionally differentiate lending stan-

dards, most recently through the "declining markets" policy of 2008, suggests that these approaches have been quickly abandoned in the face of pressure from Congress, realtors, and community groups. Other similar settings, such as the U.S. Postal Service charging the same flat rate for all first-class mail, and U.K. insurance providers pricing nationally (Finkelstein and Poterba 2013), can best be explained by understanding the political economy of these markets.

In the final part of the paper, we attempt to quantify the economic impact of the GSEs' constant interest rate policy on regional risk sharing. To do so, we build a structural spatial model of collateralized borrowing where households face both idiosyncratic and region-specific shocks. Individuals in the model can choose whether to own a home and whether to borrow against their home, as well as make usual life-cycle saving and spending decisions. We compare two scenarios, one in which interest rates respond to the local default risk within each region, and one in which a common interest applies to all regions.

After parameterizing and calibrating the structural model based on our analysis of GSE and prime jumbo mortgage rate dispersion, we find that the GSEs national interest rate policy has significant insurance value. We estimate that households would pay $150 on average to maintain this constant interest rate policy, as this policy allows households to borrow more cheaply in the event of negative regional shocks. The cumulative value of this insurance is 0.16% of aggregate consumption, or roughly $18 billion.

We also use the model to estimate the magnitude of ex-post redistribution across regions when interest rates are set using a constant national rate. The difference between the top 5% and bottom 5% outcomes in terms of regional shocks leads to an ex-post redistribution of $1000 across households. This amount is comparable in size to fiscal stimulus packages such as tax rebates and payroll tax holidays.

The results suggest that the magnitudes of redistribution we observe through the mortgage market during aggregate downturns are economically meaningful. Although there are a range of consequences to the housing and mortgage markets that are often attributed to the presence of Fannie Mae and Freddie Mac, their common national interest rate policy is one important and understudied dimension of their impact on households choices.

2 Background

Most mortgages are sold to the secondary market after origination. For example, from 2004-2006, about 80 percent of all mortgages were securitized. Conventional mortgages are mortgages that meet criteria for purchase by Fannie Mae and Freddie Mac.[1] For the most part, securitized conventional mortgages are either directly purchased by or placed into pools that are guaranteed by

[1]Specifically, conventional mortgages are mortgages where (1) the mortgage amount is lower than a set limit (e.g., in $417,000 in 2006), (2) the loan amount relative to house value is below a set limit, and (3) borrower characteristics meet certain quality thresholds based on FICO (credit) scores and borrower debt to income ratios.

Fannie Mae and Freddie Mac.[2] This guarantee ensures that the default risk is born by the GSEs.

Figure 1 shows the share of mortgages in the secondary market that were securitized or directly held by Fannie Mae and Freddie Mac during the 2000s. Over most of the time period, between 60 and 70 percent of all securitized mortgages were securitized by Fannie and Freddie. The remaining amount was securitized by Ginnie Mae and the private market. Ginnie Mae is another government sponsored agency that purchases and pools mortgages issued by the Federal Housing Authority (FHA). The private market securitizes all other loans. These loans include jumbo mortgages (those that exceed the conventional mortgage size limits), sub-prime mortgages, and Alt-A mortgages. As seen from Figure 1, during the 2004-2007 period, the share of loans securitized by the private market grew at the expense of both those loans securitized by Ginnie Mae and loans securitized by Fannie/Freddie. The substitution between private market and Ginnie Mae loans is not surprising given that the expansion of sub-prime and Alt-A mortgages targeted lower quality borrowers that normally were the domain of the FHA loans. In late 2007, the private secondary mortgage market dried up and essentially all securitization of mortgages since that time has been conducted by the GSEs.

3 Data and Empirical Specification

Our primary data sources come from Fannie Mae's Single Family Loan Performance Data and Freddie Mac's Single Family Loan-Level Data Set. The population of both data sets includes a subset of the 30-year, fully amortizing, full documentation, single-family, conventional fixed-rate mortgages acquired by the GSEs between the late 1990s and 2012.[3] The data includes both borrower and loan information at the time of origination as well as data on the loan's performance. With respect to information at the time of origination, the data includes the borrower's credit (FICO) score, the date of origination, the loan size, the loan size relative to the house value (LTV), whether the loan is an origination or a refinancing, the three digit zip code of the property, and the interest rate on the mortgage. The loan performance data is provided monthly and includes information on the loan's age, the number of month's to maturity, the outstanding mortgage balance, whether the loan is delinquent, and the number

[2] Throughout the paper we will often refer to Fannie Mae and Freddie Mac as Government Sponsored Enterprises (GSEs). The GSEs were previously privately owned despite their public charters, but were taken into conservatorship by the federal government in 2008 during the crisis.

[3] Both data sets were made available to increase transparency of loans held or guaranteed by the two agencies. Each data set can be downloaded directly from the respective GSE websites. The Data Appendix which accompanies this paper provides a more detailed discussion of this data. In particular, it provides a more detailed discussion of what types of loans were included and the specific years of coverage. Both data sets include loans originated between January 1, 1999 and the mid 2012.

of month's delinquent.[4] There is a unique loan identifier code in both data sets that allows a loan to be tracked from inception through its future performance.

For our primary analysis, we divide our data into four different time periods. The first time period includes loans originated between 2001 and 2003. The second, third and fourth time periods includes loans originated 2004 - 2006, 2007 - 2009, and 2010-2012, respectively. We view these periods as reflecting different aggregate business cycle conditions. The first period encompasses the early 2000's recession while the third period encompasses the recent recession. The 2004-2006 period is a period of high house price growth and the period spanning the two other recessions. The last period reflects the post Great Recession period.

We break the data out into these periods for two reasons. First, we want to analyze the spatial variation in interest rates and performance during periods where there is a large spatial disparity in economic conditions. As we show below, the variation in economic performance across locations was highest in during the 2001-2003 recession and the 2007-2009 recession. Given that we are going to compare the spatial variability of interest rates in the Fannie/Freddie loans to the private market loans, we need to separately analyze the 2001-2003 period from the 2007-2009 period. As seen in Figure 1, the private market essentially stopped after 2007. We separately analyze the 2010-2012 period given that for this period we can only observe the spatial variation in mortgage rates for the Fannie/Freddie loans. Given that our data stops in 2012, we cannot follow the ex-post performance of loans for two years. The two year subsequent delinquency rate is going to be our primary measure of loan performance.

Our analysis includes both new housing purchases as well as refinanced loans. Our results from pooling together housing purchases and refinancing are unchanged if we look at the two types of originations separately.[5] In total, we have roughly 19 million loans in the Fannie Mae data set and 16 million loans in the Freddie Mac data set. For ease of exposition, we pool together the loans from Fannie and Freddie into one analysis file. As we show in the online Robustness Appendix that accompanies the paper, the results are identical in essentially all respects if we only use either the Fannie Mae data or the Freddie Mac data. Appendix Table A1 provides descriptive statistics for the pooled Fannie/Freddie data sets during each of the time periods..

Our second primary data source comes from the Loan Performance database, which contains loan-level origination and performance data on the near-universe of mortgage loans sold through the private secondary market during the housing boom. Within the Loan Performance database, we only focus on what we term fixed rate "prime" jumbo mortgages. As noted above, the loans securitized by the private market include both sub-prime and Alt-A mortgages as well as mortgages that are larger than the conforming loan limit. Specifically, we want to create a set of mortgages securitized by the private market that is most similar to the mortgages in the Fannie/Freddie pool. To do that, we define our

[4]The number of month's delinquent is provided in the following bins: 30-59 days, 60-89 days, 90-119 days, 120-149 days, 150-179 days, and 180+ days.

[5]These results are shown in an Online Robustness Appendix that accompanies the paper.

"prime" jumbo set of mortgages. The mortgages in this comparison set meet the following criteria: (1) the origination value is between the conforming mortgage (which can be MSA-specific) and $700,000, (2) have a fixed interest rate, (3) have a LTV at origination of less than 100, and (4) provide full documentation at the time of origination.[6] The reason we cap the mortgage origination value at $700,000 is again we want to create a sample that is similar to the Fannie/Freddie sample. As with the Fannie/Freddie mortgages, we include both new originations and refinancings in our sample. Given that the securitization of mortgages outside of GSEs dried up after 2007, we only examine the spatial variation in mortgage rates for the prime jumbo loans during the 2001-2003 and 2004-2007 periods.

We have xxx,xxx prime jumbo loans in the LP data for the 2001-2003 period and xxx,xxx for the 2004-2006 period. Descriptive statistics for these loans can be found in Appendix Table A2. [Discuss comparison of prime jumbo loan sample with Fannie/Freddie sample]

3.1 Default and Interest Rate Residuals

Throughout the paper, we want to examine the spatial variation in mortgage rates and show how they correlate with the spatial variation in predicted future mortgage default rates. One reason interest rates and default/delinquency rates differ spatially is because borrower characteristics differ spatially. For example, borrowers with lower credit scores empirically face higher interest rates and ex-post default more. If borrower credit worthiness varies spatially, this could explain some of the spatial variation in mortgage rates and default rates. What we are after, however, is whether interest rates and predictable default rates vary spatially conditioning on borrower/loan characteristics.

