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

April 2014

PRELIMINARY - PLEASE DO NOT CITE OR QUOTE WITHOUT AUTHORS' PERMISSION

Abstract

Government sponsored mortgage enterprises (GSEs) securitize the bulk of all mortgages originated within the United States. We formally explore one of the aspects of the GSE pricing decision by studying the extent to which GSE policies result in the redistribution of resources across U.S. regions in a state contingent manner. We empirically establish that, despite large spatial variation in predictable default risk, there is essentially no spatial variation in GSE mortgage rates, conditional on borrower observables. While there is no regional risk-based pricing in the government-backed GSE market, the private market does set interest rates based in part on regional risk factors. 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. We attempt to quantify the economic impact of the GSEs' constant interest rate policy on regional risk sharing by building a structural spatial model of collateralized borrowing where households face both idiosyncratic and region-specific shocks. The model, parameterized and calibrated based on our analysis of GSE and private markets, suggests that the GSE national interest rate policy has significant ex-post redistribution consequences across regions: on an ex-post basis, it is comparable in size to fiscal stimulus packages such as tax rebates and payroll tax holidays. Although there are a range of consequences to housing and mortgage markets that are often attributed to the presence of GSEs, our paper suggests that their common national interest rate policy may be one important and understudied dimension of their impact on household choices.

*We thank seminar participants at the Bank of Italy, Kellogg, and Wharton for helpful comments and suggestions. Any remaining errors are our own.

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

I Introduction

Government sponsored mortgage enterprises (GSEs) securitize the bulk of all mortgages currently originated within the United States. In 2009, for example, the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac) owned or guaranteed more than 70 percent of new mortgages (CBO 2010).[1] Critics have long argued for dismantling Fannie Mae and Freddie Mac, and their push has intensified since the GSEs were placed into conservatorship by the U.S. government in 2008. Proponents, on the other hand, argue that GSEs serve parts of market that would not be served by private investors. In order to inform the public debate, it is necessary to quantify the costs and benefits of the GSEs on economic activity. In this paper, we take a first step towards this end and formally explore one of the aspects of the GSE pricing decisions. We do so by studying the extent to which GSE policies result in the redistribution of resources across U.S. regions in a state contingent manner.

To achieve our goal, the paper unfolds in three parts. In the first part of the paper we empirically establish that there is essentially no spatial variation in GSE mortgage rates, conditional on borrower observables. This holds across the recent housing boom and bust. As a corollary, we show that local economic conditions are unrelated to the mortgage rates on loans securitized by the GSEs. Specifically, using loan level data, we show that average interest rates in an MSA are uncorrelated with recent local house price growth, recent 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 large 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 higher rates to borrowers in regions with declining economic conditions.

We then provide some assessment of the extent to which GSE interest rates should vary spatially over time, given the large spatial variation in risk. In particular, we show that mortgage rates in the private 'prime jumbo' market, where loans are slightly larger than those made by the GSEs but comparable on many other dimensions, are strongly positively related to ex-ante predicted default probabilities across geography. Thus, while there is no regional risk-based pricing in the government-backed GSE market, 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

[1] Including loans guaranteed by the Federal Housing Administration (FHA) increases the share to over 90 percent (CBO 2010).

rates do not vary spatially over time, and conclude that political pressure is the most reasonable explanation for the patterns we observe. The GSEs face a great deal of political scrutiny. Evidence from prior efforts of the GSEs to regionally differentiate lending standards, 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 despite the presence of regional drivers of risk (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 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 markets, we find that the GSEs' national interest rate policy has some insurance value. We estimate that households would pay \$50 on average to maintain this constant interest rate policy, since it allows households to borrow more cheaply in the event of negative regional shocks. The cumulative value of this insurance is 0.06% of aggregate consumption, or roughly \$6 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 \$600 across households. This amount is comparable in size to fiscal stimulus packages such as tax rebates and payroll tax holidays.

Overall, our results suggest that the magnitudes of redistribution we observe through the mortgage market operated by the GSEs 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, our paper suggests that their common national interest rate policy may be one important and understudied dimension of their impact on household choices.

II Background

Most mortgages in the United States are sold to a secondary market after origination, rather than staying on lenders' balance sheets. For example, from 2004-2006, about 80 percent of all mortgages were securitized (Keys et al. 2013). Loans meeting the standards laid out by Fannie Mae and Freddie Mac are considered "conventional," and thus eligible for purchase by these Government-Sponsored Enterprises (GSEs).[2] These loans are purchased, packaged, and insured against loss of principal and interest in the resulting mortgage-backed securities. As a premium, investors pay a "guarantee fee" on each loan. The guarantee fee varies by some of features of the borrower and loan product, most notably the FICO credit score of the borrower and the loan-to-value (LTV) ratio of the mortgage. The "coupon" rate is the rate received by investors, and the difference between the interest rate paid by borrowers and the coupon rate is the guarantee fee, which also covers the administrative costs of Fannie and Freddie.

The alternative secondary market for mortgages is known as the non-agency or private mortgage-backed security (MBS) market. In this market, loans that do not meet the standards of the GSEs are purchased, bundled, and sold to investors in the form of securities. These investors do not receive any guarantees against losses of principal and interest of the loans underlying the securities. Thus, while investors in GSE securities are insulated from default risk, investors in the private market must accurately price both the risk of default and the risk of early prepayment. The private market began in the 1980s with loans that did not meet the conforming loan limit (i.e., "jumbo" loans), but expanded into "subprime"and "Alt-A" loans thereafter, and grew to a large share of the market during the housing boom.[3]

Regardless of whether the loan is sold and to which market it is sold, mortgage lenders must follow Regulation B, which implements the Equal Credit Opportunity Act (ECOA) of 1974.[4] The ECOA prohibits either "disparate treatment" or "disparate impact" of borrowers across any characteristic covered by the act, including age, race, sex, marital status, national origin, religion, or source of income. However, the ECOA does not prohibit differential pricing of risk, and lenders are able to price based on creditworthiness and the size of the downpayment, as well as other aspects of

[2]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. See Green and Wachter (2005) for additional details.

[3]"Subprime" is a term generally used to refer to loans made to borrowers with less than stellar credit histories, whereas "Alt-A" loans are usually loans made to borrowers who have strong credit records, but are less likely to provide full documentation of income and assets. In the MBS market, these classifications were made for investor purposes and often given at the pool level.

[4]For additional details, see http://files.consumerfinance.gov/f/201306_cfpb_laws-and-regulations_ecoa-combined-june-2013.pdf

the contract such as whether the rate is fixed or adjustable, whether the borrower provided full documentation of income and assets, and whether the borrower intends to use the property as his primary residence.

Over most of the period prior to the housing boom, 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. The non-agency category includes jumbo mortgages (those that exceed the conventional mortgage size limits), sub-prime mortgages, and Alt-A mortgages. 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.

