

Regional Redistribution 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

October 2014

Abstract

An integrated tax and transfer system together with factor mobility can help mitigate local shocks within monetary and fiscal unions. In this paper we explore the role of a new mechanism that may also be central to determining the welfare effects of regional shocks. The degree to which households can use borrowing to smooth location-specific risks depends crucially on the interest rate and how it varies with local economic conditions. In the U.S., the bulk of borrowing occurs through the mortgage market and is heavily influenced by the presence of Government Sponsored Mortgage Enterprises (GSEs). 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. In contrast, we show that the private market does set interest rates based in part on regional risk factors and provide evidence that the lack of regional variation in GSE mortgage rates is likely driven by political pressure. We quantify the economic impact of the GSEs' constant interest rate policy on regional risk by building a structural spatial model of collateralized borrowing to match various features from our empirical analysis. The model suggests that the GSEs' national interest rate policy has significant ex-post redistributive consequences, with resources transferred across regions in state-contingent ways that are comparable in size to fiscal stimulus packages such as tax rebates and payroll tax holidays.

*First draft: April 2014. We thank Heitor Almeida, David Scharfstein, Monika Piazzesi, Stijn Van Nieuwerburgh, seminar participants at the Bank of Italy, HEC, National University of Singapore, University of Chicago (Booth and Harris), University of Illinois, University of Michigan, Wharton, the FRIC 2014 conference on financial frictions, and the NBER conference on Financing Housing Capital and the 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

How are local shocks mitigated within monetary and fiscal unions? This question has gained considerable attention in recent years as large disparities in regional outcomes have occurred within both the U.S. and Europe. There is a large literature arguing that an integrated tax and transfer system together with easy factor mobility can help mitigate local shocks.¹ In this paper we explore the role of an entirely different mechanism that may also be central to determining the welfare effects of regional shocks.

The degree to which households can borrow to self-insure against local shocks depends crucially on the interest rate and how it varies with local economic conditions. In the U.S., the vast majority of such borrowing occurs through the mortgage market. In this paper we empirically document the extent to which local mortgage rates (do not) vary with local economic conditions. Government Sponsored Mortgage Enterprises (GSEs) securitize most of the loans in the U.S. mortgage market and are bound by both economic and political constraints. We establish that, despite large regional variation in predictable default risk, there is essentially no spatial variation in GSE mortgage rates (conditional on borrower observables). If mortgage rates do not respond to local economic shocks that increase ex-ante local default probabilities, individuals in those regions may face lower borrowing costs than they would otherwise. Lower borrowing costs may help to offset the negative local economic shock that increased local default probabilities. Thus, this constant interest rate “policy” followed by the GSEs results in resources being transferred across regions in state-contingent ways.² Our objective is to quantify the magnitude of this ex-post redistribution across regions when interest rates are set using a constant national rate.

Our paper unfolds in three parts. We begin by using detailed loan level data securitized by the GSEs to show that local characteristics systematically predict future local loan default even after controlling for other observable borrower and loan characteristics. For example, there is medium-run persistence in local default probabilities: Regions that experienced higher default rates yesterday are more likely to experience higher default rates tomorrow (conditional on borrower and loan characteristics). Despite this finding, we further document that interest rates on loans securitized

¹See, for example, Farhi and Werning (2013) and the citations within. Additionally, Sala-i-Martin and Sachs (1991) and Asdrubali et al. (1996) explore the role of an integrated fiscal system in smoothing income across U.S. states. For a classic example of the importance of factor mobility, see Blanchard and Katz (1992). Recent examples include Farhi and Werning (2014), Charles, Hurst and Notowidigdo (2013), and Yagan (2014). Also see Feyrer and Sacerdote (2011) for arguments that the integrated tax and transfer system as well as the ease of factor mobility are reasons for the long run stability of the monetary union across U.S. states.

²Below, we discuss how the implicit government guarantee associated with the GSEs give the GSEs a cost of funds advantage over private lenders. This cost advantage allows them to charge lower interest rates on average. It also prevents the private market from competing away any pricing decisions by the GSEs that are not profit maximizing - up to the size of the cost advantage.

by the GSEs do not vary at all with predictable loan default. These patterns hold across different time periods and are robust to many different specifications to predict local mortgage default rates. The results are striking. Even though 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 an assessment of the extent to which GSE interest rates should vary spatially over time, given the large spatial variation in default risk. To do this, we exploit a loan level data set of loans securitized by private agencies. To facilitate comparisons, we focus on a set of loans that we refer to as being “prime jumbo” loans. GSEs by law are only allowed to securitize loans smaller than some threshold size, known as the conforming loan limit. Our prime jumbo loans are larger than those made by the GSEs but comparable on many other dimensions (in particular FICO score and LTV). Unlike the interest rate on GSE loans, we document that the interest rate on prime jumbo loans is strongly positively related to ex-ante local predicted default risk. Thus, while there is no regional risk-based pricing in the government-backed GSE market, the private market did set interest rates based in part on regional risk factors during 2000s. Employing a variety of econometric techniques, including a regression discontinuity approach around the conforming limit threshold, we construct counterfactuals as to the extent that GSE mortgage rates should have varied across regions within the U.S. – during both the early 2000s and during the Great Recession – if they priced local risk similarly to the private market. Again, these results are robust to controlling for many potentially confounding factors including potentially differential pre-payment propensities.

In the second part of the paper, we explore a number of explanations for why the relationship between mortgage rates and predictable local economic risk differs between the GSE and private markets. We conclude that political pressure is the most reasonable explanation for the patterns we observe. The GSEs face a great deal of political scrutiny. We provide evidence from prior efforts of the GSEs to differentiate lending standards and/or loan fees across regions, most recently through the ‘declining markets’ policy of 2008 and the ‘state based guarantee fee’ policy of 2012. Both of these policies were quickly abandoned in the face of pressure from Congress, realtors, and community groups. The tenor of the complaints all centered on an objection to the GSEs using different standards across regions. The lack of local variation in pricing rules shows up in many pricing decisions for the U.S. government. For example, the U.S. Postal Service charges the same flat rate for all first-class mail regardless of the distance traveled. Finkelstein and Poterba (2013) also find that political economy considerations can explain why U.K. insurance providers price nationally despite the presence of local drivers of mortality risk.

Our estimates allow us to do a simple back-of-the-envelope calculation of the extent to which resources are transferred across regions through the U.S. mortgage market. In particular, we conclude that if the GSEs had priced local economic risk as in the private market, there would be a 53 basis point difference in interest rates between the most risky and least risky MSAs based on ex-ante default risk (plus or minus two standard deviations from the mean). For an average loan amount, this pricing difference results in a roughly \$1,027 per year transfer from households in the least risky MSAs to those in the riskiest MSAs. A difference of this magnitude is sizable: it is comparable to other transfer programs, such as the transfer to a U.S. household from extending unemployment benefits by one year and larger than the tax rebate checks issued by the Federal government during the 2001 and 2008 recessions.

The simple back of the envelope calculation is likely to overstate the level of transfers through the U.S. mortgage market for two important reasons: 1) The back of the envelope calculation is static in nature and so assumes that regional differences are permanent. In reality, some regions which currently suffer from poor economic conditions and receive implicit transfers will face improved economic conditions in the future and will then be subject to implicit taxes, and 2) The simple back of the envelope calculation also takes households' mortgage holdings as given and does not allow households to re-optimize in response to policy changes. For example, if the GSE pricing rule was eliminated, households in regions with poor economic conditions would likely delay entry to the housing market and reduce the size of their houses to mitigate some of the negative effect of higher interest rates.

In the final part of the paper, we provide a more rigorous quantitative assessment of the economic impact of the GSEs' constant interest rate policy on transfers across regions. In particular, we build a structural spatial model of collateralized borrowing where households face region-specific shocks to house prices and labor earnings as well as purely idiosyncratic labor earnings risk. Individuals in the model can choose whether to own a home or to rent, in addition to choosing non-durable consumption and liquid savings over their life-cycle. Owner occupied housing is subject to fixed adjustment costs but serves as collateral against which individuals can borrow to smooth non-durable consumption. The model's consumption equivalents account for both re-optimization on the part of individuals and the persistence of the regional shocks. 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 rate applies to all regions. We use the empirical work in the first part of the paper to discipline the counterfactual interest rate policy.