To formally illustrate these patterns, we purge the spatial variation in actual mortgage rates paid of spatial differences in borrower/loan characteristics. To do that, we first estimate the following equation using our loan level micro data:

$$r_{ist} = \alpha_1 + \alpha_1 X_{it} + \Gamma_1 D_t + \eta_{ist}$$

where r_{ist} is the loan-level mortgage rate for a loan made to borrower i, in location s, during period t and X_{it} is a set of control variables for borrower i in period t. Throughout the paper, our notion of space, s, is a U.S. metropolitan area (MSA). The borrower/loan controls include a quadratic in borrower FICO score and a quadratic in initial loan LTV.[7] D_t is a vector of time dummies

[6] In all MSAs prior to 2008, the conforming mortgage limit was uniform, such as $417,000 in 2006 and 2007. Starting in 2008, the GSEs implemented a new regionally varying loan limit based on the cost of housing. In some portions of high-cost MSAs – New York, Los Angeles, and San Francisco, for instance – the conventional mortgage rate was raised to as high as $729,750 for a one-unit residence.

[7] We experimented with higher order polynomials in both FICO score and LTV. Additionally, we experimented with interactions between FICO and LTV. None of these changes had any quantitative effects on our results. Given this, we used the quadratic separable specification as our base specification for ease of exposition.

based on the quarter*year of origination. The goal of these specifications is to recover η_{ist} and ζ_{ist}. $\eta_{ist}(\zeta_{ist})$ is the residual mortgage rate (ex-post default rate) for borrower i in MSA s during time t after controlling for borrower/loan characteristics and time fixed effects.

Once we have the residuals, we compute location specific averages for each of our four time periods, R_{st}. Specifically,

$$R_{st} = \frac{1}{N_{st}} \sum_{i=1}^{N_{st}} \eta_{ist}$$

where N_{st} is the number of loans in the MSA s during quarter*year t. Formally, R_{st}^m will be the average mortgage rate residual in an MSA during a given period.

One natural question is to assess how much of the variation in mortgage rates securitized by Fannie/Freddie can be explained by time effects and borrower characteristics. To answer this, we run four different versions of the above first stage regression. Each regression includes only loans originated during one of our four time periods. We run this first stage regression both including and excluding the vector of X controls. The R-squared from the above first stage regressions excluding the X vector of controls were xx, xx, xx, and xx in the 2001-2003 period, the 2004-2006 period, the 2007-2009 period, and the 2010-2012 period, respectively. These results shows that most of the variation in mortgage rates across individuals in the Fannie/Freddie data sets are due to time effects. Including both the time dummies and the X vector of controls increased the R-squared from the first stage regressions on mortgage rates to xx, yy, zz, and ww in the 2001-2003, 2004-2006, 2007-2009, and 2010-2012 periods, respectively. In summary, the vast majority of the variation in mortgage rates within the Fannie/Freddie data can be explained simply by time effects, a quadratic in FICO scores and a quadratic in LTV.

Table 1 shows the mean mortgage rate (unconditional), the standard deviation of unconditional mortgage rates (r_{ist}'s) across MSAs, and the standard deviation of conditional mortgage rates (R_{st}'s) across the MSAs for our four time periods of study. Two things are of note from Table 1. (I will fill the rest of this in once I see table 1 - discuss variation in unconditional mortgage rates and variation in conditional mortgage rates)

4 Spatial Variation in Mortgage Rates

4.1 Conventional Mortgage Rates and Current Local Economic Conditions

To examine whether conventional mortgage rates vary with local economic conditions, we need to define measures of local economic activity observable to the lender so that it could potentially be used in their pricing decisions. We define three such measures. Our primary measure of local economic activity

is the lagged delinquency rate on loans securitized by Fannie/Freddie in the MSA. Specifically, for MSA conditions in period t, we measure the fraction of loans originated in the local area between $t-k$ and $t-1$ that were at least 60 days delinquent at one time during that same period. Given that we are using quarter*year as our time interval, our base specification defines the lagged delinquency rate over the proceeding two years (quarter $t-9$ through quarter $t-1$).[8] We use lagged delinquency as our primary measure of local economic activity both because it is a summary statistic for many economic factors that could predict future default (weak local economic activity, declining house prices, etc.) and because it is easily observable by lenders. Our second and third measures of local economic activity are one-year lagged local housing price growth (between quarter $t-5$ and quarter $t-1$) and one-year lagged change in the local unemployment rate (between quarter $t-5$ and quarter $t-1$).[9] Appendix Table A3 shows the mean and standard deviation of each of our three local economic condition measures separately for our four time periods.

To describe the data, we run the following specification:

$$R_{st} = \omega_0 + \omega_1 E_{s,t-1} + \nu_{st}$$

where R_{st} is the average residualized MSA mortgage rate (defined above) and $E_{s,t-1}$ is lagged MSA measure of economic activity. The four panels of Figure 2 plot the relationship between the residualized MSA mortgage rate and the lagged Fannie/Freddie mortgage delinquency rate estimated separately for each of our four time periods. Each observation in each panel is a MSA. Given we averaged the residuals over the respective time periods, each MSA is represented only once in each figure. To fit a line through the observations, we weight each observation by the number of loans in the MSA during the relevant time period. Figure 2 shows that there is no relationship at all between lagged local default and the interest rate charged on mortgages (holding borrower and loan characteristics fixed) during the four periods. For example, during the 2001-2003 period (panel A) the scatter plot shows that a 1 percentage point change in lagged default rates increase mortgage rates by 0.2 basis points. The patterns in the 2001-2003 period are not anomalous. In all time periods, there is no economically meaningful relationship between lagged local default rates and local mortgage rates.

Table 2 reports the coefficients from the above regression where again each observation is weighted by the number of loans in the sample during the relevant time period. Each entry in Table 2 is from a separate regression. Specifically, for each of our four time periods and for each of our three local economic condition measures, we estimate a separate regression. The top entry is the estimate

[8] Throughout the paper, 60+ days delinquent is our primary measure of default. Roughly xx percent of mortgages in our Fannie/Freddie sample that were 60 days delinquent became 90 days delinquent one month later.

[9] While the lagged delinquency measure is defined within our Fannie/Freddie data, we take measures of lagged MSA house price growth from FHFA and the lagged change in MSA unemployment rate from the BLS's Local Area Unemployment Statistics. A more detailed discussion of both datasets can be found in the Data Appendix.

of ω_1, the second entry (in parentheses) is the standard error of the estimate, and the third entry (in brackets) is the marginal effect of a two standard deviation change in the relevant local economic condition measure. Again, consider the measure of lagged delinquency. During the 2007-2009 and 2010-2012 periods there is a statistically significant relationship but the magnitudes are extremely small. This is different from the 2001-2003 period and the 2004-2006 period. For example, during the 2007-2009 period a two standard deviation change in the lagged delinquency rate (about 1.4 percentage points) leads to a 3.8 basis points difference in mortgage rates. When lagged house price growth is our measure of economic activity, the results are very similar. In the 2000-2003 period and the 2004-2006 period, higher levels of house price growth were associated with slightly *higher* mortgage rates (2.6 basis points in response to a two standard deviation housing price growth change). As we discuss below, this finding is puzzling in that high house price growth predicts lower actual defaults going forward during these time periods. In the late 2000s, house price growth was associated with slightly lower mortgage rates.

To summarize our first key empirical fact there is essentially no variation between lagged local economic conditions and the local interest rate charged to borrowers on their mortgage. Overall, we show that the large economic variation across MSAs in local economic conditions was associated with at most only a 4 basis point difference in mortgage rates across MSAs. In the following section, we will show that lagged local default is a strong predict of future default - particularly in periods of aggregate recessions (2001-2003 and 2007-2009). This finding will make the lack of interest rate variation across regions seem puzzling.

Before proceeding, it is worthwhile to explore the robustness of our key results using other data sources. As mentioned earlier, one draw back of the Fannie/Freddie data is that it does not track the fees and points associated with the loan. It is possible that in places where economic conditions are worse lenders charge higher fees/points for a given loan.[10] To explore whether our results hold up accounting for points and fees, we perform two robustness exercises. First, we use data from LoanSifter, Inc. LoanSifter is a leading provider of point-of-sale pricing information for mortgage lenders. In particular, it solicits MSA-specific mortgage offer quotes for a specific loan type in from a specific lender.[11] One product they provide is the average mortgage rate being offered for a specific loan type during a given week in a given MSA. The loan type is defined by the loan amount, the loan's term, the amount of points associated with the loan, the borrower's FICO score, and the loan to value ratio at the time of origination. We have access to two such loan types during the period spanning September 2009 and November 2010. The first is all 30 year fixed rate mortgages with no points, a loan to value ratio of at least 80 percent, a loan origination amount of xxxxx, and the borrower's FICO score of 750. The

[10] Points are upfront fees associated with a mortgage that are expressed as a percentage of the loan origination amount. The existence of points and other fees drives a wedge between the mortgage rate and the APR (annual percentage rate).

[11] A more detailed discussion of the LoanSifter data is provided in the Data Appendix.

second is a ..... [12]

Figure 3 shows the relationship between the average residual mortgage rate and the average lagged default for all xxx MSAs in the LoanSifter sample. The interest rates where only residualized with respect to month of origination dummies (given the FICO score and LTV are held fixed within each mortgage type. The results confirm the patterns in Figure 2 and Table 2. There is no relationship at all between lagged economic conditions and mortgage rates even holding potential loan fees fixed. The results in Figure 3 also confirm that the spatial patterns in the actual transaction data shows up as well in the loan offer rates.