The impact of government involvement in the mortgage market has been discussed in the context of stable availability of the 30-year fixed-rate mortgage (Fuster and Vickery 2013) and in the differential pricing around the jumbo-conforming loan limit (Sherlund 2008). Fannie Mae and Freddie Mac were publicly chartered but privately operated from 1968 until they were taken into conservatorship by the federal government in 2008 during the crisis. The reasons for their failure remain the subject of some debate, but their excessive risk-taking and leverage, lack of diversified assets, and implicit government support that kept the cost of lending down have been pointed to as major determinants of their demise (Acharya et al. 2011). At the time of this writing, Fannie and Freddie are highly profitable and have paid back nearly all of the funds needed for their 2008 bailout. Proposals to reform Fannie and Freddie have focused on the tradeoff between maintaining a countercyclical source of liquidity for mortgage markets with the risk to taxpayers and lack of innovation and regulatory capture. These discussions have not noted, to the best of our knowledge, the regional redistribution that Fannie and Freddie provide through a constant interest rate policy as a un-priced benefit of the current system.

III Data and Empirical Specification

We use two main data sources for our empirical work in this paper. The first includes a sample of loans securitized by either Fannie Mae or Freddie Mac. The second includes a sample of jumbo loans securitized by the private market.

III.A Fannie Mae/Freddie Mac Sample

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.[5] 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 of month's delinquent.[6] There is a unique loan identifier code in the data sets that allows a loan to be tracked from inception through its future performance.

When creating our analysis file, we pool together data from both the Fannie Mae and Freddie Mac datasets. In doing so, we are exploring the spatial variation in interest rates for all conventional loans that are securitized by the GSEs. As we show in the Online Robustness Appendix that accompanies the paper, all of our key results are identical in all respects if we use only the Fannie Mae data or the Freddie Mac data. Finally, within our analysis sample, we include loans associated with only new housing purchases.[7] In total, our sample includes roughly 7 million loans in our combined analysis sample.

[5] 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.

[6] 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.

[7] The results are unchanged if we included refinanced loans as well. The data appendix discusses other sample restrictions as well. In particular, we only include mortgages that were originated between January 2001 and December of 2009 and only include mortgages that were originated within one of our included MSAs.

III.B Prime Jumbo Sample

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, 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 "prime" jumbo set of mortgages accordingly: (1) the origination value is between the conforming mortgage limit and two times the conforming mortgage limit, (2) have a fixed interest rate, (3) have a LTV at origination of less than 100 percent, (4) have a FICO score at origination greater than 620, and (5) provide full documentation at the time of origination.[8] As shown by Keys et al. (2012), a FICO score of 620 is a cutoff for the loan to be purchased by the GSEs. The reason we cap the mortgage origination value at twice the conforming limit is we want to create a sample that is similar to the Fannie/Freddie sample. Very expensive loans may be differ along many dimensions of loan and borrower characteristics compared to loans in the Fannie/Freddie sample. In essence, our prime jumbo loans are designed to be similar to the Fannie/Freddie loans in all respects except the origination value of the loan is slightly higher. As with the Fannie/Freddie mortgages, we only include originations for new purchases.

III.C Time Periods

Aggregate business cycle conditions differed markedly across sub-periods of the 2000s. In the early 2000s, the aggregate economy was in a mild recession. During the mid-2000s, house prices were booming and aggregate default rates on conventional mortgages were low. In the latter 2000s, house prices fell, mortgage default rates were high and the economy was in a severe recession. Given the changing business cycle conditions, we analyze the spatial patterns in interest rates and default during three distinct time periods. The first time period focuses on the years around the 2001 recession and includes loans originated between 2001 and 2003. The second time period focuses on the subsequent period of rapid house price growth between 2004 and 2006. The final period focuses on loans originated between 2007 and 2009 during the Great Recession.[9]

[8]The conforming limit was raised from $275,000 to $417,000 between 2001 and 2006. This period pre-dates the policy to vary loan limits regionally based on "high cost" areas, which began in 2008.

[9]Our analysis stops in 2009 because we want to follow loans for a full two years after origination. We discuss the rationale for this restriction below.

Given the structure of our data, we examine the spatial patterns for both the conventional mortgages and the prime jumbo mortgages during the 2001-2003 period and the 2004-2006 period. After 2006, the private market essentially stopped securitizing loans. As a result, it is impossible to compare the spatial patterns in interest rates and defaults between loans securitized by the GSEs and loans securitized by the private market during this time period. However, we still have data on loans securitized by the GSEs after 2006. For this reason, we only explore the spatial patterns in interest rates during the 2007-2009 period for these loans.

Appendix Table A1 provides descriptive statistics for the pooled Fannie/Freddie loans and the prime jumbo loans during each of the time periods. Given the restrictions we have placed on mortgage types to maintain comparability with GSE originations, the prime jumbo sample is substantially smaller than the sample of Fannie and Freddie loans. In both samples, the median LTV ratio is 80%. However, the Fannie/Freddie loans are initiated with less equity in the home. For example, during the 2004-2006 period, roughly one-quarter of all loans in the sample were originated with an LTV greater than 0.8. The comparable number in the prime jumbo sample was 14 percent. Conversely, the loans originated by the GSEs have relatively higher credit scores compared to this special subgroup of the private market.

III.D Controlling for Borrower and Loan Characteristics

Throughout the paper, we want to examine the spatial variation in mortgage rates and show how the variation correlates with the spatial variation in predicted future mortgage default rates. One reason interest rates and default/delinquency rates could differ spatially is because borrower and loan characteristics could 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 the predictable component of default rates vary spatially conditioning on borrower/loan characteristics. A borrower with a given credit score and LTV ratio may be more likely to default in one area relative to another because overall economic conditions differ across the areas.

To formally illustrate these patterns, we purge the spatial variation in actual mortgage rates paid of spatial differences in borrower and loan characteristics.[10] To do that, we first estimate the

[10] We discuss how we purge default rates of borrower and loan characteristics in the following section.

following equation using our loan level micro data:

$$r^j_{ist} = \alpha^j_1 + \alpha^j_1 X_{it} + \Gamma^j_1 D_t + \eta^j_{ist}$$

where r^j_{ist} is the loan-level mortgage rate for a loan made to borrower i, from sample j, 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). Sample j refers to whether we use individuals from the GSE sample or the private jumbo sample. We run these regressions separately using data from each of our two samples. D_t is a vector of time dummies based on the quarter*year of origination. The borrower/loan controls include detailed FICO and LTV controls. Specifically, all regressions include dummies for whether the loan was originated with an LTV within a 5 unit range (70-74, 75-79, 80-84, etc.) and whether the borrower had a FICO score within a 20 unit bin (620-639, 640-659, etc.). We also allow the FICO score controls to be interacted with the time dummies. The goal of these specifications is to recover η^j_{ist} . η^j_{ist} is the residual mortgage rate for borrower i in MSA s during time t for loans in sample j after controlling for borrower/loan characteristics and time fixed effects. For ease of interpretation in some of our results below, we add back in the average mortgage rate for the relevant sample/time period.[11]

One natural question is how much of the variation in mortgage rates within each sample can be explained by time effects and borrower characteristics. To answer this, we run the above regression using data from each sample/time period, both including and excluding the vector of X controls. For the Fannie/Freddie data, the R-squared from the above regressions excluding the X vector of controls were 0.71, 0.46, and 0.68 in the 2001-2003 period, the 2004-2006 period, and the 2007-2009 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 0.73, 0.50, and 0.72 in the 2001-2003, 2004-2006, and 2007-2009 periods, respectively. As expected, higher FICO scores and lower LTV are associated with statistically significant and economically meaningfully lower mortgage rates. For example, within the 2001-2003 sample, those borrowers with a FICO score of 620 paid a 0.25 percentage point higher mortgage rate, respectively, than those borrowers with a FICO score of 760, all else equal.