We find that in the full structural model, the GSE constant interest rate policy generates large

transfers across regions. In our benchmark calibration, the GSE pricing policy generates a one-time \$680 tax on a region with a two-standard deviation increase in employment and generates a one-time subsidy of \$670 for a region with a two-standard deviation decrease in employment. This \$1350 one-time net transfer is equivalent to a \$120 annual transfer. Again, these transfers are of a similar size to those that arise from various other regional transfer policies. However, they are only one-eighth of the transfers implied by the simple back-of-the-envelope calculation. Thus, accounting for the fact that household behavior is not static and regional shocks are not permanent matters substantially for policy analysis.

Our model also allows us to explore more subtle distributional consequences that cannot be assessed in the reduced form calculation. In particular, we show that the GSE pricing policy has a much larger effect on middle-aged individuals than on young individuals. This is because the young mostly choose to rent and so are less sensitive to the local mortgage rate. In contrast, we show that if the young do not have access to these housing rental markets, then they are affected quite dramatically by the GSEs' constant interest rate policy, so that the option to rent has important welfare consequences.

Overall, our results suggest that during aggregate downturns the magnitudes of redistribution we observe through the mortgage market operated by the GSEs 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.

Our work relates to a few existing literatures. First, there is a small body of work looking at the extent to which risk is shared across U.S. states through credit markets. For example, Asdrubali et al. (1996) examine risk sharing across U.S. states and suggest that credit markets smooth about 23 percent of regional shocks. In that paper, the key mechanism is general borrowing and lending across regions. Our paper complements this finding by highlighting a direct mechanism by which the credit market serves to insure regional shocks. Our results suggest that the GSEs play a very large role in the extent that credit markets share risks across regions because interest rates do not vary with local ex-ante predictable default risk. This mechanism, as far as we can tell, is a novel addition to the regional risk sharing literature. Lustig and Van Nieuwerburgh (2010) directly explore the role of housing equity in supporting regional risk sharing within the U.S.. As housing equity increases, households are better able to borrow. The increased ability to borrow relaxes local liquidity constraints allowing local residents to better insure themselves to local shocks. Lustig and

Van Nieuwerburgh find that the extent of regional risk sharing varies with the state of the aggregate housing market.

Second, our work contributes to the recent surge in papers that have exploited regional variation to highlight mechanisms of importance to aggregate fluctuations. For example, Mian and Sufi (2014), Mian, Rao, and Sufi (2013), and Midrigan and Philippon (2011) have exploited regional variation within the U.S. to explore the extent to which household leverage has contributed to the Great Recession. Nakamura and Steinsson (2014), Shoag (2014) and Suarez Serrato and Wingender (2010) used sub-national U.S. variation to inform the size of local government spending multipliers. Blanchard and Katz (1992), Autor et al. (2013), and Charles et al. (2014) use regional variation to measure the responsiveness of labor markets to labor demand shocks. Beraja et al. (2014) develop a procedure to link regional variation in prices, wages and employment together with aggregate time series variation in those variables to infer the type of shocks that the U.S. economy experienced during the Great Recession. Stroebl and Vavra (2014) document a strong causal response of local retail prices to changes in housing wealth. Our work contributes to this literature by showing that the institutional features of the U.S. mortgage market help to support local lending in depressed regions by keeping interest rates low in those regions.

Additionally, our work contributes to the growing literature emphasizing that housing finance has important implications for the U.S. economy. Recent papers in this literature include Agarwal et al. (2012), Keys et al. (2014), Lustig and Van Nieuwerburg (2005), Mian, Sufi, and Trebbi (2014), Mian, Rao, and Sufi (2014), Mayer et al. (2009), Piazzesi et al (2007), and Scharfstein and Sunderam (2013).

Finally, while not the primary focus of our paper, our work speaks to the cost and benefits of the GSEs. 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 the 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. There is very little academic work on this topic.³ In our paper, we take a first step towards this end and formally explore one of the aspects of the GSEs' pricing decisions - the constant interest rate policy across U.S. regions. However, our paper is silent on the fact that the implicit subsidy to the GSEs may distort the allocation of capital towards the housing market and away from other productive resources. As we discuss in the conclusion, a more thorough analysis is needed to examine the overall impact of the GSEs on the U.S. economy.

³One exception is the recent book by Acharya et al. (2011).

This paper takes the GSEs as given and explores the consequences of their policies for regional risk sharing.

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).⁴ These loans are purchased, packaged, and insured against loss of principal and interest in the resulting mortgage-backed securities. As a premium, lenders pay a "guarantee fee" on each loan. 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.⁵

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.⁶

Figure 1 shows the share of mortgages in the secondary market that were securitized or directly held by the GSEs relative to the non-agency market during the 2000s. Over most of this 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

⁴Specifically, 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.

⁵The guarantee fee varies by some of features of the borrower (FICO score) and loan product (loan to value ratio), known as a "Loan Level Price Adjustment."

⁶"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.

Authority (FHA). The private market securitizes all other loans. The non-agency category includes jumbo mortgages (those that exceed the conventional mortgage size limits), subprime mortgages, and Alt-A mortgages. As seen from Figure 1, during the 2004-2006 period, the share of loans securitized by the private market grew at the expense of those loans securitized by the GSEs. In late 2007, the private secondary mortgage market dried up and essentially all securitization of mortgages since that time through late 2011 has been conducted by the GSEs.

Why do the GSEs dominate the conventional mortgage market? Researchers have estimated that the government's implicit guarantee to keep Fannie and Freddie solvent makes the cost of funds cheaper for the GSEs relative to the private market.⁷ Estimates suggest that mortgage rates for conventional mortgages are between 20 to 40 basis points lower than mortgage rates for otherwise similar jumbo mortgages.⁸ This difference is attributed to both the implicit guarantee and the scale of the GSE market. This cost differential makes it difficult for the private market to undo pricing mistakes made by the GSEs. If political constraints prevent the GSEs raising interest rates in declining markets and lowering interest rates in relatively strong markets, the cost of funds differential prevents private markets from competing with lower interest rates in relatively stronger markets. The cost differential, however, does bound how large the interest rate subsidy can be.

As discussed above, the GSE market is much bigger than the private market. This could, in part, contribute to their cost advantage. However, that cost advantage cannot explain the lack of pricing ex-ante regional default risk by the GSEs. One may also worry that the constant rate policy may be due to GSE loans being securitized which allows for better diversification of idiosyncratic and regional risk. Notably, our comparison will be with loans in the private market that are also pooled together and securitized. Thus, securitization per se cannot explain the absence of regional risk based pricing in one market and not in the other. Put another way, the focus of our empirical analysis will be to establish that, in contrast to loans sold in the private market, loans in the GSE market do appear to price systematic predictable expected defaults.

Finally, it is worth discussing who ultimately holds these securities and bears the risk of the mis-pricing. Although institutional investors may hold both GSE-backed and private mortgage-backed securities, only the private securities investments face default risk. In contrast, the GSEs guarantee the principal and interest payments of their mortgage-backed securities. Thus, the GSEs directly bear the risk of mis-pricing. From the investors' perspective, they only face the risk of early prepayment in GSE-backed mortgage securities. When the GSEs were publicly traded, their

⁷See, for example, Sherlund (2008).

⁸For a recent discussion of this literature, see Lehnert et al. (2008). The conclusion in the literature is that a large portion of the interest rate differential can be attributed to the existence of the GSEs.

shareholders also bore the risk that the GSE pricing model was not accurate. As experienced in 2008, the housing bust precipitated putting the GSEs into government conservatorship, and ultimately their losses were borne by taxpayers. In sum, the costs from failing to price local default risk are first borne directly by the GSEs, who fully insure securities-holders against default risk, and then indirectly by taxpayers who implicitly provide a government backstop.

III Data

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 1999 and 2012.⁹ 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 ratio), whether the loan is originated for purchase or 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 months to maturity, the outstanding mortgage balance, whether the loan is delinquent, the number of months delinquent, and whether the loan pre-paid.¹⁰ There is a unique loan identifier code in the data sets that allows a loan to be tracked from inception through its subsequent 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 conventional loans that are securitized by both the GSEs. As we show in the Online Robustness Appendix that

⁹The GSEs also securitized ARMs, especially during the recent housing boom (see Keys et al. 2013). However, since these loans account for a smaller part of their loan portfolio, information on such loans was not made available. Both the data sets were recently made available to increase the transparency of loans held or guaranteed by the two agencies. Each data set can be downloaded directly from the respective GSE websites.

¹⁰The 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.