As a second robustness exercise, we explored the patterns in the Freddie Mac Primary Mortgage Market Survey (PMMS). The strength of the PMMS data set is that it measures the actual APR inclusive of points and fees at the state level (as opposed to just the mortgage contract rate). The downside of this data set is that it does not include controls for borrower FICO or LTV. In the online appendix that accompanies this paper, we show that APR does not vary with local economic conditions once we control for state differences in FICO scores and LTV. Again, these results further suggest that our results are robust to controlling for any potential spatial difference is points and fees.

4.1.1 Jumbo Mortgage Rates and Current Local Economic Conditions

Above, we show that conventional mortgage rates do not vary with local economic conditions. One natural question is to ask whether mortgage rates not securitized by Fannie/Freddie vary with local economic conditions. To do this, we focus on our sample of prime jumbo mortgages. Figure 4 shows the spatial relationship between lagged default rates and residualized mortgage rates for the sample of prime jumbo mortgages during the 2001-2003 period.[13] The lagged default rate is based off the sample of Fannie/Freddie loans. Given that, the variation in the x-axis across MSAs for Figure 4 is identical to the variation in the x-axis for panel A of Figure 2. The prime jumbo mortgage rates are residualized in exactly the same manner described above for the Fannie/Freddie mortgage rates.

Figure 4 shows that the residualized jumbo mortgages rates are strongly increasing in lagged local default rates during the 2001-2003 period. In partic-

[12] We used the Federal Reserve Board of Governors' subscription to this data for these robustness exercises. We had access to this data when one of the authors was working at the Board of Governors.

[13] To conserve space, we relegate the discussion of spatial variation in the prime jumbo mortgage rate during the 2004-2006 period to the online appendix. In summary, there was very little variation in mortgage rates during this period. This is not surprising given that there was also very little relationship between lagged defaults and future defaults during this time period. The link between lagged economic activity and future defaults is highest in both the Fannie/Freddie data and the prime jumbo data during periods of aggregate recessions (2001-2003 and 2007-2009). As noted above, given that the prime jumbo mortgage dried up after 2007, we cannot perform the same spatial analysis during either the 2007-2009 period or the 2010-2012 period.

ular, the relationship shows that a 1 percentage point increase in lagged default increases the jumbo mortgage rate by 32 basis points (p-value ¡ 0.01) all else equal. [xx discuss regressions for lagged unemployment rate and lagged house price growth]. These results stand in sharp contrast relative to the loans securitized by Fannie/Freddie. During the 2001-2003 period, lagged local economic activity is strongly predictive of local mortgage rates. These results already control for differences in borrower credit score and LTV which may differ across the types of loans in the different areas. The main difference between the Fannie/Freddie loans and the prime jumbo loans is that one is securitized by the private market and the other is securitized by the GSEs.

5 Should Local Mortgage Rates Vary With Lagged Local Economic Conditions?

5.1 Do Lagged Economics Conditions Predict Future Mortgage Default Rates?

Why would interest rates vary across regions in a way that is correlated with lagged economic activity? If there is medium run persistence in local economic activity, lagged local economic activity will predict future local economic activity. If a household's decision to default on their mortgage is correlated with local economic activity, mortgage rates should potentially vary with lagged local economic activity. This was the case for the jumbo prime mortgages but not the loans securitized by the GSEs.

To examine whether future mortgage default rates are predicted by lagged local economic activity, we estimate the following specification:

$$y_{ist} = \alpha_2 + \alpha_2 X_{it} + \Gamma_2 D_t + \zeta_{ist}$$

where y_{ist} is the ex-post default rate for a loan made to borrower i, in location s, during period t and where X_{it} and D_t are defined as above. Our measure of default is whether the loan becomes 60 days delinquent at any time during the subsequent two years after origination. Akin to what we did with mortgage rates, we run this regression separately using the Fannie/Freddie data for three separate time periods: 2001-2003, 2004-2006, and 2007-2009. We do not run if for the 2010-2012 period because we were not able to follow the loans forward for two years to track their delinquency behavior. For the prime jumbo loans, we only run this regression for the 2001-2003 period and for the 2004-2006 period. Y_{st} is the average residual default rate for each state during each of the the time periods such that:

$$Y_{st} = \frac{1}{N_{st}} \sum_{i=1}^{N_{st}} \zeta_{ist}$$

Table 3 shows the mean ex-post default rate (unconditional), the standard deviation of unconditional ex-post default rates (y_{ist}'s) across MSAs, and the

standard deviation of conditional mortgage rates (Y_{st}'s) across the MSAs for each time period of study. [xxx-Add something of interest from Table 3??]

Figure 5 shows the relationship between lagged local default rates and future residual default rates for the Fannie/Freddie data. The definition of the lagged local default rates is identical to the definition used in Figures 2 and 4. Each panel of Figure 5 is for a different time periods. As above, each MSA is in each figure only once and we weight the MSA by the number of loans in the sample when fitting a line. The key finding is that during both the 2001-2003 period and during the 2007-2009 period, lagged local default is a strong predictor of subsequent ex-post residual default. A 1 percentage point increase lagged default increased the ex-post residual default rate by 0.7 percentage points during 2001-2003 and 0.4 percentage points during 2007-2009 (p-value of both ¡ 0.01). Given that the actual average ex-post default rate for loans originated in 2001-2003 and 2007-2009 were x percent and y percent, respectively, the estimates imply that a one-standard deviation change in lagged local default results in a roughly x percent and y percent increase in actual ex-post local default, all else equal. Notice, during the 2004-2006 period, there is no relationship between lagged default and ex-post future default. This was, in part, because lagged defaults were quite low on average. These results suggest that in periods when defaults are higher (like during aggregate recessions), there is more persistence in local default rates.

Figure 6 shows the same relationship between lagged local default and ex-post residual default rates for the prime-jumbo sample of loans for the 2001-2003 and 2004-2006 periods. Like above, the measure of lagged local default is computed using the Fannie/Freddie data so that the x-axis of Figure 6 and Figure 5 are identical for common time periods. During the 2001-2003 period, lagged local default rates are highly predictive of future defaults in the jumbo market. [finish describing results when we see the results]

Table XXX is similar in structure to Table 2. Each entry in the table represents a different estimate of ψ_2 from the following regression:

$$Y_{st} = \psi_0 + \psi_1 E_{s,t-1} + \varepsilon_{st}$$

The different entries in the table represent different measures of $E_{s,t-1}$, different time periods, and different samples (Freddie/Fannie sample or prime-jumbo sample). As seen in Table XXX, during the periods of aggregate recessions (2001-2003 and 2007-2009) lagged local economic activity - whether it be lagged defaults or lagged house price growth - is a strong predictor of future local defaults.

5.1.1 How Much Should Conventional Mortgage Rates Vary With Current Local Economic Conditions?

In this sub-section, we use two methods to assess how much systematic spatial variation in mortgage rates we should have expected to see in the Fannie/Freddie loans given that local conditions predict future default. Given that during periods of aggregate recessions there is some medium run persistence in local

economic activity, default risk varies spatially in a predictable way. Our first method to assess how much GSE interest rates should have varied with lagged local economic activity is to use the prime-jumbo loans as a benchmark.

Specifically, we estimate:

$$Y_{st}^j = \psi_0^j + \psi_1^j E_{s,t-1} + \varepsilon_{st}^j$$

where Y_{st} and $E_{s,t-1}$ are defined as a above. The j super-script refers to the sample on which the regression is run. Formally, $j = \{F, P\}$ where F designates the Fannie/Freddie sample and P designates the prime jumbo sample. We run this regression only for the 2001-2003 time period which is a period where data exists for both samples and where lagged economic activity strongly predicts future defaults. After running the above regression, we form a predicted measure of local residual defaults for both samples such that:

$$\hat{Y}_{st}^j = \psi_0^j + \psi_1^j \hat{E}_{s,t-1}$$

Our key specification then runs local residual mortgage rates on predicted measures of local residual default:

$$R_{st}^j = \beta_0^j + \beta_1^j \hat{Y}_{st}^j + \epsilon_{s,t}^j$$

Table YYY shows the estimates of β_1^F and β_1^P. [xxx Describe Results xxx]

The difference between β_1^F and β_1^P tell us the differential response between the prime jumbo loan market and the Fannie/Freddie market to local differences in predicted future default rates. [Describe counterfactual].

As a robustness exercise, we perform a separate counterfactual to explore how mortgage rates on loans secured by the GSEs should have responded to observed patterns in predicted default. We consider a loan made in perpetuity. We further suppose that the loan defaulted quickly after origination. Suppose the borrower put down 20 percent at the time of the loan. We also assume that the recovery rate on the the loan, conditional on default, is either 70 percent or 50 percent.[14] If this loan defaulted with a 3.0 percent probability (roughly a two standard deviation in the cross-MSA predicted default rates), the default premium on this loan should between 38 basis points (with a 70 percent recovery rate) and 106 basis points (with a 50 percent recovery). These estimates are in the ball-park of the differences in mortgage rates across MSAs in the prime-jumbo market in response to a x.x percent probability of predicted default. (xxx - use the same x.x as above).