The results are broadly similar for the sample of prime jumbo loans securitized by the private

[11] This just affects the scaling of the residuals. To calibrate one of our key parameters, we run log of the interest rate residuals on log of predicted default. To do this, we need all of the residuals to be non-negative and centered around the sample mean.

market. For these loans, the R-squared from the above regressions excluding the X vector of controls were 0.45 and 0.37 during the 2001-2003 period and the 2004-2006 period. Including the X vector of controls raises the R-squared to 0.55 and 0.43, respectively. In other words, the FICO and LTV controls explain more of the variation in the prime jumbo sample than in the Fannie/Freddie sample. Like with the Fannie/Freddie sample, borrowers with a low LTV or a high FICO score paid lower mortgage rates, albeit with a far steeper relationship. For example, within the 2001-2003 sample, those borrowers with a FICO score of 620 paid a 1.25 percentage point higher mortgage rate relative to those borrowers with a FICO score of 760, all else equal.

Once we have the residuals from the above regressions with the full set of controls, we compute location specific average mortgage rates , R^j_{st}. We do this separately for each time period and for each sample. Specifically,

$$R^j_{st} = \frac{1}{N^j_{st}} \sum_{i=1}^{N^j_{st}} \eta^j_{ist}$$

where N^j_{st} is the number of loans in the MSA s during quarter*year t within each sample. Formally, R^j_{st} will be the average mortgage rate residual in an MSA during a given period for a given sample.

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^j_{st}'s) across the MSAs for our three time periods of study and for each sample. Note that there is no information on the prime jumbo loans in the 2007-2009 period because, as discussed above, this market effectively ceased operations in mid-2007.

A few things are of note from Table 1. First, the unconditional cross-MSA standard deviation in interest rates is high for both the prime jumbo and Fannie/Freddie loans. For example, during the 2001-2003 period there was, on average, a 58 basis point and 96 basis point standard deviation in unconditional interest rates across MSAs in the GSE and prime jumbo samples, respectively. Second, the unconditional spatial variation in interest rates was much higher during the 2001-2003 period than the 2004-2006 period for loans in both samples. This is not surprising given that the spatial variation in economic conditions was much higher during the 2001-2003 period than the 2004-2006 period. We formally illustrate this fact below. Third, the conditional cross-MSA standard deviation is close to zero for the Fannie/Freddie sample in both time periods. In other words, the unconditional cross-sectional dispersion in mortgage rates across MSAs is driven by cross-sectional variation in LTV and FICO scores. Conditional on LTV and FICO scores, there is essentially no cross-sectional dispersion in mortgage rates. Finally, even after controlling for LTV and FICO score, there is still substantial cross-sectional variation in mortgage rates in the prime-jumbo sample. For example, there was a 33 basis point standard deviation in conditional mortgage rates across MSAs during the 2001-2003 period.

IV Spatial Variation in Mortgage Rates

In this section, we continue to explore the spatial variation in mortgage rates within our two samples. In particular, we document (1) the extent to which mortgage rates vary systematically across regions and (2) whether the spatial variation in mortgage rates is related to variation in local economic conditions.

IV.A 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, within each MSA s 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 $t-1$).[12] 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 (e.g., weak local economic activity, declining house prices) 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 $t-1$).[13] Appendix Table A2 shows the mean and standard deviation of each of our three local economic condition measures separately along with the correlation between the measures for each of our time periods.

To describe the data, we run the following specification:

$$R^j_{st} = \omega^j_0 + \omega^j_1 E_{s,t-1} + \nu^j_{st}$$

[12] 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.

[13] 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.

where R_{st}^j is the average residualized MSA mortgage rate (defined above) and $E_{s,t-1}$ is lagged MSA measure of economic activity. ω_1^j represents the responsiveness of current average mortgage rates within an MSA to lagged local economic conditions. The three panels of Figure 1 plot the relationship between the residualized MSA mortgage rate and the lagged Fannie/Freddie mortgage delinquency rate estimated separately for each of our three time periods focusing only on our sample of GSE loans. Each observation in the figure is an MSA*quarter observation. 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 1 illustrates one of the key empirical results within our paper: there is no relationship between mortgage rates and local economic activity for loans securitized by Fannie/Freddie conditional on borrower and loan observables. For example, during the 2001-2003 period (panel A) the scatter plot shows that a 1 percentage point increase in lagged default rates *reduces* mortgage rates by 0.21 basis points. The standard deviation of lagged default rates during this period is 0.4 percentage points. So, a one standard deviation change in lagged default rates across MSAs lowers the MSA interest rate by 0.0008 percentage points. This is both statistically and economically zero. The patterns in the 2001-2003 period are not anomalous. During the 2004-2006 period and the 2007-2009 period, a one standard deviation increase in lagged default increased MSA mortgage rates by 0.0004 percentage points and 0.009 percentage points.[14] In all time periods, there is essentially no economically meaningful relationship between lagged local default rates and local mortgage rates.

The relationship between local interest rates and local economic activity differs in the sample of prime jumbo loans. These results are illustrated in Figure 2. The figure shows how the residualized jumbo mortgages rates in an MSPA respond to lagged local default rates during the 2001-2003 (Panel A) and 2004-2006 (Panel B) periods. The lagged default rate is computed using the sample of Fannie/Freddie loans. Given that, the variation in the x-axis across MSAs for panels A and B of Figure 2 are identical to the variation in the x-axis for panels A and B of Figure 1. Panel A of Figure 2 shows that a 1 percentage point increase in lagged default increases the jumbo mortgage rate by 0.31 percentage points (p-value < 0.01) during the early 2000s, all else equal. The estimate of ω_1^j is 147 times larger for the prime jumbo sample during the 2001-2003 than for the Fannie/Freddie sample. Panel B of Figure 2 shows that a 1 percentage point increase in lagged default increases jumbo mortgage rates by 26 basis points during the 2004-2006 period.

Table 2 reports the coefficients from the above regression where again each observation is weighted

[14] A 0.009 percentage point increase in mortgage rates implies that a mortgage rate of 5.0 percent would rise to 5.009 percent. Equivalently, a 0.009 percentage point increase is a 0.9 basis point increase.