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 both loans associated with a new purchase mortgage or a refinancing.¹¹ In total, our sample includes roughly 13 million loans in our combined analysis sample which were originated during the 2001-2006 period.

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 in the year of origination, (2) the mortgages 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.¹² As discussed in Keys et al. (2010), a FICO score of 620 is a cutoff above which loans are more likely 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 GSE sample. Very expensive loans may differ along many dimensions of loan and borrower characteristics compared to loans in the GSE sample. In essence, our prime jumbo loans are designed to be similar to the Fannie/Freddie loans in all respects except that the origination value of the loan is slightly higher. As with GSE mortgages, we include both originations for new purchases and refinancings. Finally, we restrict the sample to only include observations where there are at least 5 loan originations in an MSA and quarter of year cell. Our unit of analysis for exploring spatial variation in mortgage rates is at the MSA level. This restriction ensure that there will be a minimum amount of loans for each MSA-quarter cell. In total, our “prime jumbo” sample includes 70,560 loans originated during the 2001-2006 period.

¹¹The results are unchanged if we exclude refinance loans. The data appendix discusses other sample restrictions as well. In particular, we only include mortgages that have a FICO score at origination of at least 620 – the bulk of GSE data –, that were originated between January 2001 and December of 2009, and only include mortgages that were originated within one of our included MSAs.

¹²The 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.

III.C Time Periods

We focus our main analysis on the time period between 2001-2006. We start our data in 2001 given that we want to create a measure of lagged local default during the prior two years before origination. As discussed below, this will be our primary measure of local economic activity. Given that the GSE data begin with originations in 1999, our analysis sample must therefore start in 2001. We end our main analysis sample in 2006. The reason for this, as noted above, is because the private secondary market effectively ceased operation in 2007. The 2001-2006 period allows us to compare directly the spatial patterns in mortgage rates within both the GSE and prime jumbo samples.

We evaluate the robustness of the GSE results by also studying the pricing behavior during the recession. Given the much larger local variation in both actual and predicted defaults during the Great Recession, we explore if spatial patterns and pricing were amplified during the 2007-2009 period.

III.D Additional Sample Restrictions

Table 1 provides descriptive statistics for both our GSE sample (column 1) and our prime jumbo sample (column 4) without any further sample restrictions during the 2001-2006 period. A few things are of note about the GSE sample relative to the prime jumbo sample. First, borrower quality looks higher in the GSE sample despite our initial restrictions on the prime jumbo sample. In the full GSE sample, the average FICO score of borrowers is 728. The comparable number in the prime jumbo sample is only 656. Second, the GSE data covers 374 distinct metropolitan statistical areas (MSAs). However, prime jumbo loans are only in 106 distinct MSAs. This is not surprising given that the origination amount on a prime jumbo loan has to exceed a relatively large value. For many MSAs in the U.S., it is rare for a property to transact above the conforming loan limit. As average property values in the MSA increase, the probability that loans exceed the conforming loan threshold also increases.

To further facilitate comparison between the GSE data and the prime jumbo data, we make two additional sets of restrictions to the GSE data. First, we restrict the GSE data to only include loans for the 106 MSAs where we have at least five observations of prime jumbo data. Specifically, we ensure that the MSA-quarter coverage between the two samples is identical. This reduces the sample size of GSE loans from 13.1 million loans to 8.1 million loans. Descriptive statistics for this sample are shown in column 2 of Table 1. This restriction does not alter the borrower quality comparisons at all: It is still the case that the MSA-matched GSE sample had higher FICO scores

than the prime jumbo sample.

Our second set of restrictions is more substantial. Here we restrict the GSE sample to match the prime jumbo sample on three dimensions. First, we restrict to the same MSA-quarter coverage (as discussed in the prior paragraph). Second, we restrict the sample so that the sample size matches exactly.¹³ This is important given that when we measure the variability of interest rates and default rates across MSAs, we want to ensure we have similar power within the two samples. Third, we restrict the GSE sample so that it replicates the FICO distribution of the prime jumbo sample.¹⁴ As a result, the distribution of borrower quality as measured by FICO score will not differ between the two samples.¹⁵ We refer to this sample as the “matched” GSE sample where the matching occurs on MSA-quarter, FICO score and sample size. Descriptive statistics for the matched GSE sample are shown in column 3 of Table 1. Given the matching procedure, it is not surprising that the median FICO variation and the MSA coverage match exactly with the prime jumbo sample. Notably, the average LTV of loans at the time of origination is also similar, despite not being an attribute on which the samples were directly matched. This matched GSE sample will be our main analysis sample going forward.

Table 1 also shows the average interest rate on the loans within each sample. Consistent with the literature, the unconditional interest rate on the GSE loans during this period was about 36 basis points lower than the rate on prime jumbo loans (6.30% vs. 6.66%). Throughout the paper, 60+ days delinquent will be our primary measure of default. Table 1 measures the fraction of loans that became 60+ days delinquent at some point during the two years after origination. Unconditionally, 2.9% of the GSE loans in the matched sample become delinquent in the two years after origination while only 2.1% of the prime jumbo loans become delinquent. As we show below, conditioning on the date of origination and focusing on loans originated around the conforming limit cutoff, we show that ex-post delinquency measures are nearly identical between the two samples.

III.E 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

¹³In the data appendix that accompanies this paper, we go into detail about the sample restrictions.

¹⁴Approximation error in randomly sampling similar FICO loans resulted in sample sizes being slightly different. Our final sample includes 70,560 loans in the GSE matched sample and 70,680 loans in the prime jumbo sample.

¹⁵All of these sample restrictions were made to ease comparison of the two samples. However, given that all of our estimation procedures also include controls for observable loan and borrower characteristics, the matching did not make much difference. In many of our tables, we show the results with and without restricting the samples to be similar in size and FICO distribution. The results are nearly identical across the specifications. In the data appendix that accompanies the paper, we go through in detail the exact selection criteria for our main sample to facilitate the replication of our results.

interest rates and delinquency rates could differ spatially is because borrower and loan characteristics could differ spatially or because borrowers in the two samples originated their loans in different time periods. 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. Of course, matching the two samples on FICO score mitigates some of this concern. What we are after, however, is whether interest rates and the predictable component of default rates vary spatially conditioning on borrower and loan characteristics. A borrower with a given credit score and LTV ratio may be more likely to default in one region relative to another because overall economic conditions differ across regions.

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

$$r_{ist}^j = \alpha_0^j + \alpha_1^j X_{it} + \Gamma_1^j D_t + \Gamma_2^j D_t \cdot X_{it} + \eta_{ist}^j$$

$$y_{ist}^j = \varphi_0^j + \varphi_1^j X_{it} + \Psi_1^j D_t + \Psi_2^j D_t \cdot X_{it} + \nu_{ist}^j$$

where r_{ist}^j is the loan-level mortgage rate for a loan made to borrower i , in MSA s , during period t and y_{ist}^j is an indicator variable for whether the loan made by borrower i , in MSA s , during period t defaulted at some point during the subsequent 24 months. X_{it} is a set of control variables for borrower i in period t . 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 of origination. The borrower/loan controls include detailed FICO and LTV controls. Specifically, all regressions include quadratics in FICO and LTV, and each of these terms are fully interacted with quarter of origination dummies.¹⁶ The goal of these specifications is to recover η_{ist}^j and ν_{ist}^j , the residual mortgage rate and residual ex-post delinquency rate, respectively, for borrower i in MSA s during time t for loans in sample j after controlling for borrower/loan characteristics and time fixed effects.

Once we have the residuals from the above regressions with the full set of controls, we compute location specific average mortgage rates, R_{st}^j , and location specific average ex-post default rates,

¹⁶In results not shown, we have also experimented with 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 these LTV and FICO score controls to be interacted with the time dummies. Other specifications included interactions between FICO and LTV. The results are quantitatively and qualitatively similar.

Y_{st}^j . We do this separately for each time period and for each sample. Specifically,

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

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

where N_{st}^j is the number of loans in the MSA s during quarter*year t within each sample. Formally, R_{st}^j (Y_{st}^j) will be the average mortgage rate residual (ex-post delinquency residual) in an MSA for loans originated during a given period for a given sample.