A key parameter in our empirical exercise below is the extent to which mortgage rates should very with local economic activity. [discuss how we calibrate this parameter]

[14] Put in some cites showing that these are plausible recovery rate assumptions.

5.2 Why Conventional Mortgage Rates Do Not Vary With Current Local Economic Conditions?

Why does the mortgage rates on loans sold to the private market vary with local economic conditions but the mortgage rates on loans sold to GSEs do not? The quasi public nature of the GSEs may impose political economy constraints on the extent to which they can vary mortgage rates across space.[15] There is some evidence for this political economy story. In early 2008, the GSEs attempted to implement a mortgage policy that restricted credit differentially across U.S. locations. The policy was dubbed a "Declining Market Policy". The policy required more equity at the time of origination in markets for which house prices were declining. In non-declining markets, Fannie/Freddie would purchase mortgages where the initial LTV was lower than 0.95. However, in declining markets, Fannie/Freddie would only purchase mortgages where the initial LTV was lower than 0.9.[16] The policy did not affect interest rates, it only affected underwriting standards.

The policy was announced in December of 2007 and was implemented in mid-January of 2008. After receiving large amounts of backlash from a varied set of constituents, the policy was abruptly abandoned in May of 2008. Consumer advocacy groups rallied against the policy arguing that it was a form of space based discrimination.[17] Real estate trade organizations used their political clout to protest the policy because it was hurting business. For example, the *Wall Street Journal* summarized the GSEs abandoning the declining market policy by saying "The change [in GSE policy] comes in response to protests from vital political allies of the government-sponsored provider of funding for mortgages, including the National Association of Realtors, the National Association of Home Builders and organizations that promote affordable housing for low-income people".[18] The Washington Post reported that "Critics, including the National Association of Realtors and consumer advocacy groups, had charged that Fannie Mae's policy served to further depress sales and real estate values in areas tainted as declining".[19] Even though it may have been profitable to require different downpayments in different areas, Fannie Mae and Freddie Mac succumbed to political pressure and abandoned the policy.

Even though the policy focused on imposing spatial variation in downpayments, it can shed light on reasons why the GSEs do not raise interest rates in

[15] Many government agencies do not engage in space based pricing. For example, it costs the same to mail a letter through the U.S. Post Office from New York to New Jersey as it does to mail the same letter from New York to Alaska. If that letter was mailed using private companies like FedEx or UPS the price from New York to Alaska would exceed the price from New York to New Jersey.

[16] Fannie/Freddie had slightly different definitions for what was a declining market. Roughly, declining markets were defined as locations where house prices were declining over the last two to four quarters.

[17] See "Fannie Mae Sets New Loan Boundaries" from *NPR's Marketplace* on April 18, 2008. http://www.marketplace.org/topics/world/fannie-mae-sets-new-loan-boundaries

[18] See "Fannie Is Poised to Scrap Policy Over Down Payments", *Wall Street Journal*, May 16, 2008.

[19] See "Looser Credit on the Way in 'Declining' Markets?" *Washington Post*, May 24, 2008.

riskier markets during the recession. The source of the pushback on charging different interest rates across locations would likely have been the same. Interestingly, the argument against the declining market policy was that it would further hurt the depressed areas by further reducing mortgage activity. This is exactly the mechanism we will highlight and quantify in the next sections for the constant interest rate policy. By foregoing profit maximizing behavior and charging a constant interest rate across all regions despite different levels of predictable default risk, the GSEs are redistributing resources towards the declining markets.

Before continuing, we want to highlight that even among private firms there is a reluctance to use space based pricing when it is profit maximizing to do so. Finkelstein and Poterba (2013) document a lack of spatial risk-based pricing in UK insurance markets. They show that geography predicts longevity but geography is not priced in British annuities. They speculate that fear of consumer/regulatory backlash is the main reason the insurance companies do not use geography in their pricing decisions.

6 Model description

We now assess how the lack of regional variation in interest rates affect households living in different locations by building a multi-region household model with life-cycle consumption and housing choices.

In the model, households are subject to region specific shocks in addition to purely idiosyncratic income shocks. Incomplete markets imply that these shocks are not fully insured. Since the previous sections have shown that GSE mortgage rates do not respond to regional shocks, we initially assume that there is no regional variation in mortgage rates and calibrate the model to match various features of the data. We then use the model to explore what would happen if the constant interest rate policy was removed so that mortgage rates can vary with local economic conditions.

Our model allows for regional variation in mortgage rates to affect welfare through two key channels: 1) We assume that households are able to borrow against their houses subject to holding some minimum equity. 2) Households typically borrow all but the required down payment when purchasing houses. If interest rates rise when local conditions deteriorate, the first channel lowers welfare by making it more difficult to smooth consumption. In addition, the second channel means that households in bad regions delay purchasing housing and reduce the sizes of their eventual purchases.

6.1 Demographics and Location

The economy is characterized by a continuum of households indexed by i. Household age is indexed by $j = 1, ...J$. Households enter the labor-force

at age 25 and retire at age 60. After retirement, households face stochastic mortality risk but live to a maximum age of 85.

Households live in specific regions indexed by k, and in our baseline results we assume that households never move. Regional economic activity in region k at period t is given by

$$\log \gamma_{k,t} = \rho_\gamma \log \gamma_{k,t-1} + \varepsilon_{k,t}.$$

The effect of this regional shock $\gamma_{k,t}$ on other aspects of the model will be made concrete as we describe those pieces.

6.2 Preferences and Household Choices

Households receive flow utility

$$U_{ijk} = \frac{\left(c_{ijk}^{\alpha} h_{ijk}^{1-\alpha}\right)^{1-\sigma}}{1-\sigma}$$

from non-durable consumption c_{ijk} and housing services h_{ijk}. See Piazessi et al 2007 for evidence that this Cobb-Douglas specification is reasonable. Households discount expected flow utility over their remaining lifetimes with discount factor β.

6.3 Income Shocks

Time t household labor earnings for working age households are given by

$$\begin{aligned} \log y_{ijk,t} &= \chi_j + z_{i,t} + \phi^y \gamma_{k,t} \\ \log z_{i,t} &= \rho_z \log z_{i,t-1} + \eta_{i,t} \end{aligned}$$

where χ_j is a deterministic age profile common to all households, $z_{i,t}$ is a purely idiosyncratic persistence income shock and $\phi^y \gamma_{k,t}$ is a region-specific shock to income. ϕ^y is a parameter that governs the sensitivity of household income to local economic conditions.

When retired, households receive Social Security benefits. We describe the computation of these benefits in the calibration section of the model.

6.4 Housing Markets and Interest Rates

Housing services can be obtained from owner-occupied housing or through a rental market. Housing can be purchased at price $p_{k,t} = (\gamma_{k,t})^{\phi_h}$ or rented at price $p_{k,t} r^f$.[20] We denote owner occupied houses as $h_{i,t}$ and rented houses as $h_{i,t}^f$.

[20] Note that this specification imposes a constant price-rent ratio.

Owning is subject to a fixed adjustment cost:

$$F_{ik,t} = \begin{cases} FP_{k,t} \text{ if } h_{i,t+1} \neq h_{i,t} \\ 0 \text{ if } h_{i,t+1} = h_{i,t} \ . \end{cases}$$

In addition to this cost, owning has two benefits over renting. First, households can borrow against houses subject to a minimum equity requirement.

$$m_{ik,t} \leq (1-\theta)\, p_{k,t} h_{i,t},$$

where θ is the minimum down payment or equity that must be held in the house. Second, we assume that the rental stock depreciates at rate $\delta^f > \delta^h$. In a competitive equilibrium the rental price of housing must be equal to the risk free rate plus the rate of depreciation of the rental stock:

$$r^f = r + \delta^f.$$

See Luengo Prado and Diaz 2010. Thus, $\delta^f > \delta^h$ implies that the imputed rental price of owner occupied housing is cheaper than that of renting.

The interest rate paid on mortgages is equal to the risk free rate plus a risk premium

$$r^m_{k,t} = r + \Psi_{k,t},$$

where the risk-premium is declining in regional economic activity:

$$\log \Psi_{k,t} - \log \overline{\Psi} - \phi^r \log \gamma_{k,t}.$$

In addition to borrowing through mortgages and saving through the purchase of durable housing, households can save in a one-period bond b with risk-free rate r.

6.5 Household Problem

The household model is solved recursively. Let $\mathbf{s}_{jk} = (b_j, m_j, h_j, z_j; \gamma_{jk})$ be the state vector for a household, where b_j, m_j, h_j reflect start of the period values before household's decisions. Note that h_j is the housing stock at the start of the period before depreciation, and that we make a timing assumption that households get utility from the house that they choose today rather than the house that they start the period with. For notational convenience, we index time solely by age. Since there are no aggregate shocks, separately tracking t does not change the solution.

In each period prior to retirement, the household solves

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

with

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

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

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

and a household that chooses to sell its current house[21] and rent has value function

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

[21] If 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.

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. A computational appendix discusses the numerical solution of the model.

7 Calibration

Our benchmark calibration strategy proceeds in two parts: 1) For parameters that do not depend on regional economic activity, we calibrate parameters to standard values from the literature or to match standard moments from wealth data. 2) For parameters that do vary regionally, we calibrate to match estimates from the previous section. Our model period is one-year and we calibrate the model accordingly.