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 three time periods, for each of our three local economic condition measures, and for each of our samples, we estimate a separate regression. The top panel focuses on the GSE sample while the bottom panel focuses on the prime jumbo sample. Within each panel, results are shown using lagged GSE default as the measure of economic activity (top rows), lagged house price growth as the measure of economic activity (middle rows), and lagged unemployment changes as the measure of economic activity (bottom rows). For each regression, we provide estimates of ω_1^j, the standard error of ω_1^j (in parentheses), and the marginal effect of a one standard deviation change in the relevant local economic condition measure (in brackets).

Table 2 documents that lagged default is a good summary statistic for local economic activity at the time of origination in that prime jumbo mortgage rates responded to it in both the 2001-2003 period and 2004-2006 period. Local mortgage rates in the prime jumbo loans also respond to differences in the lagged unemployment rate. However, the interest rate response is much smaller than with the lagged local default rate. Given this, going forward we focus our analysis on the lagged local default rate as our measure of lagged local economic activity.[15]

To summarize, we document that for loans securitized by Fannie/Freddie (1) there is little spatial variation in mortgage rates across U.S. MSAs and (2) the variation that exists is unrelated to local economic conditions. These patterns stand in stark contrast to the patterns for the prime jumbo loans. These loans, we argue, are a good comparison for the Fannie/Freddie loans. In the following section, we will show that lagged economic activity is a strong predictor of future default for *both* the GSE loans and the prime jumbo loans. This finding makes the lack of interest rate variation across regions for the Fannie/Freddie loans even more stark.

[15] In the Online Robustness Appendix that accompanies the paper, we confirm our GSE findings using two alternate data sets. 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 mortgage offer quotes for a specific loan type from a specific lender. 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. Both loan types focus on 30 year fixed rate mortgages with no points with a loan to value ratio of at least 80 percent and a borrower's FICO score of 750. The loans we have access to differ by origination value. Using this data, we confirm the patterns in the GSE data used in our primary analysis. During the period covered by the loan sifter data, we find no relationship at all between local economic conditions and local mortgage rates being offered by lenders. These results provide confidence that our results are not biased because we do not have access to points/fees in the GSE data.

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

In the prior sections, we used argued that the prime jumbo loans are a good comparison set for the Fannie/Freddie loans with respect to exploring the spatially variation in interest rates. However, the prime jumbo loans may default more in response to variation in local economic conditions than the Fannie/Freddie loans in response to that same variation in local economic conditions. In that case, it is not surprising that the interest rates vary more in the prime jumbo sample. If that is true, we will need to adjust the interest rate variation for differences in the predicted variation in default.

To start to explore these relationships, we estimate:

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

where $E_{s,t-1}^j$ is defined as a above and Y_{st}^j is the residualized ex-post default rate for loans originated in period t in MSA s from sample j over the subsequent two years. We residualize the ex-post default rates to control for the fact that some types of borrowers and some types of loans are more likely to default than others.[16] Our measure of the ex-post default rate is the 60 day delinquency rate. Given that, our ex-post actual default measure is what fraction of the loans originated in period t ever became 60 days delinquent during the subsequent two years. Specifically, the above regression just assesses for each sample the extent to which the lagged Fannie/Freddie default rate in the MSA (our preferred measure of $E_{s,t-1}^j$) has predictive power for future default rates for loans issued in that period with a given MSA. We run this regression only for the 2001-2003 and 2004-2006 time period which is a period where data exists for both samples.

After running the above regression, we form a predicted measure of local residual defaults for both samples such that:

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

The predicted measure is made MSA by MSA. Our key specification then runs local residual

[16] Formally, we obtain Y_{st}^j by first running a regression of ex-post default of a given loan on FICO, LTV, Time, and FICO*Time controls using the micro date of loans. We get the residuals from this regression and then add back in the average ex-post default rate during the time period (so the residuals are centered around the average default rate). We then compute the average of residuals (centered around the average default rate) for each MSA, for each time period, and for each sample. This procedure is exactly the same as the procedure discussed above to compute the residualized mortgage rates.

mortgage rates on predicted measures of local residual default:

$$\ln(R^j_{st}) = \beta^j_0 + \beta^j_1 \ln(\hat{Y}^j_{st}) + \epsilon^j_{s,t}$$

Table 3 shows the estimates of ψ^j_1 and β^j_1 for both the Fannie/Freddie sample and the prime jumbo sample for both the 2001-2003 time period and the 2004-2006 time period. The difference between β_1's for the Fannie/Freddie sample and the prime jumbo sample tell us the differential local mortgage rate response between the two market to local differences in predicted future default rates. All of the coefficients in the table are from log-log specifications and thus can be interpreted as elasticities in a straightforward manner.

The first row of Table 3, Panel A, shows the estimates of psi_1 for the GSE market. There we see that recent defaults in an MSA are strongly predictive of subsequent defaults, with a one percent increase in recent defaults predicting a 0.45 percent increase in subsequent defaults during the 2001-2003 period. Similarly for the first row of Panel B, a one percent increase in recent GSE defaults in an MSA predicts a 0.86 increase in subsequent non-GSE defaults during this period. That is, the recent performance of loans backed by the GSEs in an MSA is an excellent predictor of subsequent loan performance in that MSA across both GSE and private markets.

Estimates of the impact of predicted defaults on local mortgage rates (β_1) are provided in the second row of Table 3 in both panels. The table shows that local mortgage rates are orders of magnitude more responsive in the prime jumbo segment of the private market than in the GSE market. While a one percent increase in predicted defaults in an MSA has effectively no impact on GSE mortgage rates (the coefficient is a precisely estimated zero), a similar increase raises mortgage rates by 0.046 percent. To determine whether this magnitude is economically meaningful, we turn to a model in the sections below, but first we discuss potential explanations for why mortgage rates in the GSE market do not vary with even predictable economic conditions.[17]

[17] As a robustness exercise, we performed a separate counterfactual exercise 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. 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 estimated from Table 3.

V.A Why Do Conventional Mortgage Rates Not Vary With Current Local Economic Conditions?

Why do 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.[18] 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 that had an initial LTV lower than 0.95. However, in declining markets, Fannie/Freddie would only purchase mortgages where the initial LTV was lower than 0.9.[19] 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.[20] 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."[21] 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."[22] 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

[18] 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.

[19] Fannie and 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.

[20] 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

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

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

on reasons why the GSEs do not raise interest rates in 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 (2014) 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.

VI 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 in the manner in which they do for the prime jumbo market.

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.

VI.A 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.

VI.B 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 Piazzesi 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.

VI.C Income Shocks

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

$$\log y_{ijk,t} = \chi_j + z_{i,t} + \phi^y \gamma_{k,t}$$

$$\log z_{i,t} = \rho_z \log z_{i,t-1} + \eta_{i,t}$$

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.

VI.D 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 $p_{k,t} r^f$.[23] We denote owner occupied houses as $h_{i,t}$ and rented houses as $h^f_{i,t}$.

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 Diaz and Luengo-Prado (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.

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

VI.E 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:

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

and a household that chooses to sell its current house[24] and rent has value function

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

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.

[24] 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.

VII 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.