The bottom rows of Table 1 show the standard deviation of unconditional mortgage rates (r_{ist} 's) across MSAs and the standard deviation of conditional mortgage rates (R_{st}^j 's) across the MSAs for our matched GSE sample and our prime jumbo sample during the 2001-2006. Similarly, we show the standard deviation of unconditional and conditional delinquency rates across MSAs. The cross-MSA variation in interest rates is reduced dramatically once conditioning on borrower, loan and time controls. Additionally, the conditional cross-MSA standard deviation of mortgage rates is twice as high in the prime jumbo sample relative to the matched GSE sample while the conditional cross-MSA standard deviation of delinquency rates are similar between the two samples.

IV Local Mortgage Rates and Predictable Local Default Risk

In this section, we document our key empirical facts. As we will illustrate, GSE mortgage rates do not vary at all with measures of local default risk while the prime jumbo rates do vary with the local default risk measures.

IV.A A Metric For Local Economic Activity

In order to examine whether mortgage rates vary with local economic conditions, we need to define measures of local economic activity observable to lenders that could potentially be used in their pricing decisions. Our primary measure of local economic activity is the lagged delinquency rate on loans securitized within each sample. Specifically, within each MSA s in period t , we measure the fraction of loans originated during the prior two year period that defaulted at some time between their origination and period $t-1$. Given our unit of analysis is a quarter, our measure is the fraction of all loans originated between 9 quarters prior and 1 quarter prior that became 60 day delinquent

by the current quarter.¹⁷ We refer to this measure as $E_{s,t-1}^j$ where $E_{s,t-1}$ denotes lagged economic activity in location s prior to the current period. We index this measure by j because we could measure lagged defaults in either the GSE sample or in the prime jumbo sample. 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 labor markets, declining house prices) and because it is easily observable by lenders.¹⁸

To describe the data, Figure 2a show a simple scatter plot of local mortgage rates residuals for the GSE loans, R_{st}^{GSE} , in the matched sample against lagged local GSE default rates, $E_{s,t-1}^{GSE}$, during the 2001-2006 period. Figure 2b presents the same result for the sample of MSAs present in both datasets. Figure 2c analogously shows the scatter plot of local mortgage rates residuals for the prime jumbo loans, R_{st}^{jumbo} , against lagged local GSE default rates, $E_{s,t-1}^{GSE}$, during the same time period. For these pictures, we define the measure of local economic activity as being the lagged GSE default rate. This is done for illustrative purposes so that the variation on the x-axis is the same in both pictures. Below, we show how these measures of lagged default can be used to create measures of predicted default for both samples. Each observation in the figures is a MSA-quarter pair.

Figures 2a and 2b shows that there is no relationship between lagged local GSE default rates and average local mortgage rates. Column (2) of Table 2 summarizes the regression line of the scatter plot in Figure 2b. A one percentage point increase in lagged GSE default is associated with an increase in local GSE mortgage rates by only 2.4 basis points (i.e., from 6.000 to 6.024).¹⁹ The standard deviation of lagged GSE default across the MSAs is 0.7 percentage points. This implies that a two standard deviation increase in lagged default is associated with only a 1.8 basis point increase in local GSE mortgage rates. This estimate is essentially a precise zero. The patterns in Figure 2b are for the restricted GSE sample during the 2001-2006 time period. Columns (1) and (3) of Table 2 show the same estimated association in the full sample of GSE loans and in the sample of GSE restricted to include only the MSAs covered in the prime jumbo sample and matched to the prime jumbo FICO score distribution. In all samples, there is no economically

¹⁷We have experimented defining our measure of local economic activity as the delinquency rate on loans originated over the prior 4 quarters and the prior 6 quarters. All results in the paper were essentially unchanged using these alternate definitions of local economic activity.

¹⁸We also used both the lagged local unemployment and lagged housing price growth as our measure of local economic activity. Results were generally similar. The one difference was that lagged local house price growth during the early 2000s negatively predicted local mortgage default while lagged local house price growth during the mid 2000s positively predicted local mortgage default. The latter result was driven by the fact that local house price growth during the mid 2000s predicted local house price declines during the late 2000s and households are more likely to default when house prices decline. Given this, we focused on using lagged default as our measure of local economic activity.

¹⁹When fitting a line through the scatter plot, we weight each observation by the number of loans originated during the MSA-quarter. As a result, larger MSAs with more loans are weighted more when fitting the line. All results in the paper are weighted in a similar manner.

meaningful relationship between lagged GSE default and GSE mortgage rates. Finally, columns (5) and (6) show that the 2001-2006 patterns persisted through the 2007-2009 period. During the Great Recession, there was also no statistical relationship between lagged local mortgage defaults and local mortgage rates.

The pattern in Figure 2c is in stark contrast to those in Figures 2a and b. Figure 2c shows that there is a strong positive correlation between lagged GSE default rates and local interest rates for prime jumbo loans. MSAs which had larger GSE defaults in the prior year originate loans with higher interest rates conditional on borrower and loan characteristics. Column (4) of Table 2 shows that a one percentage point increase in lagged local GSE default rates was associated with a 31 basis point increase in local prime jumbo mortgage rates. This is 10 times larger than the effect on GSE mortgage rates. This implies that a two standard deviation increase in lagged default is associated with a 23 basis point increase in local GSE mortgage rates.

IV.B Relationship between current local economic activity and future default

The results in Figure 2 and Table 2 suggest that prime jumbo mortgage rates are more responsive to local economic activity than are GSE mortgage rates. However, this finding may be due to the fact that default rates on prime jumbo loans are more responsive to lagged local economic conditions than are the default rates on GSE loans. In this sub-section, we assess the extent to which lagged local economic conditions predict subsequent actual default. In the next sub-section, we will address the concern and relate local interest rates on both GSE loans and prime jumbo loans to local measures of predicted loan default.

We refer to predicted local default for loans in each sample j , in each location k during each time period t as \hat{Y}_{kt}^j . We calculate three measures of predicted default. Our first and primary measure predicts the conditional relationship between future delinquency and current economic conditions. In particular, we regress the following on both the GSE and prime jumbo samples:

$$y_{ikt}^j = \theta_0^j + \theta_1^j X_{it} + \Psi_1^j D_t + \Psi_2^j D_t \cdot X_{it} + \lambda^j E_{k,t-1}^{GSE} + \nu_{ikt}^j$$

where y_{ikt}^j , X_{it} , D_t and $E_{k,t-1}^{GSE}$ are defined above. In particular, we use the underlying micro data to see whether lagged GSE default rates predict subsequent mortgage default (conditional on loan and borrower observables). We use the lagged GSE default rate for both samples so that we can examine the actual default rate responsiveness to similar underlying lagged economic conditions.

For the matched GSE sample during the 2001-2006 period, our estimate of λ^{GSE} is 1.75 (standard error = 0.30). For the prime jumbo sample during the same period, our estimate of λ^{jumbo} is 2.51 (standard error = 0.31). Using the above relationship, we define our first measure of predicted local mortgage default as:

$$\hat{Y}_{kt}^j = \lambda^j E_{k,t-1}^{GSE}$$

One may wonder if the relationship between lagged GSE default and future default is an artifact of the period we studied. We explored this possibility by re-running the above relationship for the GSE data during the 2001-2003 period, the 2004-2006 period, and the 2007-2009 period. In all three sub-periods, lagged GSE default positively and significantly predicted future default rates. We also examined the relationship for the prime jumbo sample during the 2001-2003 period and the 2004-2006 period. Again, in both sub-periods, lagged GSE default predicted future loan default. Collectively the results illustrate that during the 2000s, there was consistent medium-run persistence in mortgage default rates.

For robustness, we also explored two additional measures of lagged default. The first we refer to as our “random walk” forecast such that:

$$\hat{Y}_{kt}^j = E_{k,t-1}^j$$

This specification says that the best forecast of today’s loan default rate is yesterday’s default rate. Notice, for each sample, the lagged default rate is sample specific. This differs from the first predicted default measure where both the future default rates of loans in the GSE sample and the prime jumbo sample depended on the lagged GSE default rate. This allows for lagged default rates on the prime jumbo sample to have better predictive properties for loans in the prime jumbo sample than would lagged GSE default rate.²⁰

Second, we examine a “perfect foresight” prediction of future default such that:

$$\hat{Y}_{kt}^j = Y_{k,t}^j$$

In this specification, the best prediction of future default in a given sample in a given location conditional on observables is the actual future default rate for a given sample in a given location.

Above, we define this latter variable as $Y_{k,t}^j$.

²⁰Lagged local jumbo default rates are also highly predictive of future jumbo default rates. Conditional on loan and borrower observables, a 1 percentage point increase in lagged local default increased the probability of future jumbo default by about 1 percentage point during the pooled 2001-2006 period.