7.1 Standard Parameters

Following Floden and Linde (2001), we set $\rho_z = 0.91$ and $\sigma_\eta = 0.21$ to match the annual persistence and standard deviation of residual earnings in the PSID.

During retirement, households receive social security benefits which we calculate using the method of Guvenen and Smith (2013). In reality Social Security Benefits are a function of life-time earnings, but this would substantially complicate the solution of the model as these life-time earnings would become a state variable. However, a relatively accurate measure of lifetime earnings can be imputed from earnings in the final period of working life given the persistence of the income process. Thus we forecast life-time income given income in the final period of working and then apply the actual benefits ratios from Social Security charts to this imputed lifetime income.

As is standard in the risk-sharing literature, we set $\sigma = 2$ to generate an intertemporal elasticity of substitution of 1/2. Our model period is annual and we set the risk-free rate to $r = 0.03$ to roughly match the average one-year treasury bill rate in the 2000s. In addition we calibrate an average risk-premium of 0.01. We calibrate $\delta_h = 0.03$ to match the average ratio of residential investment to the residential stock in BEA data. We set $\theta = 0.25$ so that households are required to have a minimum 25% down payment. (See discussion in Diaz and Luengo Prado for justification). We pick $F = .05$ so that there is a 5% transaction cost from adjusting housing.

We jointly pick β, r^f and α to match various wealth and home ownership targets.[22] In particular, we target a home-ownership rate of 69% as in SCF data. We also target the median wealth-to-income ratio of 1.52 from SCF data. (cite Kaplan and Violante discussion). Finally, using BEA data, we target a ratio of durable expenditure to housing expenditures of 15. These targets yield $\beta = 0.93$, $r^f = 0.072$ and $\alpha = 0.88$. Since $r^f = r + \delta^f$, this rental rate for

[22]Note that we pick parameters to match these targets under the assumption that $\phi^r = 0$, which corresponds to the data generating process under current policy.

housing implies $\delta^f = .032$ so that the rental stock depreciates roughly 7% faster than the owner occupied housing stock. This is similar to the numbers in Diaz and Luengo Prado.

7.2 Calibrating Regional Variation

In addition to these more standard parameters, we must calibrate parameters that vary with regional economic conditions. Our calibration strategy uses lagged-default as a broad measure of economic activity in a region. By measuring $\log \gamma$ using lagged-default[23], we can then compute all necessary elasticities by computing the empirical elasticities with respect to lagged-default. We set $\rho_\gamma = 0.788$ and σ_ε to match the behavior of lagged-default. To calculate ϕ^h we regress log MSA house prices on lagged-default with MSA and quarter fixed effects to arrive at a value of 1.67. Similarly, we calibrate a value of $\phi^y = 2.02$ by performing a similar regression using the unemployment rate. We show robustness results with respect to both of these parameters.

The key elasticity for our policy experiment is ϕ^r. That is, how much would interest rates respond to lagged-default in a world with no GSE intervention? Regressing jumbo mortgage rates on lagged-default in the early period yields a coefficient of 32.32. Applying this same relationship to GSE mortgages in 2007-2010 implies that a one-standard deviation increase in lagged-default should increase mortgage rates by 6.7%. This means that an increase in lagged-default of 0.0138 should increase r^m from 0.04 to 0.04268 ($= 1.0670 \times 0.04$). Since we assumed that $r^m_{k,j} = r + \overline{\Psi}\gamma^{-\phi_r}_{k,j}$ this implies that $\phi_r = 17.2$.[24] Again, we provide robustness results for both larger and smaller ϕ_r and discuss alternative counterfactuals based on FICO scores in the following section.

8 Results

For ease of discussion we label regions with low economic activity and high lagged-default "bad" regions and regions with high economic activity and low lagged-default "good" regions. We assume that in the absence of intervention from GSEs, mortgage rates would move with regional economic activity so that good regions would have low rates and bad regions would have high rates.

This implies that the constant interest rate policy will tend to make households in the bad regions better off and households in the good regions worse off. To assess the quantitative size of this "transfer" we ask how much households in a given region would be willing to pay to change from a variable interest rate policy to a constant interest rate policy. In particular, we solve the model with the variable interest rate and calculate how much additional consumption we would have to give households today to make them indifferent between the variable interest rate and the constant interest rate policy.

[23] Note that logγ=−lagged default as there should be a negative relationship.

[24] $0.03 + .01\left(e^{.0138}\right)^{17.2} = .04268$

Formally, let $V_j^{\text{constant r}}(\mathbf{s}_{jk})$ be the indirect utility obtained from solving the household problem with state $\mathbf{s}_{jk}$ in a world with $\phi_r = 0$. Similarly, let $V_j^{\text{variable r}}(\mathbf{s}_{jk})$ be the indirect utility obtained from solving the model in a world with $\phi_r > 0$, and let $\widetilde{c_{jk}}, \widetilde{h_{jk}}$ be the choice for non-durable consumption and housing services that obtain this maximal value. Finally, let $E_{\gamma,z,j}$ denote the expectation of these value functions over values of the idiosyncratic shock and age[25], conditional on living in a region with economic activity γ. We then solve for λ so that:

$$E_{\gamma,z,j} V_j^{\text{constant r}}(\mathbf{s}_{jk}) = E_{\gamma,z,j}^{\text{variable r}} \left\{ U(c\,(1+\lambda)\,, h) + \beta E_j V_{j+1}^{\text{variable r}}(\mathbf{s}_{jk}) \right\}.$$

That is, we compute the one-time percentage change in consumption that, in expectation, makes households indifferent between being in a world with constant r^m and a world with variable r^m.

Table 5 shows the implied values of λ for various regions. Clearly the worst regions (on the left side of the table) are made substantially better off by the constant interest rate policy while the best regions (on the right side of the table) are made substantially worse off.

The second row of Table 5 displays a back-of-the-envelope calculation that converts these one-time consumption equivalents to dollar values. To do this, we first estimate average consumption in dollars per household by dividing real non-durable consumption from the BEA by the number of households in the U.S.[26] This calculation gives that average household non-durable consumption is just over \$63,000. While this number represents the average consumption per household, in the model households in bad regions consume less than households in good regions. This implies that the same λ in a bad region represents a smaller amount of consumption in dollars than in a bad region, so we account for these differences.

Overall our results imply fairly large transfer payments across regions. Thus, this constant interest rate policy has significant redistributional consequences across regions. We will show shortly that this conclusion is quite robust to a range of alternative model parameters. In contrast, the overall welfare effects of the policy are more sensitive to the particular calibration of the model. With concave utility, if the variable interest rate resulted in a pure mean preserving

[25] In all calculations unless otherwise noted, we focus on the utility of working age households since we want to understand how the mortgage market interacts with household risk. Retired households face no such labor market risk. Furthermore, our model abstracts from many features which are important for understanding end of life behavior. The effects of variable interest rates on the retired are quite sensitive to the behavior of households in the terminal period, and our model is ill-suited for matching end of life behavior. Reassuringly, solving the model with retirement periods of various different lengths remains the conclusions for working age households unchanged as the behavior in this final period is essentially irrelevant for working age households.

[26] Total non-durable consumption in 2012 was \$7335.9 billion = (2296.8 spending on non-durable goods +6982.7 spending on services - 1943.6 spending on housing services). The census bureau estimates that there are approximately 115 million households in the U.S. Dividing the first number by the second gives \$63,790 per household.

spread in consumption, it would necessarily lower ex-ante welfare. Evidently, this is not the case. In some parameterizations, the overall welfare effect of the constant interest rate policy when adding across regions is positive while in other parameterizations it is actually negative. For some parameter values, bad regions are made more worse-off by the variable interest rate than good regions are made better-off. However, for other parameter values this reverses. For our benchmark results, we find an overall consumption equivalent of -.04% so that the constant interest rate policy very mildly reduces welfare for the average household in the economy.[27]

Given these results, we now return to analyzing regional redistribution. In our model, when mortgage rates rise with deteriorating economic conditions this lowers welfare through two main channels. 1) Housing equity represents a significant fraction of overall household wealth. 2) Most households borrow all but the required down payment when purchasing their first house. If interest rates rise when local conditions deteriorate, the first channel lowers welfare by making it more difficult to smooth consumption. In addition, the second channel means that households in bad regions delay purchasing housing and reduce the sizes of their eventual purchases.

Interestingly these two channels interact differently with households of different ages. Young households typically have little housing equity, so the first channel is largely irrelevant. At the same time, the second channel is most relevant for the young as they consider purchasing their first houses. Conversely, the first channel is substantially more important for the middle-aged. Table 6 redoes the calculation in Table 5 but conditions on age. That is, we look at how much young households and middle aged households in different regions are willing to pay to move to a constant interest rate policy.[28]

Overall we find that the welfare of the middle aged is substantially more sensitive to eliminating the constant interest rate policy. This suggests that in general, the first channel is more important than the second. Important for this result is that we allow for a rental market.

Table 7 shows results when we shut down the rental market and force all households to purchase houses. Without a rental market, the previous results reverse and the young in poor regions are worse off with a variable interest rate than are the old. This is because without a rental market, all households are forced to purchase houses, which are still subject to transaction costs. If the interest rate varies with economic conditions this makes the young particularly bad off: they are forced to purchase when their income is low and interest rates are high.