VII.A 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 the 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.[25] 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 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 (2010).

[25] 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.

VII.B 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, we can then compute all necessary elasticities by computing the empirical elasticities with respect to lagged-default.[26] 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 regionally-based pricing? Using the estimates from Table 3, where we regressed log jumbo mortgage rates on log lagged-default in the early 2001-2003 period, and applying those estimates to the variation in GSE mortgages in the post-boom 2007-2010 period, implies that a two-standard deviation increase in lagged-default should increase mortgage rates by 5.9%. During this post-boom period, GSE mortgage rates averaged 6 percent, so we benchmark ϕ^r such that a two-standard deviation increase in default leads to a 36 basis-point increase in rates. This initial benchmark is intended to capture the positive relationship between defaults and interest rates as shown in the prime jumbo market (Figure 6). We provide robustness results for both larger and smaller ϕ^r and discuss alternative counterfactuals based on FICO scores in the following section.

VIII 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 lower rates and bad regions would have higher 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

[26] Note that $\log\gamma = -$lagged default as there should be a negative relationship.

give households today to make them indifferent between the variable interest rate and the constant interest rate policy.

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, conditional on living in a region with economic activity γ.[27] 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 4 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 4 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.[28] 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 re-distributional consequences across regions. We will show shortly that this conclusion is quite robust to a range of alternative model parameters. In contrast,

[27] 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.

[28] Total non-durable consumption in 2012 was $7335.9 billion = (2296.8 spending on nondurable 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.

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 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 0.06% so that the constant interest rate policy very mildly increases welfare for the average household in the economy.[29]

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 5 redoes the calculation in Table 4 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.[30]

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 6 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

[29] Note that this number is not the sum of the numbers in row 1 of Table 4 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.

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

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 worse off: they are forced to purchase when their income is low and interest rates are high.

In Tables 7 and 8, 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.

IX Conclusion

In 2008, when placed into conservatorship, the Federal National Mortgage Association (Fannie Mae) and the Federal Home Loan Mortgage Corporation (Freddie Mac) owned or guaranteed roughly half of the U.S.'s $12 trillion mortgage market. While there have always been proposals to scale back or 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 and quantify the costs and benefits of the GSEs' impact on local economic activity.

In this paper, we explore one of the understudied aspects of the GSEs' pricing decisions. We first empirically establish that there is essentially no spatial variation in GSE mortgage rates, conditional on borrower observables, in any recent period. The lack of spatial variation in mortgage pricing by the GSEs is surprising given the significant extent of spatial variation in predictable default risk.

We then ask the extent to which agency mortgage rates should 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.

Using a structural spatial model of collateralized borrowing where households face both idiosyncratic and region-specific shocks, we find that the GSEs' national interest rate policy has some insurance value, and households would pay $50 on average to maintain this constant interest rate policy. We 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 roughly $600 across households, an amount comparable in size to recent fiscal stimulus programs such as tax rebates and tax holidays.

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. By distributing resources across U.S. regions in a state contingent way, in addition to providing counter-cyclical liquidity to the mortgage market, Fannie Mae and Freddie Mac provide meaningful insurance during aggregate downturns. We hope to better understand the impact of this particular policy on housing market activity and house prices in future work.

X References

Acharya, Viral V., Matthew Richardson, Stijn Van Nieuwerburgh, and Lawrence J. White, "Guaranteed to Fail: Fannie Mae, Freddie Mac, and the Debacle of Mortgage Finance," Princeton, NJ: Princeton University Press, 2011.

Congressional Budget Office, "Fannie Mae, Freddie Mac, and the Federal Role in the Secondary Mortgage Market," December 2010.

Diaz, Antonia, and Maria Jose Luengo-Prado, "The Wealth Distribution with Durable Goods," International Economic Review, 51(1), February 2010.

Finkelstein, Amy, and James Poterba, "Testing for Asymmetric Information using "Unused Observables" in Insurance Markets: Evidence from the U.K. Annuity Market," Journal of Risk and Insurance, forthcoming 2014.

Floden, Martin and Jesper Linde, "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" Review of Economic Dynamics, 4(2), 2001.

Fuster, Andreas and James Vickery, "Securitization and the Fixed-Rate Mortgage," Federal Reserve Bank of New York Staff Report No. 594, January 2013.

Green, Richard K., and Susan Wachter, "The American Mortgage in Historical and International Context," Journal of Economic Perspectives, 19(4), 2005.

Guvenen, Fatih, and Anthony Smith, "Inferring Labor Income Risk and Partial Insurance from Economic Choices," Federal Reserve Bank of Minneapolis Research Department Staff Report No. 485, June 2013.

Kaplan, Greg and Giovanni L. Violante, "How Much Consumption Insurance Beyond Self-Insurance?" American Economic Journal: Macroeconomics, 2(4), October 2010.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," Quarterly Journal of Economics, 125(1), February 2010.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig, "Mortgage Financing in the Housing Boom and Bust," in E. Glaeser and T. Sinai, eds., "Housing and the Financial Crisis," Chicago, IL: University of Chicago Press, 2013.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, "Housing, Consumption, and Asset Pricing," Journal of Financial Economics, 83, March 2007.

Sherlund, Shane M., "The Jumbo-Conforming Spread: A Semiparametric Approach," Federal Reserve Board Finance and Economics Discussion Series 2008-01, January 2008.

Figure 1: Relationship between Interest Rates and Lagged Local Default, GSE Loans, 2001-2003 and 2004-2006