IV.C Local Mortgage Rates and Predictable Default

To examine whether the mortgage rates on GSE loans and the mortgage rates on prime jumbo loans respond similarly to predicted local default, we estimate the following equation separately for each sample during the 2001-2006 period:

$$r_{ikt}^j = \alpha_0^j + \alpha_1^j X_{it} + \Gamma_1^j D_t + \Gamma_2^j D_t \cdot X_{it} + \beta^j \hat{Y}_{kt}^j + \eta_{ist}^j$$

The regression is nearly identical to the ones above explaining mortgage rate variation aside from the addition of the predicted default variable. The coefficients of interest are β^{GSE} and β^{jumbo} . Column (1) of Table 3 shows our estimates of β^{GSE} for our three predicted default measures while the second column shows our estimates of β^{jumbo} . Columns (3) and (4) shows the difference between the coefficients ($\beta^{jumbo} - \beta^{GSE}$) as well as the p-value of the difference.

Regardless of the measure of predicted default, the responsiveness of mortgage rates in the prime jumbo market to predicted default measures is much higher than the responsiveness of mortgage rates in the GSE sample. Even controlling for the fact that the responsiveness of actual default rates to lagged economic conditions may differ between the two types of loans, interest rates respond much more to local economic conditions in the prime jumbo sample. One further claim can be made. It is not only that the prime jumbo interest rates respond more, it is also that the GSE interest rates do not respond in any economically meaningful way to the predicted default measures. A one percentage point change in local predicted default will only increase local GSE mortgage rates by 2 basis points. In the next section, we will show that all three of the lagged default specifications yield similar differential variations in interest rates between the two samples once scaled appropriately by the underlying variation in the predicted default metric.

Given the nature of our data, we can also explore the responsiveness of local mortgage rates to measures of local predicted default around the conforming loan threshold. This allows for even better comparisons relative to our earlier matched sample approach. Specifically, we can get our estimate of ($\beta^{jumbo} - \beta^{GSE}$) using a regression discontinuity approach. Specifically, we estimate:

$$r_{ikt}^j = \alpha_0 + \alpha_1 X_{it} + \Gamma_1 D_t + \Gamma_2 D_t \cdot X_{it} + (\tilde{\alpha}_1 X_{it} + \tilde{\Gamma}_1 D_t + \tilde{\Gamma}_2 D_t \cdot X_{it}) D_{it}^{jumbo} + \Upsilon_2 Bin_{it} + \Upsilon_2 Bin_{it} \cdot \hat{Y}_{kt}^j + \eta_{ist}^j$$

For this regression, we pool together both the 2001-2006 prime jumbo sample and the 2001-2006 matched GSE sample. D^{jumbo} is a dummy variable indicating that the loan is from the prime jumbo sample. As seen from the regression, we allow the responsiveness of mortgage rates to observables

(FICO, LTV) and time effects to differ across the two samples. For tighter comparison with the GSE sample, we only included loans up to two times the conforming threshold in the prime jumbo sample.

The key addition to this specification is the variables Bin_{it} and $Bin_{it} \cdot \hat{Y}_{kt}^j$. For each loan, we compute a metric of the mortgage size relative to the conforming loan threshold. Loans above the conforming threshold will have a metric that ranges from 1 to 2. These loans will all be from the prime jumbo sample. Loans below the conforming threshold will have a metric that ranges from 0 to 1. As noted above, we also created our matched GSE sample so that it has a similar distribution of loan sizes below the conforming threshold as the prime jumbo sample has above the conforming threshold. The variable Bin_{it} is an indicator variable for the extent to which the loan size differs from the conforming threshold. Specifically, the Bin_{it} variable is defined in 0.2 unit intervals of the ratio of the loan size to the conforming loan limit (e.g., 0.8-1, 1-1.2, 1.2-1.4, etc.). For example, loans in the 1-1.2 bin have an origination value that is between the conforming limit and 20 percent greater than the conforming limit. The regression includes dummy variables for all 10 bin values and allows the responsiveness of local interest rates to our measures of local predicted default to differ across the bins.

One concern for us would be if there was a large amount of selection for the loans above the conforming threshold. Even though these loans may be similar on observables, they may differ on unobservables. This type of selection would not be surprising given the large financial benefit in terms of lower average interest rate for GSE loans relative to prime jumbo loans. As a result, better borrowers may migrate to the GSE sample by choosing to put up more equity and take out a loan smaller than the conforming threshold. We explore these issues in Figure 3. Figure 3 shows that there is no discrete change in FICO scores across the conforming threshold. This is not surprising given that the samples were matched on FICO scores.

Figure 4 explores whether there is selection on unobservables at the conforming threshold. It does so by comparing the default rates of the GSE loans right below the threshold with the default rates for the prime jumbo loans right above the threshold. If there was selection, one would imagine that better borrowers put up more cash so that they secure a loan lower than the conforming threshold. Figure 4 shows that there is a very slight increase in default probabilities for jumbo loans in the first bin above the conforming threshold relative to the first bin below the threshold (differential actual default probability = 0.004 with a standard error of 0.001). While the difference in actual default rates is small, it does appear that some selection is taking place. However, the second bin above the threshold shows no differential ex-post default probabilities relative to the GSE loans just below

the threshold. The differential default probability between GSE loans close to conforming limit and loans in the second bin above the threshold is close to 0.001 with a standard error of 0.001. Similar results hold for the third, fourth and fifth bins above the threshold. So while there may be a small amount of selection occurring within the first bin above the threshold, there does not seem to be any evidence of selection in the other bins that is correlated with actual loan performance.²¹ Given the potential of selection in the first bin above the threshold, we also examine the sensitivity of our estimates of $(\beta^{jumbo} - \beta^{GSE})$ when we compare the first GSE bin below the threshold to both the first and second prime jumbo bin above the threshold. As we show below, the results are nearly identical under the two procedures. This is not surprising given that the implied differential default probabilities just above the threshold appears quite small.

Figure 5 shows our estimates of β for each of the 10 bins using our three default measures. The results are, again, striking. The responsiveness of local mortgage rates to local predicted default rates is essentially zero for all bins below the conforming threshold, regardless of our definition of predicted default. However, for the bins directly above the conforming thresholds, there is a strong positive relationship between local default probabilities and local mortgage rates. The estimated responsiveness is nearly identical in the second, third, and fourth bins above the threshold. The results, combined with the actual default analysis in Figure 4 show that the pricing behavior of mortgages with respect to local default risk changes discretely between the GSE and prime jumbo samples. Column (5) of Table 3 shows our RD estimates of the differences in responsiveness for our three measures of predicted default. Our RD estimates are very similar to the regression based estimates shown in column (3) of Table 3.

IV.D Robustness of Main Results

Before turning towards interpreting these results, it is useful to explore the robustness of our results to alternate specifications and controls. In particular, we examine whether our results hold controlling for prepayment risk, MSA fixed effects, and points and other fees paid borrowers.

IV.D.1 Prepayment Risk

One concern with the empirical work above is that we did not account for other potential local risks that could affect local loan pricing. In particular, aside from default risk, the biggest risk lenders face is prepayment risk. If prepayment risk differs dramatically between GSE loans and

²¹This is not surprising given that the second bin has a loan value that is, on average between \$40,000 and \$80,000 above the threshold. It is hard for most households buying a \$500,000 home to easily provide that amount of money to reduce the loan balance so that it could be securitized by the GSEs.

prime jumbo loans in a way that is correlated with local default risk, the lack of variation in GSE mortgage rates with local default risk may not be surprising.

To this end, we created a measure of predicted local prepayment risk. We followed the same procedures as above for local default risk in that we created three different measures: the regression based approach for both samples using lagged local GSE prepayment rates, a perfect foresight model where the predicted prepayment rate was the actual local prepayment rate for each sample, and random walk model where the predicted prepayment rate was the actual lagged local prepayment rate for each sample. Our first finding is that the predicted prepayment rates, conditional on loan and borrower observables, were very similar between the GSE loans the prime jumbo loans. For example, using our RD approach, predicted prepayment rates were only 2 percentage points lower for prime jumbo loans above the conforming threshold relative to the GSE loans below the threshold.