[27] Note that this number is not the sum of the numbers in row 1 of Table 5 for two reasons: 1) There are not equal number of households in regions of each type. Households are more likely to be in an "average" region than in an extremely good or an extremely "bad" region. 2) As previously mentioned, the average level of consumption differs across regions, so the curvature of the utility function and thus how sensitive utility is to changes in consumption varies across regions.

[28] We define young as 25-35 as this is the primary range for first-time home purchases, and we define middle aged as 36-60.

In the Appendix, we discuss the robustness of our findings to different elasticities and choices of parameters. The overall conclusion from this robustness analysis, as can be seen, is that our findings are stable across reasonable range of elasticities.

9 Conclusion

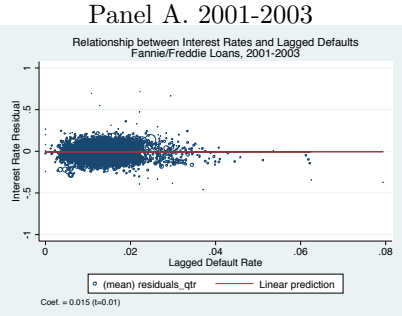
To be added

Figure 1: Share of Mortgage Securitized

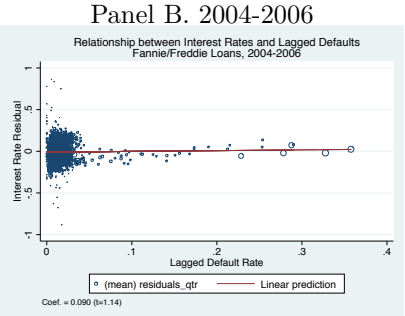
Note: Figure shows the proportion of mortgages securitized by type for years xxxx-yyyy.
Source: Inside Mortgage Finance.

Figure 2: Relationship between Interest Rates and Lagged Local Default, GSE Loans, 2001-2003 and 2004-2006

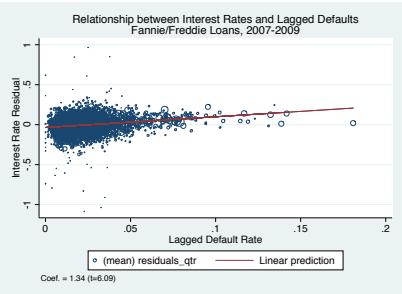
Panel A. 2001-2003



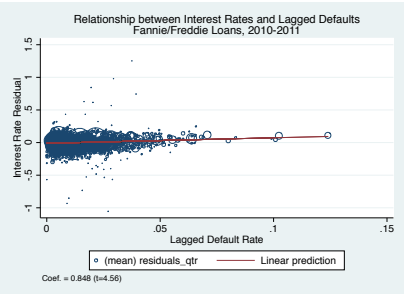
Panel B. 2004-2006



Panel C. 2007-2009

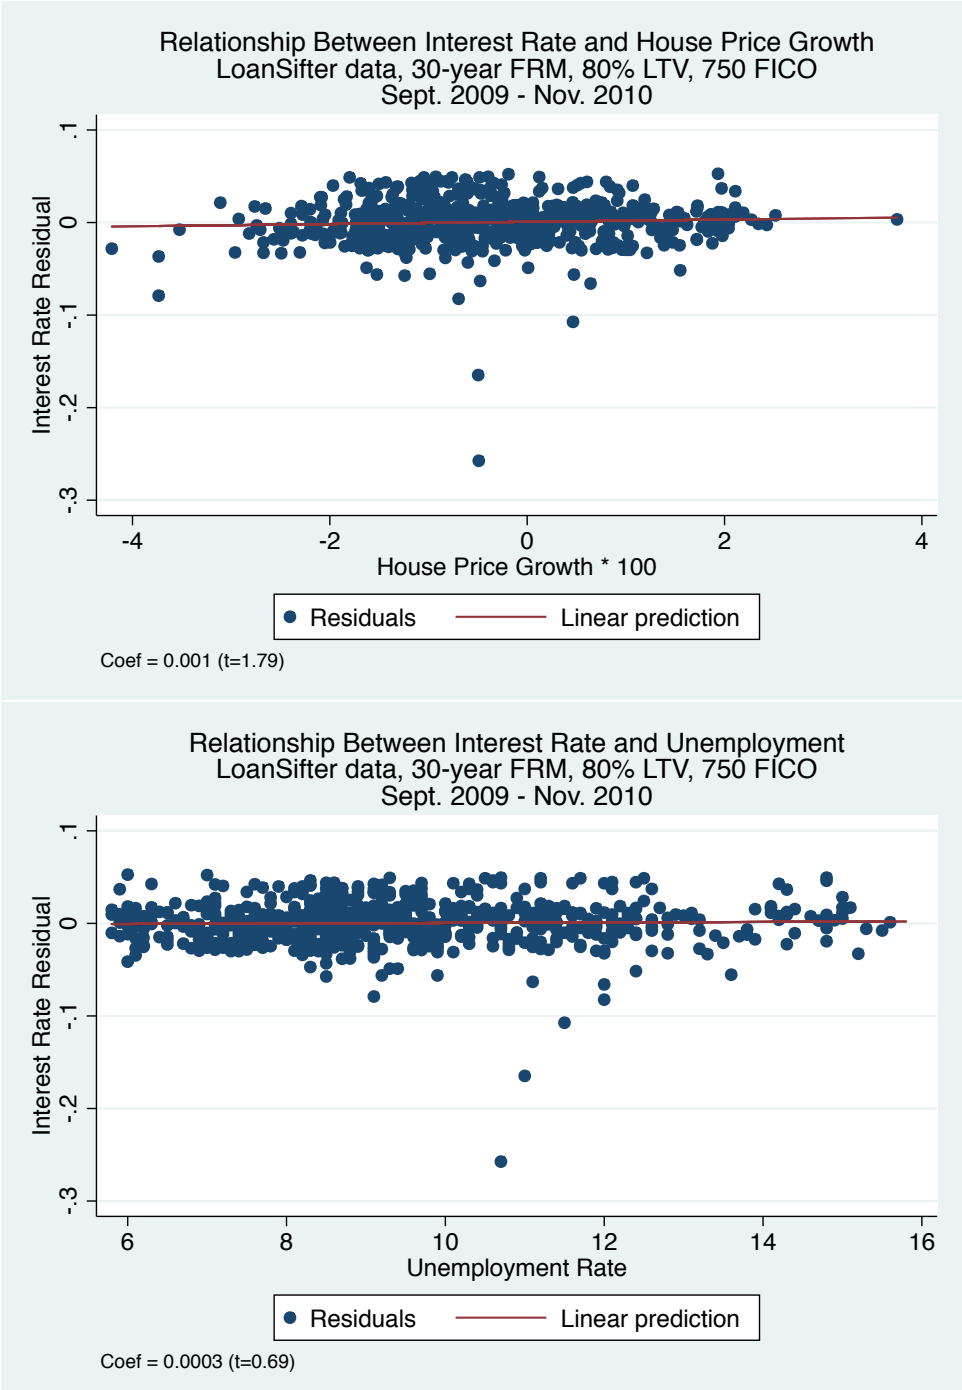


Panel D. 2010-2012



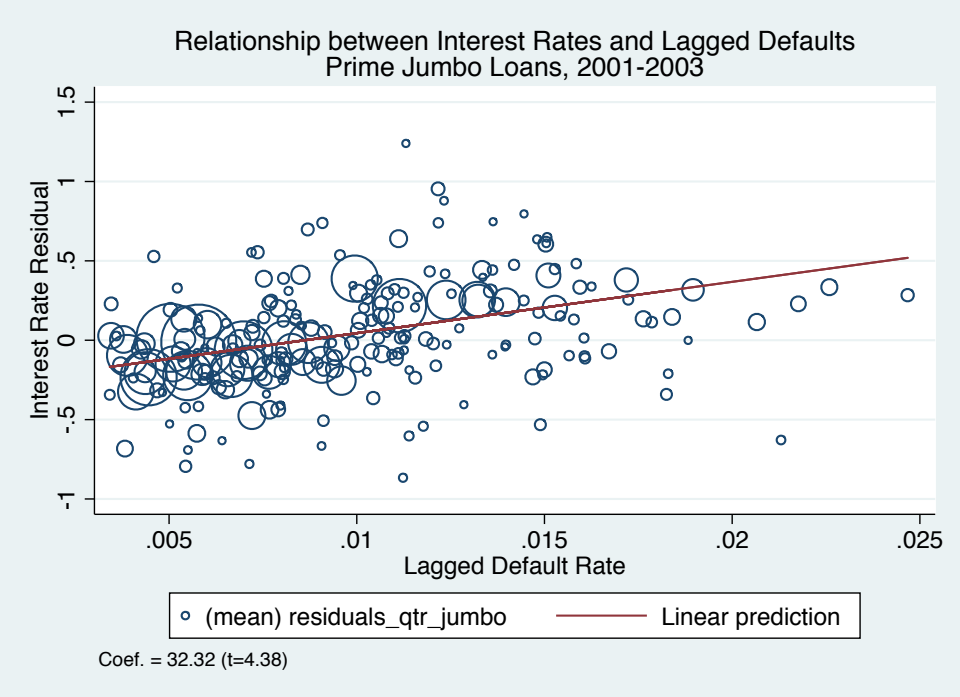
Note: Figure shows the relationship between residualized interest rates and lagged MSA-level default of loans originated within the last two years for four time periods. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 3: Relationship between Offered Interest Rates and Local Economic Conditions, 2009-2010