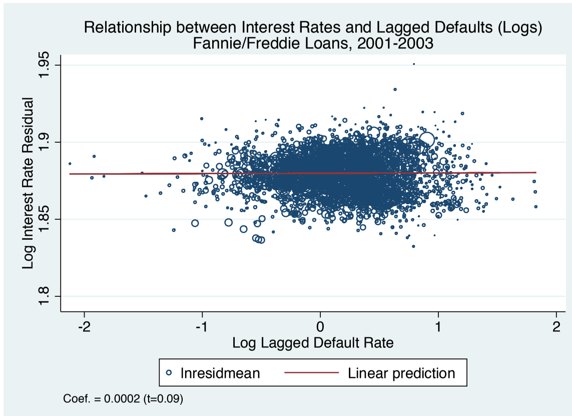

(a) 2001-2003

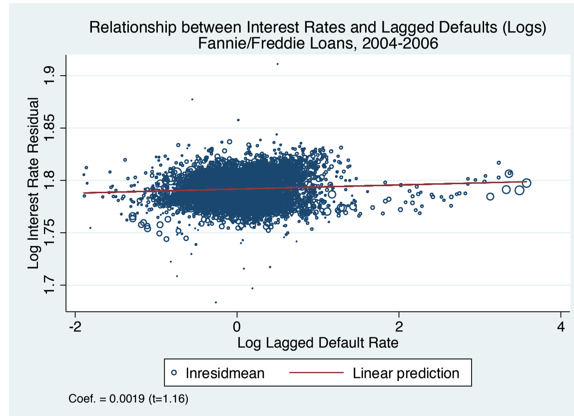

(b) 2004-2006

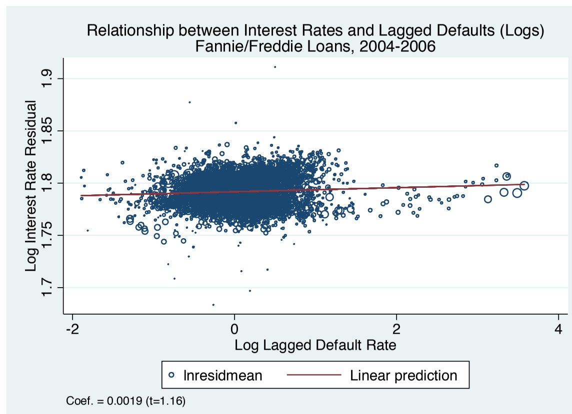

(c) 2007-2009

Note: Figure shows the relationship between residualized interest rates and lagged MSA-level default of loans originated within the last two years for three time periods. Adjusted residual removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 2: Relationship between Interest Rates and Lagged Local Default, Non-GSE "Prime Jumbo" Loans, 2001-2003

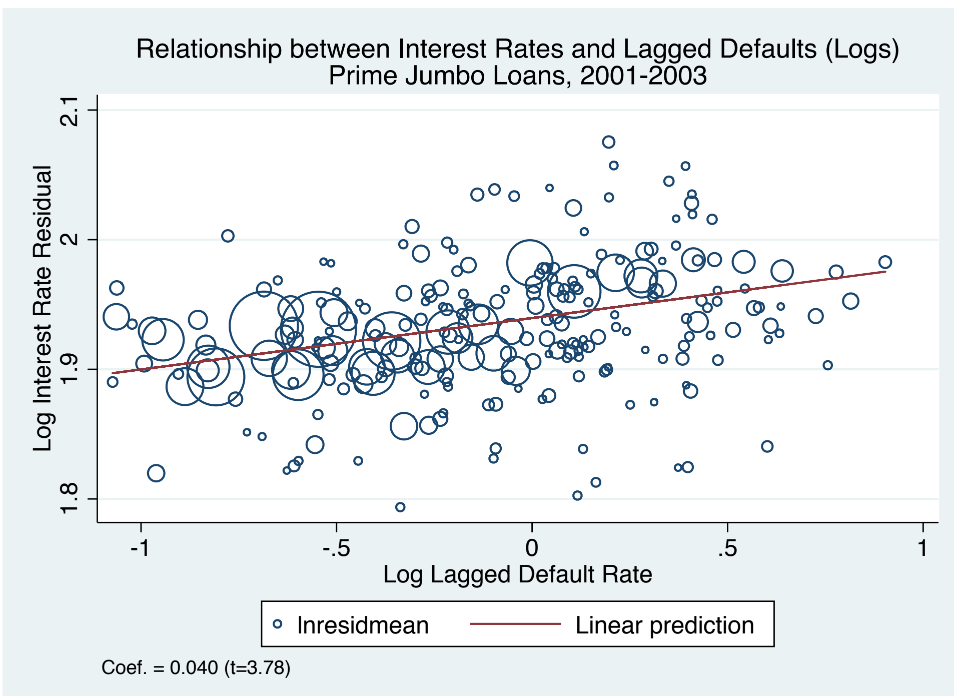

Note: Figure shows the relationship between Prime Jumbo residualized interest rates and lagged MSA-level default of loans originated within the last two years for 2001-2003. Source: Authors' calculations using Loan Performance Database.

Figure 3: Relationship between Default Rates and Lagged Local Default, GSE Loans, 2001-2003

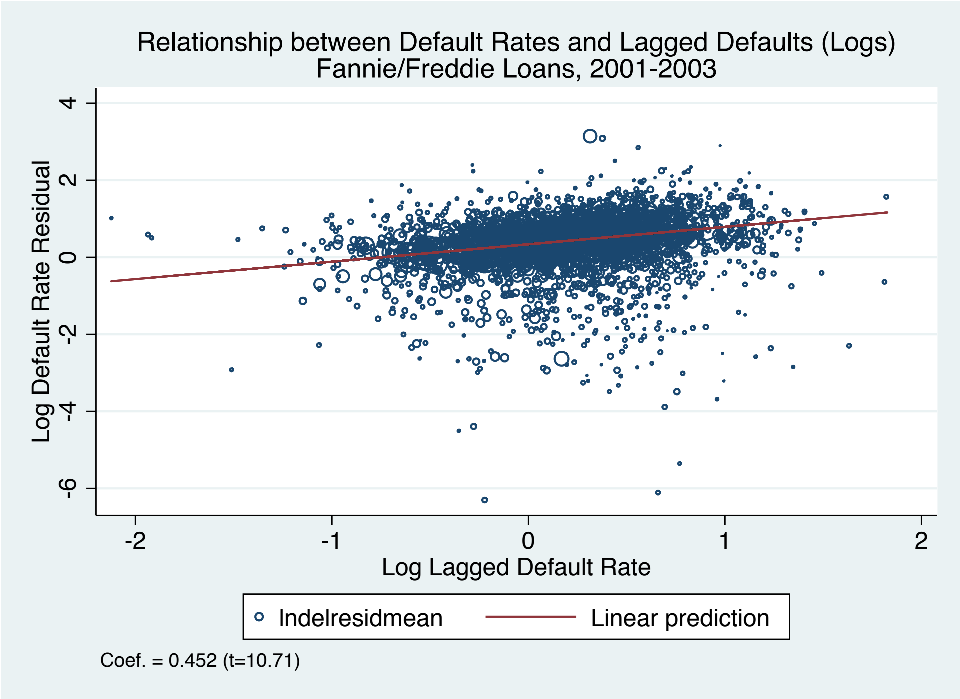

Note: Figure shows the relationship between residualized default rates and lagged MSA-level default of loans originated within the last two years for 2001-2003. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