Given there was a slight difference in prepayments between the two samples, we added the local predicted prepayment rate to both the bivariate figures (shown in Figure 6a and 6b), and our base RD specification when estimating $(\beta^{jumbo} - \beta^{GSE})$. The results of this exercise are shown in Table 4. We focus on our main specification where predicted local default is based on a regression of actual default on lagged GSE default and controls for both samples. Column (1) of Table 4 redisplayes the base RD results from column (5) of Table 3). Column (2) shows the results when we add the regression based measure of predicted prepayments to the RD regression. Controlling for predicted prepayment risk did not change the RD estimates in any meaningful way. Again, this is not surprising given the fact that conditional prepayment probabilities hardly differed between the samples.

IV.D.2 MSA Fixed Effects

Another potential concern with our results above is that we may be identifying our results across MSAs where the composition of GSE loans and prime jumbo loans differ. To account for this potential, we re-estimated all our specifications including MSA fixed effects. This allows us to compare prime jumbo loans within an MSA to GSE loans within the same MSA. Column (3) of Table 4 controls for MSA fixed effects while column (4) controls for both MSA fixed effects and local prepayment risk. As can be seen from the table, the estimated difference in interest rate responsiveness to local default risk, $(\beta^{jumbo} - \beta^{GSE})$, is essentially unchanged in all the specifications.

IV.D.3 Points and Other Fees

Finally, we have not examined regional variation in points paid or other loan fees. It may be the case that mortgage rates do not vary across MSAs in the GSE sample, but points and other fees do vary with local default risk. To address this concern, we secured additional data from LoanSifter. LoanSifter is a company that collects loan quotes from various lenders about the interest rate they charge for a given loan type where loan type is defined as a function of FICO score, fixed interest rate, points charged, and initial LTV. From them, we were able to secure fixed-rate loan quotes from banks for a given size loan (\$300,000), a given LTV (80%), no points, and three FICO score levels (750, 680, 620) during the period of September 2009 through September 2010. The key for this data is that points are held fixed across all loan quotes. Given the loan size, all quotes were for loans all eligible for securitization by the GSEs. Within this data, we found no relationship between quoted mortgage rates for a given contract at the local level and local measures of default risk.

The limitation of the LoanSifter data is that it covers only a few loan types during the 2009-2010 period. To examine whether our results were skewed because we did not have data on points and fees, we used the Federal Housing Finance Agency (FHFA) Mortgage Interest Rate Survey (MIRS), which collects data from lenders on initial fees and charges and effective interest rates, by state, from 1978 through 2012.²² In results not shown, we find no evidence of either interest rates, points and fees, or effective interest rates systematically varying with local measures of default risk at the state level.

V Interpreting the Results

In this section, we perform two related analyses. First, we construct a counterfactual of how much GSE interest rates should have varied across regions if local risk was priced similarly to the prime jumbo sample. Second, we start to assess how the lack of risk based pricing transferred resources across regions of the U.S. during the late 2000s.

V.A How Much Should Have GSE Loan Rates Varied With Predictable Default?

Table 5 shows the standard deviation of predicted default for our three default measures. The first and second columns examine the standard deviation of predicted default for our matched GSE

²²This data, found in Table 15 of the MIRS annual reports, can be obtained from the FHFA website, www.fhfa.gov.

sample and our prime jumbo sample during the 2001-2006 period. The last column examines the predicted default measures a sample of GSE loans restricted to the same MSAs as the prime jumbo loans, but during the 2007-2009 period.²³

Table 6 is our key counterfactual table. Given the standard deviation of predicted default rates (Table 5), Table 6 computes how much GSE interest rates should have varied across regions in response to a two standard deviation change in predicted default. We use our baseline RD coefficients to perform the counterfactual (column (1) of Table 4). Given the stability of coefficient estimates across the specifications, it matters little to our analysis if we use the baseline RD coefficients, the RD coefficients with additional controls, or our the regression based estimates shown in column (3) of Table 3. Table 6, therefore, computes the counterfactual by multiplying our estimate of $(\beta^{jumbo} - \beta^{GSE})$ by two times the relevant standard deviation of predicted default. Our preferred estimates (row 1 of Table 6 which uses the regression measure of predicted default) suggests that a two standard deviation shock to predicted default should have resulted in a 16 basis point variation in GSE mortgage rates across regions during the 2001-2006 period and a 29 basis point variation in GSE mortgage rates across regions during the 2007-2009 period.

The other specifications of lagged default give roughly similar estimates. In our modeling section below, we calibrate the regional transfers that occurred through the mortgage market due to the GSE's constant interest rate policy when the economy experiences a shock the size of the Great Recession. We will choose parameters such that a two standard deviation shock in local economic activity across regions would result in a 25 basis point variation in mortgage rates across regions if the GSEs allowed mortgage rates to adjust to local default risk similar to those in the prime-jumbo market. Given our counterfactual estimates for the other predicted default measures shown in Table 6, we examine the robustness of our model results when a two standard deviation shock causes only a 15 basis point variation or a 35 basis point variation in mortgage rates across regions.

V.B A Back of the Envelope Calculation for Regional Redistribution

Our estimates allow us to make a naïve back-of-the-envelope calculation of the extent to which resources are transferred across regions through the U.S. mortgage market. To illustrate this calculation, we focus on predicted defaults using lagged GSE default rates as a measure of local economic conditions. To make the results comparable to the structural analysis that follows, we focus on the top and bottom 2.5 percent of the predicted default distribution (that is, two standard deviations

²³When predicting local default during the 2007-2009 period using lagged default, we use the coefficients from the 2001-2006 sample.

above and below the mean). This allows us to focus on transfers between regions which experienced (relatively) good economic conditions to those that experienced (relatively) bad conditions during the Great Recession period of 2007-2009.

On average, the 2.5 percent of mortgages with the lowest default risk experienced a default rate of 0.5% over the prior two years, whereas the top 2.5 percent of mortgages saw default rates of 4.8% in the prior period. We then take the coefficients from the estimated relationship between lagged GSE defaults and interest rates in each market and apply them to the default risk distribution during 2007-2009. The estimated relationship between predicted defaults and interest rates in the GSE market is 1.51, while the relationship in the prime jumbo market is 11.94. Using these estimates of the relationship, we can apply the prime-jumbo counterfactual to determine the extent of regional transfer. In the lowest 2.5 percent of the predicted default distribution, had GSE pricing accounted for regional risk like in the prime-jumbo market, the regional risk adjustment would be 4.5 basis points. Similarly, in the highest 2.5 percent of the predicted default distribution, the regional risk adjustment would be 58 basis points. Thus the difference in interest rates between top and bottom tails of the default distribution, had GSE pricing accounted for regional risk as in the prime-jumbo market, would be 53 basis points.

In order to obtain a dollar equivalent of the regional risk sharing premia in this counterfactual scenario of GSE regional risk-based pricing, we can multiply these estimates by the average loan amount in each quintile. Doing so yields a regional risk adjustment of \$117 per loan in the lowest default tail, while the adjustment in the upper default tail is \$1,144. Comparing across the top and bottom tails of the risk distribution, we find that the difference in annual dollar costs in the presence of a model that prices regional risk would be \$1,027.

The simple back of the envelope calculation is likely to overstate the level of transfers through the U.S. mortgage market for two important reasons. First, this calculation is static in nature and thus assumes that regional differences are permanent. In reality, some regions which currently suffer from poor economic conditions and receive implicit transfers will face improved economic conditions in the future and will then be subject to implicit taxes. Second, and more importantly, the naïve calculation also takes households' mortgage holdings as given and does not allow households to re-optimize in response to policy changes. For example, if the GSE pricing rule was eliminated, households in regions with poor economic conditions would likely delay entry to the housing market and reduce the size of their houses to mitigate some of the negative effect of higher interest rates. We will deal with these issues directly when we provide a more rigorous quantitative assessment of the economic impact of the GSEs' constant interest rate policy on transfers across regions.

V.C 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.²⁴ 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 95%. However, in declining markets, Fannie/Freddie would only purchase mortgages where the initial LTV was lower than 90%.²⁵ The policy did not affect interest rates, it only affected underwriting standards.

The Declining Market 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.²⁶ 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.”²⁷ 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.”²⁸ Even though it may have been profitable to require different downpayments in different areas, Fannie Mae and Freddie Mac succumbed to political pressure and quickly abandoned the policy.

In September of 2012, the Federal Housing Finance Authority (FHFA), which now oversees the

²⁴Many 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.

²⁵Fannie 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.