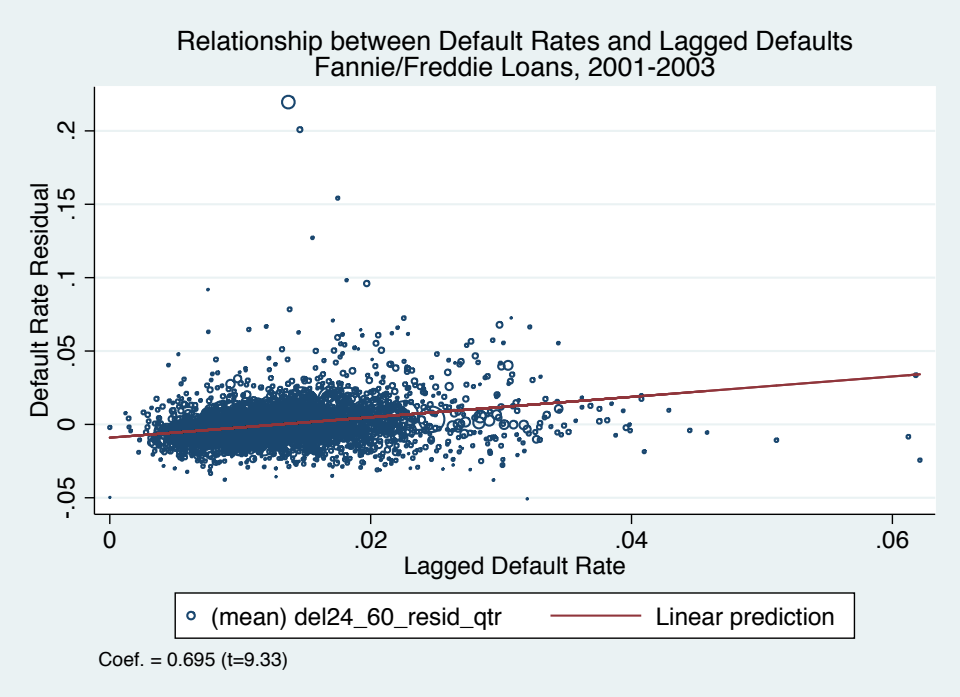
Note: Figure shows the relationship between residualized interest rates and lagged MSA-level economic conditions, recent house price growth and recent unemployment rates. Source: Authors' calculations using LoanSifter Offer Database.

Figure 4: Relationship between Interest Rates and Lagged Local Default, Non-GSE "Prime Jumbo" Loans, 2001-2003



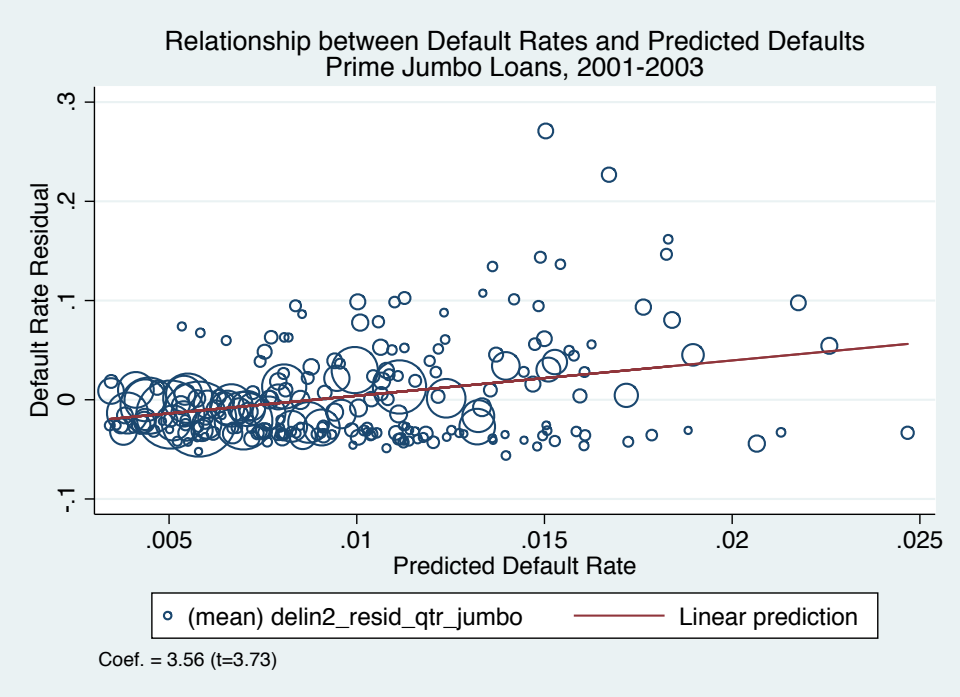
Note: Figure shows the relationship between Prime Jumbo residualized interest rates and lagged MSA-level default of loans originated within the last two years for two time periods: 2007-2009 and 2010-2012. Source: Authors' calculations using Loan Performance Database.

Figure 5: Relationship between Default Rates and Lagged Local Default, GSE Loans, 2001-2003



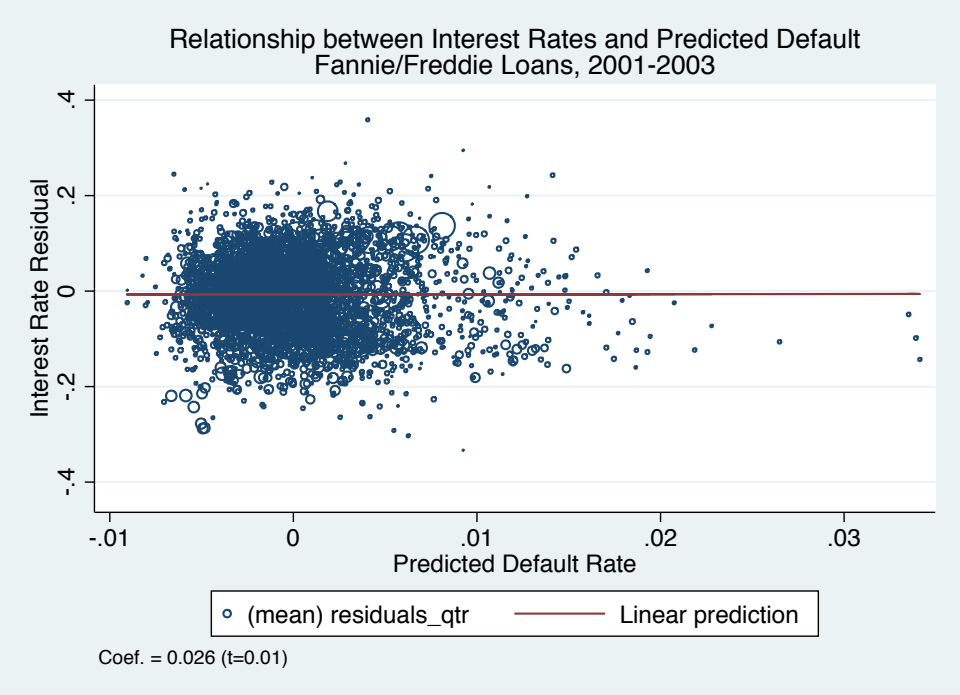
Note: Figure shows the relationship between residualized default rates and lagged MSA-level default of loans originated within the last two years for 2001-2003. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 6: Relationship between Default Rates and Lagged Local Default, Non-GSE Prime Jumbo Loans, 2001-2003



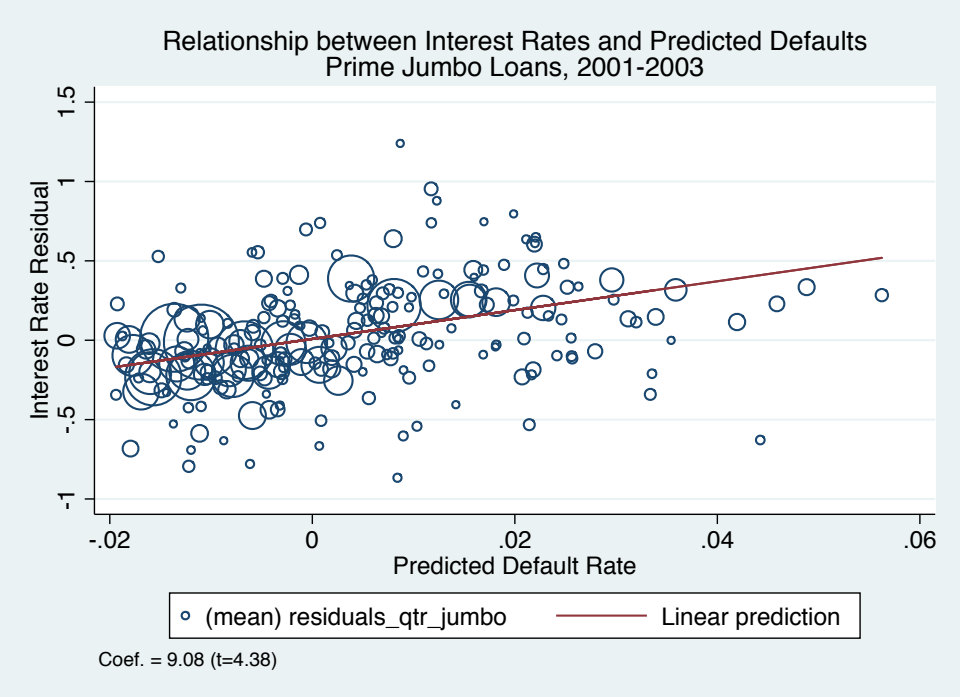
Note: Figure shows the relationship between residualized default rates and lagged MSA-level default of loans originated within the last two years for 2001-2003. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 7: Relationship between Interest Rates and Predicted Default Rates, GSE Loans, 2001-2003