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

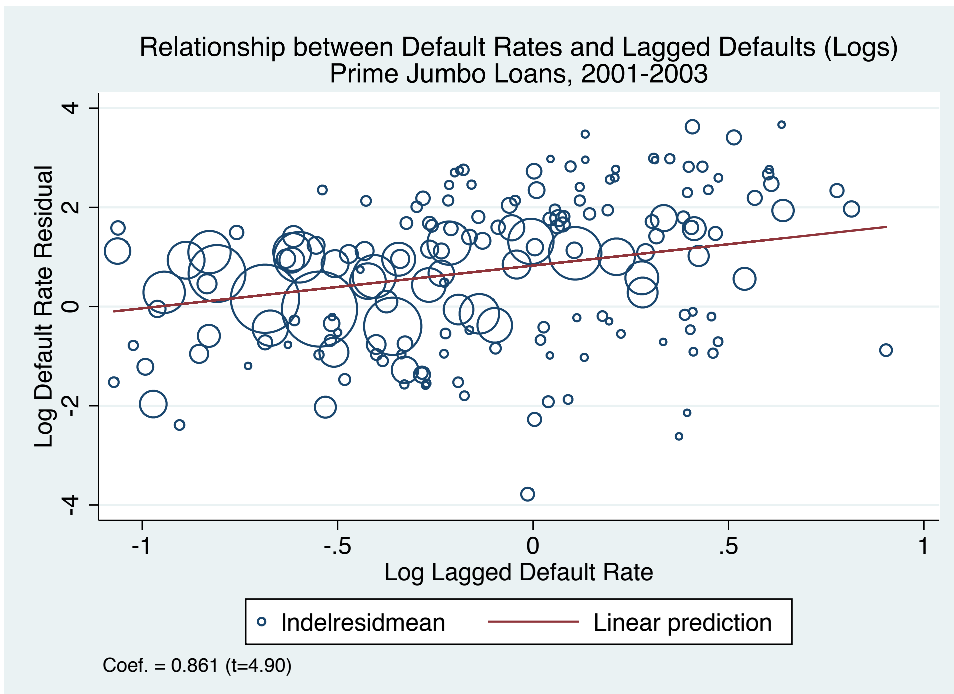

Note: Figure shows the relationship between residualized default rates and lagged MSA-level default of loans originated within the last two years for 2001-2003. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 5: Relationship between Interest Rates and Predicted Default Rates, GSE Loans, 2001-2003

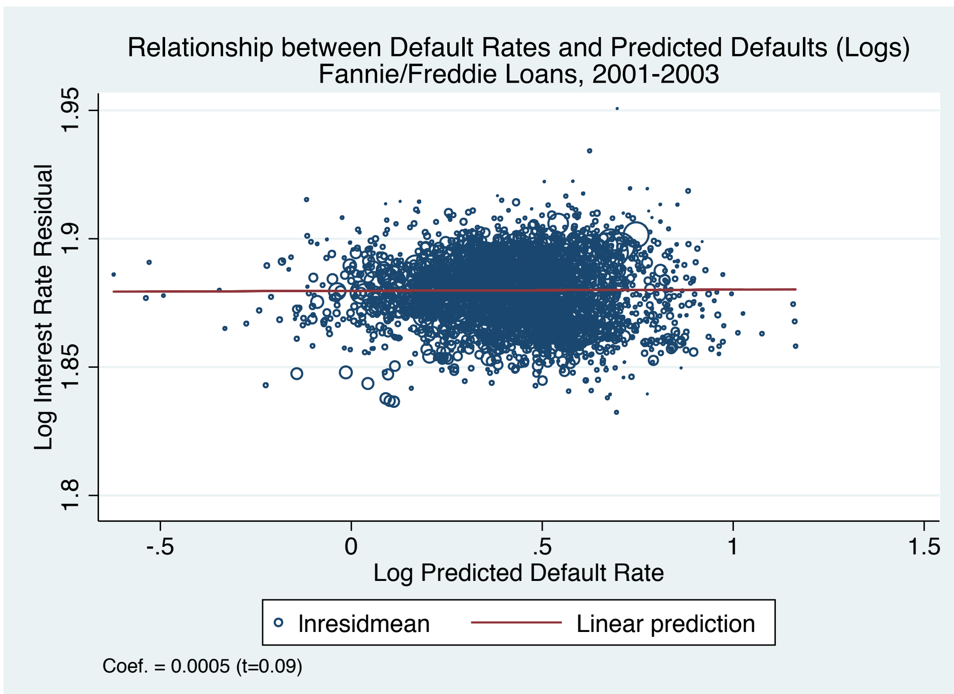

Note: Figure shows the relationship between residualized interest rates and predicted default rates for loans originated in GSE market, 2001-2003 period. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases.

Figure 6: Relationship between Interest Rates and Predicted Default Rates, Non-GSE "Prime Jumbo" Loans, 2001-2003

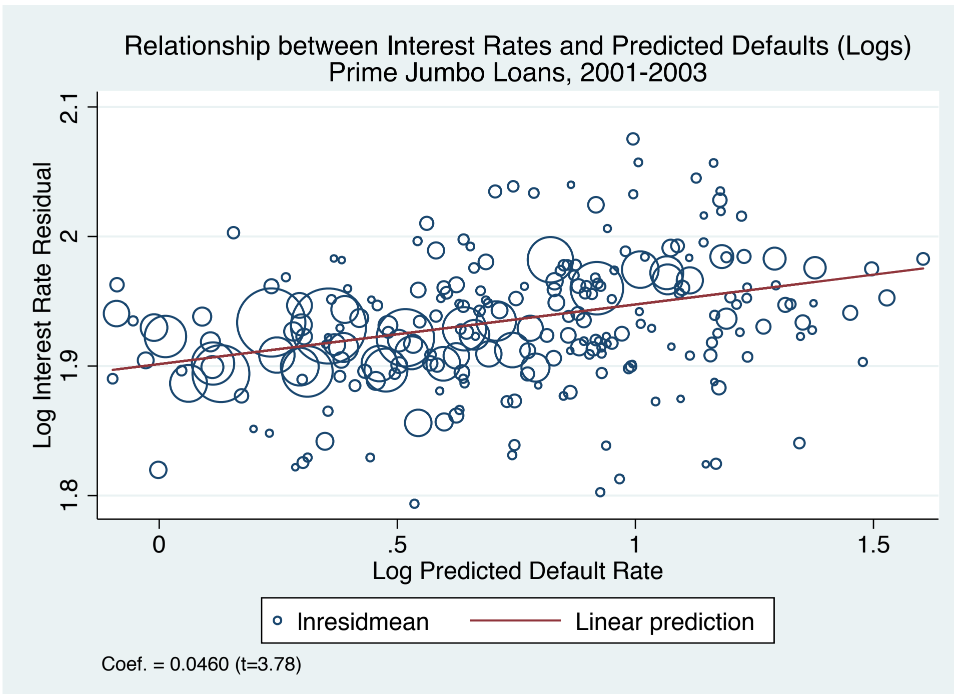

Note: Figure shows the relationship between Prime Jumbo residualized interest rates and predicted default rates for loans originated in non-GSE market, 2001-2003 period. Source: Authors' calculations using Loan Performance Database.

Table 1: Conditional and Unconditional Cross-MSA Dispersion in Mortgage Rates

	Time Period			
	All	2001-2003	2004-2006	2007-2009
<u>Fannie/Freddie Loans</u>				
Mean Interest Rate	6.24	6.56	6.01	5.90
Unconditional Cross-MSA Standard Deviation	0.589	0.578	0.300	0.556
Conditional Cross-MSA Standard Deviation	0.074	0.077	0.070	0.073
<u>Prime Jumbo Loans</u>				
Mean Interest Rate	6.89	6.89		
Unconditional Cross-MSA Standard Deviation	0.963	0.963		
Conditional Cross-MSA Standard Deviation	0.329	0.329		

Note: The table shows mean interest rates in the two samples across time periods, as well as cross-MSA standard deviations in interest rates, both unconditionally and conditionally. Conditional measure of standard deviation removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases (top panel) and LoanPerformance (bottom panel).