²⁶See “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>

²⁷See “Fannie Is Poised to Scrap Policy Over Down Payments,” *Wall Street Journal*, May 16, 2008.

²⁸See “Looser Credit on the Way in ‘Declining’ Markets?” *Washington Post*, May 24, 2008.

GSEs, proposed a new 25 basis point fee at the time of origination that differed across locations. The fee was tied to states that had long judicial delays in foreclosures. The rationale was that institutional features in those states tied up defaulted properties in the foreclosure process for longer periods of time which increased GSE losses in those states. At the time of its original announcement, the fee would only have been applied to loans originated in New York, California, Florida, Connecticut, and Illinois. These were the states with the longest foreclosure delays. In late 2012, FHFA invited comments on the proposal from the public. Like the Declining Market Policy, this policy received a large amount of public backlash. In particular, the governor of Illinois wrote a detailed public comment against the new fee. In December of 2013, the FHFA announced that despite the backlash they were going to implement the fee increase in the previously announced state (excluding Illinois) in April of 2014. In January of 2014, after another round of political pressure, the FHFA announced that the policy to charge differential state-based guarantee fees has been delayed indefinitely.

Even though these policies focused on imposing either spatial variation in downpayments or spatial variation in loan guarantee fees, they can shed light 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 wish to highlight and quantify. By foregoing profit-maximizing behavior and charging a constant interest rate across all regions despite different levels of predictable default risk, the GSEs redistribute resources towards markets with weaker economic activity and greater default risk.

VI Model description

We now turn to a more formal quantitative assessment of the welfare effects of the GSEs' constant interest rate policy. As discussed earlier, the back-of-the-envelope calculation described above provides a jumping off point for quantifying the size of the implied transfers from low default regions to high default regions, but is likely to be overstated because it holds two important objects fixed, which are not fixed in reality. The calculation takes current local economic conditions and the regional distribution of household mortgage holdings as given, and computes the change in a location's total implied mortgage payments if interest rates varied with local default instead of remaining constant. However, differences in local default rates are not permanent, and household

mortgage holdings will endogenously change in response to changes in interest rate policy. A more credible assessment of the size of regional transfers induced by the GSE pricing rule requires a formal model that can account for both transitory local shocks and household re-optimization in response to changes in the economic environment.

In this section, we lay out a quantitative spatial model, which captures various salient factors of the U.S. housing market. In this multi-region life-cycle consumption savings model, households face region-specific shocks to house prices and labor earnings as well as purely idiosyncratic labor earnings risk. Individuals in the model can choose whether to own a home or to rent, in addition to choosing non-durable consumption and liquid savings. Owner occupied housing is subject to fixed adjustment costs but serves as collateral against which individuals can borrow to smooth non-durable consumption.

Beyond providing more credible measures of the quantitative impact of the GSE pricing rule, this model features life-cycle households – and thus allows us to assess the effects of the constant interest rate policy separately for young and middle-aged households. Since housing portfolios differ dramatically by age, the GSE policy should differentially affect households in different parts of their life-cycle. Since our micro data does not include information on household age, calculating distributional effects requires our structural model.

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 vary with local economic conditions in the manner in which they did in 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 and 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 $j = 1, \dots, J$. 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

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

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.

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 $p_{k,t}r^f$.²⁹ We denote owner occupied houses as $h_{i,t}$ and rented houses as $h_{i,t}^f$. Buying or selling an owner occupied house requires paying a fixed cost which is proportional to the current value of the house. That is, the fixed fraction lost for household i takes the following form:

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

Offsetting this disadvantage, 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

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

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

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

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 .

²⁹Note that this specification imposes a constant price-rent ratio.

VI.E Household Problem

The household model is solved recursively. Suppressing the household index i , let $\mathbf{s}_{jk} = (b_j, m_j, h_j, z_j; \gamma_{jk})$ be the state vector for a household. 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

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

when households choose to adjust the size of their owner-occupied house. The value function for

non-adjusters is given by:

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

and a household that chooses to sell its current house³⁰ and rent has value function

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

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

³⁰If 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) We calibrate parameters that do not depend on regional economic activity to standard values from the literature together with standard moments from wealth data. 2) For parameters that vary with regional activity, 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 real 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 $\theta = 0.20$ so that households are required to have a minimum 20% down payment. 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.³¹ 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. (see Kaplan and Violante 2010). Finally, using BEA data, we target a ratio of durable expenditure to housing expenditures of 15. These targets yield $\beta = 0.92$, $r^f = 0.07$ 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 values in Diaz and Luengo-Prado (2010).

³¹Note 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 relatively standard parameters, we must calibrate parameters that vary with regional economic conditions. Our baseline calibration uses local employment as our measure of economic activity. We estimate an annual AR(1) process for log MSA employment³², which yields $\rho_\gamma = 0.947$ and $\sigma_\varepsilon = 0.018$. For simplicity, we assume that local labor earnings move one-for-one with local employment so that $\phi^y = 1$, but we assess the importance of this assumption below. To estimate ϕ_h , we regress log MSA house prices on log MSA employment, which yields $\phi_h = 0.48$. We pick the key policy elasticity ϕ^r so that the regional variation in interest rates in our model when the GSE pricing policy is removed is consistent with that predicted in our data. In particular, we pick ϕ^r so that a two-standard deviation in γ increases mortgage rates by 25 basis points.³³ We provide robustness results for both larger and smaller ϕ^r and discuss alternative counterfactuals in the following section.

VII.C Model Fit

How well does our model fit non-targeted moments? Figure 7 shows the life-cycle profiles in our model compared to the data. Overall the model qualitatively replicates life-cycle patterns in the data. We do a good job of matching the hump-shaped profile of non-durable consumption as well as the increasing homeownership rate across time. We begin our model at age 25 since we do not model schooling decisions, and as a result we underpredict the homeownership rate at age 25, but this is a feature shared by other life-cycle housing models. In addition, the model overpredicts total savings over the lifetime and predicts more borrowing early in life than is observed in the data. This reflects the fact that in the model, households face a deterministic life-cycle profile of income while in the data, these trends are more uncertain (see Guvenen and Smith 2013). Furthermore, in the model households retire deterministically at age 60 and are not able to extend their working life beyond this age. If the household retirement was less sharp, then the level of savings accumulated at retirement would be reduced. Overall, we think the model provides a close enough fit to the data that we are comfortable using it to assess the counterfactual effects of changes in GSE interest rate policy. Our goal is not to provide a precise welfare evaluation of GSEs and is instead to gauge the rough size of their impact on the economy. We think that household reoptimization in response to changing policy is a first order effect that must be modeled in order to get the impact of the GSE

³²We remove permanent differences in employment across MSAs by including MSA fixed effects, and we remove aggregate business cycle effects by including year fixed effects. We include these same fixed effects when calculating the elasticity of house prices to local employment.

³³The implied elasticity of the total borrowing rate ($r + \bar{\Psi}\gamma$) to γ is 0.54.

policy roughly correct. We are less concerned that the modest departures between our model and data will dramatically affect our policy conclusions.

VIII Results

For ease of discussion we label regions with low economic activity and high predicted default “bad” regions and regions with high economic activity and low predicted 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 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 .³⁴ We then solve for λ so that:

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

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 7 shows the implied values of λ for various regions. Clearly the worst regions (on the left

³⁴In 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 lengths leaves the conclusions for working age households unchanged, as the behavior in this final period is essentially irrelevant for working age households.

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 7 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.³⁵ 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 good region, so we account for these differences.

Overall our results imply large transfer payments across regions. Thus, this constant interest rate policy has significant re-distributional consequences across regions. In our benchmark model, the net transfer from good regions to bad regions is nearly \$1400. This one-time transfer is equivalent to an annual transfer of \$120 per household over the households' working life or aggregate annual transfers of almost \$14 billion. For comparison, the Department of Labor forecasts that total unemployment insurance benefits paid in 2014 will equal \$49 billion.³⁶ Thus, the implicit transfers generated by the GSE constant interest rate rule are of a similar order of magnitude to those induced by the unemployment insurance system, but have clearly not obtained similar attention. While this \$120 annual household transfer is quite large, it is important to note that, as predicted, it is substantially smaller than the \$363 transfer implied by the back-of-the-envelope calculation. Ignoring household responses to changing interest rate policy and the less than permanent nature of regional shocks clearly leads to an overstatement of the welfare consequences of constant interest rates.