Note: Figure shows the relationship between residualized interest rates and predicted default rates for loans originated in GSE market, 2001-2003 period. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 8: Relationship between Interest Rates and Predicted Default Rates, Non-GSE "Prime Jumbo" Loans, 2001-2003



Note: Figure shows the relationship between Prime Jumbo residualized interest rates and predicted default rates for loans originated in non-GSE market, 2001-2003 period. Source: Authors' calculations using Loan Performance Database.

Table 1: Determinants and Dispersion of Mortgage Rates, GSE Loans, 2001-2012

	All years	2001-2003	2004-2006	2007-2009	2010-2012
Regression with year*qtr fixed effects and interactions for LTV, LTV^2, FICO, and $FICO^2$					
R^2	0.8882	0.7313	0.4972	0.7239	0.5583
N	9591457	3332022	2113663	1714095	866301
Regression with year*qtr fixed effects					
R^2	0.8781	0.7117	0.4591	0.6799	0.4973
N	9648766	3359362	2119977	1716193	866805
Mean Interest Rate	6.09	6.58	6.02	5.91	4.75
SD Across MSAs and quarters	0.732	0.579	0.3	0.557	0.275
SD across MSAs adjusted*	0.073	0.077	0.07	0.073	0.057

Note: The table shows r-squared and sample sizes for regressions exploring the relationship between mortgage rates, time dummies (quarterly), LTV, and FICO scores. The lower panel of the table shows how much variation exists across MSAs, both after adjusting for time dummies and not. Adjusted residual removes year*quarter fixed effects and quadratic controls for FICO and LTV. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Table 2: Relationship between Interest Rates and Lagged MSA-level Default, GSE and Prime Jumbo Loans

Dependent variable: Interest rate residual					
	GSEs				Non-GSE
Time Period	2001-2003	2004-2006	2007-2009	2010-2012	2001-2003
Lagged Delinquency	0.015	0.090	1.34	0.848	32.32
	(1.40)	(0.079)	(0.221)	(0.186)	(7.39)
MSAs	374	374	374	374	43
N	2989	4482	4484	2982	224

Note: The table shows regression coefficients from a regression of residualized interest rates on lagged MSA-level default rates. Adjusted residual removes year*quarter fixed effects and quadratic controls for FICO and LTV. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and Loan Performance Database.

Table 3: Determinants and Dispersion of Default Rates, GSE Loans, 2001-2012

	All years	2001-2003	2004-2006	2007-2009	2010-2012
Regression with year*qtr fixed effects and interactions for LTV, LTV^2, FICO, and $FICO^2$					
R^2	0.0547	0.0396	0.0388	0.0804	0.011
N	9557438	3319226	2101734	1713411	866192
Regression with year*qtr fixed effects					
R^2	0.0081	0.0004	0.0008	0.0134	0.0008
N	9614660	3346519	2108036	1715509	866696
Mean Default Rate	0.024	0.022	0.022	0.041	0.004
SD across MSAs and quarters	0.026	0.011	0.023	0.041	0.005
SD across MSAs adjusted*	0.020	0.009	0.022	0.033	0.004

Note: The table shows r-squared and sample sizes for regressions exploring the relationship between default rates, time dummies (quarterly), LTV, and FICO scores. Default is defined as ever being 60 or more days delinquent within the first 24 months after origination. The lower panel of the table shows how much variation exists across MSAs, both after adjusting for time dummies and not. Adjusted residual removes year*quarter fixed effects and quadratic controls for FICO and LTV. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Table 4: Relationship between Interest Rates and Predicted Default, GSE and Prime Jumbo Loans

Dependent variable = Interest Rate Residual		
	GSE	Non-GSE
Time Period	2001-2003	2001-2003
Predicted Default	0.026	9.08
	(2.02)	(2.08)
MSAs	373	43
N	2961	224

Note: The table shows regression coefficients from a regression of residualized interest rates on predicted default rates. Adjusted residual removes year*quarter fixed effects and quadratic controls for FICO and LTV. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and Loan Performance Database.

Table 5: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates

Regional Lagged Default Rate	+2 SD	+1 SD	0	-1 SD	-2 SD
One time % change in C for indifference ($\lambda \times 100$)	1.12%	0.56%	-0.06%	-0.70%	-1.26%
Dollars per household for indifference $\left(\lambda \times 63790 \times \frac{\overline{c}_{region}}{\overline{c}_{overall}}\right)$	\$677	\$347	-\$38	-\$460	-\$854

Table 6: Redistribution by Age

Regional Lagged Default Rate	+2 SD	+1 SD	0	-1 SD	-2 SD
Overall: Consumption Equivalent ($\lambda \times 100$)	1.12%	0.56%	-0.06%	-0.70%	-1.26%
Young: Consumption Equivalent ($\lambda \times 100$)	0.54%	0.26%	0%	-0.36%	-0.76%
Middle Aged: Consumption Equivalent ($\lambda \times 100$)	1.46%	0.74%	-0.10%	-0.90%	-1.56%

Table 7: Redistribution by Age (No Rental Market)

Regional Lagged Default Rate	+2 SD	+1 SD	0	-1 SD	-2 SD
Overall: Consumption Equivalent ($\lambda \times 100$)	3.80%%	2.02%	0.20%	-1.50%	-2.96%
Young: Consumption Equivalent ($\lambda \times 100$)	3.96%	2.26%	0.40%	-1.34%	-2.88%
Middle Aged: Consumption Equivalent ($\lambda \times 100$)	3.68%	1.86%	0.06%	-1.60%	-2.98%

Table 8: Sensitivity to different values of ϕ_r

Regional Lagged Default Rate	+2 SD	+1 SD	0	-1 SD	-2 SD
Benchmark: Consumption Equivalent	1.12%	0.56%	-0.06%	-0.70%	-1.26%
1/2 Elasticity (Roughly FICO estimate)	0.62%	0.30%	0%	-0.32%	-0.62%
1/4 Elasticity	0.34%	0.18%	0.02%	-0.14%	-0.28%

Table 9: Other Elasticity Robustness

Regional Lagged Default Rate	+2 SD	+1 SD	0	-1 SD	-2 SD
Benchmark: Consumption Equivalent	1.12%	0.56%	-0.06%	-0.70%	-1.26%
No Regional Income Variation	1.20%	0.58%	-0.06%	-0.66%	-1.14%
No Regional House Price Variation	0.54%	0.10%	-0.54%	-1.34%	-2.20%

10 References

Diaz and Luengo Prado (2010)
Finkelstein and Poterba (2013)
Floden and Linde (2001)
Guvenen and Smith (2013)
Kaplan and Violante

11 Appendix

How robust are things to different elasticities? Reasonably so, discuss Table 8. Basically linear in elasticity.

Robustness to other things:

Direct comparisons across rows of Table 9 is somewhat tricky because the variation in consumption across states varies across rows. For example, in the benchmark economy in the worst state, average consumption is .888 and in the best state is .996. In the economy with no regional income variation, in the worst state average consumption is .938 and in the best state, average consumption is .944. Since there is diminishing marginal utility of consumption, it takes a smaller increase in consumption to increase utility by the same amount when consumption is low. This explains why the consumption equivalent in the benchmark economy is larger than the no regional income variation economy in the good state but smaller in the bad state.

Thus, the main takeaway from Table 9 should not be the direct comparison of the different economies but rather the fact that the large transfers across regions are quite robust to the assumptions about exactly what is correlated with regional income shocks.

When we turn off regional house price variation, there is a smaller effect on the average level of consumption so comparing rows is more relevant. In this case, variable house prices amplify the cost of variable mortgage rates in the bad state (since variable house prices mean low value of collateral in the bad state) but also reduces benefit of low r in good state (since it means that there is a low interest rate but high house prices when households want to buy).

This interaction is quite powerful in the good state. In the benchmark economy, households in the good state are willing to pay 1.26% of consumption to move to a variable interest rate while in an economy with constant house prices they are willing to pay 2.20% of consumption. This is because when house prices are constant they can really take advantage of low interest rates and buy bigger houses. In contrast, in the benchmark economy interest rates are low but house prices are high, which dampens the interest rate effect.

12 Computational Appendix

Discuss that no simultaneous positive b and m so reduce to one state variable. Then discuss renormalization to voluntary equity to make state space rectangular. Then discuss actual solution, grid sizes etc.