Table 2: Relationship between MSA Mortgage Rates and Lagged Economic Activity

	Time Period		
	2001-2003	2004-2006	2007-2009
Panel A: Fannie/Freddie Loans			
Measure of Lagged Economic Activity			
Lagged GSE Default Rate	-0.002	0.0009	0.013
	(0.014)	(0.0007)	(0.002)
	[0.008]	[0.004]	[0.009]
Lagged House Price Growth	0.42	0.18	-0.14
	(0.12)	(0.05)	(0.04)
	[0.022]	[0.016]	[0.008]
Lagged Unemployment Rate Change	0.005	0.003	0.003
	(0.004)	(0.005)	(0.001)
	[0.004]	[0.001]	[0.001]
Panel B: Prime Jumbo Loans			
Measure of Lagged Economic Activity			
Lagged GSE Default Rate	0.312		
	(0.078)		
	[0.125]		
Lagged House Price Growth	0.46		
	(0.74)		
	[0.024]		
Lagged Unemployment Rate Change	-0.027		
	(0.017)		
	[0.020]		

Note: The table shows regression coefficients from regressions of residualized interest rates on lagged MSA-level default rates, house price growth, and changes in the unemployment rate. Standard errors, clustered at the MSA level, are in parentheses. Adjusted residual removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases (top panel) and LoanPerformance (bottom panel).

Table 3: Lagged Local Default, Future Local Default, and Local Mortgage Rates

	Time Period	
	2001-2003	2004-2006
Panel A: Fannie/Freddie Loans		
Estimate of ψ_1 (Local Default on Future Default)	0.452 (0.042)	0.246 (0.033)
Estimate of β_1 (Local Mortgage Rate on Predicted Default)	0.0005 (0.0056)	0.0080 (0.0069)
Panel B: Prime Jumbo Loans		
Estimate of ψ_1 (Local Default on Future Default)	0.861 (0.176)	
Estimate of β_1 (Local Mortgage Rate on Predicted Default)	0.0460 (0.0121)	

Note: The table shows regression coefficients from regressions as specified, in log-log format. Local mortgage rate is measured as an adjusted residual, removing year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. "Future default" is defined as ever being 60 or more days delinquent within the first 24 months after origination. "Lagged default" measures the share of GSE loans originated in the MSA in the previous two years that are 60 or more days delinquent. "Predicted default" is the prediction from the regression of future default on lagged default. Standard errors, clustered at the MSA level, are in parentheses. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases (top panel) and LoanPerformance (bottom panel).

Table 4: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates

	Regional Lagged Default Rate				
	+2 Standard Deviation	+1 Standard Deviation	0	-1 Standard Deviation	-2 Standard Deviation
Percent consumption gain (λ * 100)	0.54%	0.30%	0.06%	-0.18%	-0.42%
Dollar per household effect ($\lambda \times 63790 \times \frac{\bar{c}_{region}}{\bar{c}_{overall}}$)	\$326	\$185	\$38	-\$118	-\$285

Note: The table shows the percent consumption gain and the dollar per household effect of the constant interest rate policy based on the level of the region's lagged default rate. See text for details.

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

	Regional Lagged Default Rate				
	+2 Standard Deviation	+1 Standard Deviation	0	-1 Standard Deviation	-2 Standard Deviation
Percent consumption gain (λ * 100): Overall	0.54%	0.30%	0.06%	-0.18%	-0.42%
Percent consumption gain (λ * 100): Young	0.18%	0.08%	-0.02%	-0.16%	-0.40%
Percent consumption gain (λ * 100): Middle Age	0.72%	0.40%	0.08%	-0.19%	-0.56%

Note: The table shows the percent consumption gain of the constant interest rate policy based on the level of the region's lagged default rate, by age. See text for details.

Table 6: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates, No Rental Market, by Age

	Regional Lagged Default Rate				
	+2 Standard Deviation	+1 Standard Deviation	0	-1 Standard Deviation	-2 Standard Deviation
Percent consumption gain (λ * 100): Overall	1.38%	0.72%	0.04%	-0.64%	-1.24%
Percent consumption gain (λ * 100): Young	1.52%	0.84%	0.12%	-0.58%	-1.22%
Percent consumption gain (λ * 100): Middle Age	1.28%	0.64%	-0.02%	-0.68%	-1.30%

Note: The table shows the percent consumption gain of the constant interest rate policy based on the level of the region's lagged default rate, by age, in the case without a rental market. See text for details.

Table 7: Sensitivity to Different Values of ϕ_r

	Regional Lagged Default Rate				
	+2 Standard Deviation	+1 Standard Deviation	0	-1 Standard Deviation	-2 Standard Deviation
Percent consumption gain (λ * 100): Benchmark	0.54%	0.30%	0.06%	-0.18%	-0.42%
1/2 Elasticity	0.28%	0.16%	0.02%	-0.10%	-0.22%
1/4 Elasticity	0.08%	0.00%	-0.08%	-0.16%	-0.24%

Table 8: Sensitivity to Other Elasticities

	Regional Lagged Default Rate				
	+2 Standard Deviation	+1 Standard Deviation	0	-1 Standard Deviation	-2 Standard Deviation
Percent consumption gain (λ * 100): Benchmark	0.54%	0.30%	0.06%	-0.18%	-0.42%
No Regional Income Variation	0.56%	0.24%	-0.02%	-0.30%	-0.50%
No Regional House Price Variation	0.20%	0.02%	-0.26%	-0.62%	-0.84%

Table A-1: Sample Description

	Time Period			
	All	2001-2003	2004-2006	2007-2009
Fannie/Freddie Loans				
Number of Loans	6,889,662	3,174,665	2,037,158	1,677,839
Median LTV	0.80	0.80	0.80	0.80
Fraction LTV > 0.8	32.8%	39.3%	24.3%	30.8%
Fraction FICO: 620-679	17.8%	20.8%	18.4%	11.5%
Fraction FICO: 680-719	20.7%	22.5%	20.7%	17.3%
Fraction FICO: >720	61.5%	56.7%	60.9%	71.2%
Prime Jumbo Loans				
Number of Loans	12,677	12,677		
Median LTV	0.80	0.80		
Fraction LTV > 0.8	37.7%	37.7%		
Fraction FICO: 620-679	63.9%	63.9%		
Fraction FICO: 680-719	18.2%	18.2%		
Fraction FICO: >720	17.8%	17.8%		

Table A-2: Mean and Standard Deviation of Lagged Local Economic Shocks

Lagged Economic Activity Measure	Time Period 2001-2003	2004-2006	2007-2009
Lagged GSE Default Rate			
Mean	0.834	0.675	1.062
Standard Deviation	0.381	0.398	0.650
Lagged House Price Growth			
Mean	0.093	0.194	0.025
Standard Deviation	0.052	0.087	0.054
Lagged Unemployment Rate Change			
Mean	0.485	-0.643	-0.316
Standard Deviation	0.750	0.394	0.471