We focus on the ex-post redistributional consequences of the GSE pricing rule, since this makes our results more comparable to existing studies of fiscal transfers. For example, studies of state transfers arising from the federal income tax system focus on the transfers from states with high income to those with low income rather than on the ex ante consequences of the tax system behind the "veil of ignorance." Similarly, unemployment benefits typically look at their effect on individuals who actually become unemployed rather than their ex ante consequences. Nevertheless, it is straightforward to calculate the ex ante welfare effects of the GSE constant interest rate policy. With concave utility, if the variable interest rate resulted in a pure mean preserving spread

³⁵Total 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.

³⁶See http://www.dol.gov/_sec/media/reports/annual2013/2013annualreport.pdf, page 123

in consumption, it would necessarily lower ex-ante welfare. 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 0.12% so that the constant interest rate policy very mildly increases welfare for the average household in the economy.³⁷ While these ex ante welfare effects are quite small and sensitive to the model parameterization, the ex post transfers induced by the constant interest rate policy are extremely robust.

Why does regional variation in mortgages rates engender ex post transfers? 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 should 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 8 uses our model to explore the age-specific effects of the constant interest rate policy more formally. This table redoes the calculation in Table 7 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.³⁸

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. We can use our model to show that if young households did not have access to housing rental markets, they would be made dramatically worse off by variable interest rates. In particular, Table 9 shows results when

³⁷Note that this number is not the sum of the numbers in row 1 of Table 7 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.

³⁸We 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.

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

Finally, we can explore the sensitivity of our results to various parameter choices. In our benchmark calibration, we pick ϕ_r to match the variation in interest rates observed in the jumbo market. However, if the jumbo market piggybacks off of the GSE policy in picking interest rates, this variation might still be smaller than that in a world with no intervention. Conversely, our empirical section controlled for various observables around the conforming threshold and argued that observable default varies smoothly across the threshold. For this reason, we do not believe sorting on unobservables is a major concern for our empirical estimates, but the presence of sorting could lead our benchmark estimates of ϕ_r to be overstated. Table 10 shows implied transfers under various values of ϕ_r . Unsurprisingly, the level of implied transfers is increasing in ϕ_r , but the transfers remain quite large even when reducing the value of ϕ_r nearly in half.

Table 11 explores how the overall level of regional risk interacts with the GSE constant interest rate rule. That is, does reducing the amount of income or price risk across locations affect the importance of transfers? The second row shows that reducing ϕ_y in half mildly reduces the welfare effects of a constant interest rate, but that they remain quite large. The third row shows that reducing ϕ_h has an even milder dampening effect on the importance of constant interest rates. Thus, we conclude that our results are not particularly sensitive to reasonable changes in the level of regional risk facing households. As long as we continue to generate the same variation in mortgage rates across regions that we observe in the data it is not particularly important whether the regions facing the highest mortgage rates suffer from a 5% decline in employment or a 10% decline in employment.

In our benchmark results, we assume for simplicity that regional economic conditions do not depend on interest rate policy. However, when interest rates rise, local economic activity and house prices should fall. This endogenous feedback from variable interest rates to local economic conditions will in turn affect the level of transfers implied by the GSEs constant interest rate policy. That is, if the constant interest rate policy is abandoned, then interest rates will rise in locations with poor economic conditions. This increase in interest rates will then exacerbate the already poor economic conditions in that location. To assess the importance of this feedback channel, we assume that local

house prices and income decline exogenously when the interest rate increases. This amounts to assuming that ϕ_h and ϕ_y are increasing in ϕ_r . We pick the response of income and house prices to interest rates to match VAR evidence on the response of house prices and GDP to FFR innovations in Christiano, Eichenbaum and Evans (1999) and Vargas-Silva (2008). In particular, we calibrate the feedback so that a 25 basis point increase in interest rates generates a 0.40% decline in house prices and a 0.20% decline in GDP. The fourth row of Table 11 shows that allowing for feedback from interest rates to local economic conditions amplifies the degree of transfers implied by the constant interest rate policy. If the constant interest rate policy was abandoned, the level of regional risk would be amplified and households would be made worse off.

IX Conclusion

Recent business cycles have yielded dramatic disparities in regional outcomes within the United States. While prior research has carefully studied the role of tax and transfer systems in mitigating local shocks, we propose an entirely different mechanism through which federal policy may provide regional redistribution. In this paper we empirically document the extent to which local mortgage rates (do not) vary with local economic conditions. The United States is unique in the extensive role that government institutions play in the mortgage market. 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. If mortgage rates do not respond to local economic shocks, individuals in economically depressed regions may face lower borrowing costs than they would otherwise. Thus, this constant interest rate policy followed by the GSEs results in resources being transferred across regions in state-contingent ways.

The degree to which households can borrow to self-insure against local shocks depends crucially on the interest rate and how it varies with local economic conditions. We establish empirically that, despite large regional variation in predictable default risk, there is essentially no spatial variation in GSE mortgage rates (conditional on borrower observables). In contrast, 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 estimate the magnitude of ex-post redistribution across regions when interest rates are set using a constant

national rate. The difference between the top and bottom outcomes in terms of regional shocks leads to an ex-post redistribution of roughly \$1350 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.

We conclude by noting an important caveat. Throughout, our benchmark for how much mortgage rates should vary with ex-ante default probabilities in the GSE market is the variation we observe in our sample of prime jumbo loans. We feel this is a good comparison group, particularly when we match on factors like MSA, FICO score, documentation type, fixed-rate, and 30-year term. However, we realize that political economy considerations may also limit the extent to which interest rates could vary spatially in private markets. Additionally, discussions with securitizers of private mortgages suggest that they often attempt to use the same mortgage pricing platforms as the GSEs to increase the transparency of their pricing models to the secondary market investors. Both of these factors may lead us to understate the true spatial variation we should observe in the mortgage market with respect to ex-ante differences in local default probabilities. Under this condition, our estimates of the state-contingent transfers across regions will be a lower bound. It is worth noting, however, that using our model, we can easily re-compute how resources are redistributed across regions under alternative assumptions about how mortgage rates would vary with local predicted default probabilities across regions.


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.
- Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, “Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program,” NBER Working Paper 18311, August 2012.
- Asdrubali, Pierfederico, Bent E. Sorensen, and Oved Yosha, “Channels of Interstate Risk Sharing: United States 1963-1990,” *Quarterly Journal of Economics*, 111(4), November 1996.
- Autor, David H., David Dorn, and Gordon H. Hanson, “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *American Economic Review*, 103(6), 2013.
- Beraja, Martin, Erik Hurst, and Juan Ospina, “The Regional Evolution of Prices and Wages During the Great Recession,” University of Chicago Working Paper, 2014.
- Blanchard, Olivier J., and Lawrence F. Katz, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992:1, pp. 1-75, 1992.
- Charles, Kerwin Kofi, Erik Hurst and Matthew J. Notowidigdo, “Masking and Aggregate Employment Changes: Housing Booms and Manufacturing Declines During the 2000s,” University of Chicago Working Paper 2013.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, “Monetary Policy Shocks: What Have We Learned and To What End?” in Taylor, J. and Woodford, M., eds., *Handbook of Macroeconomics*, edition 1, volume 1, pp. 65-148, Elsevier Publishing, 1999.
- 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.
- Farhi, Emmanuel, and Ivan Werning, “Fiscal Unions,” Harvard University Working Paper, March 2013.
- Farhi, Emmanuel, and Ivan Werning, “Labor Mobility within Currency Unions,” Harvard University Working Paper, April 2014.
- Feyrer, James, and Bruce Sacerdote, “Did the Stimulus Stimulate? Real Time Estimates of the Effects of the American Recovery and Reinvestment Act,” NBER Working Paper 16759, February 2011.
- 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.
- Güvenen, 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.
- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vincent Yao, "Mortgage Rates, Household Balance Sheets, and the Real Economy," NBER Working Paper 20561, October 2014.
- Lehnert, Andreas, Wayne Passmore, and Shane M. Sherlund, "GSEs, Mortgage Rates, and Secondary Market Activities," *Journal of Real Estate Finance and Economics*, 36(3), 2008.
- Lustig, Hanno, and Stijn Van Nieuwerburg, "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *Journal of Finance*, 2005.
- Lustig, Hanno, and Stijn Van Nieuwerburgh, "How Much Does Household Collateral Constrain Regional Risk Sharing?" *Review of Economic Dynamics*, April 2010.
- Mayer, Christopher J., Karen Pence, and Shane Sherlund, "The Rise in Mortgage Defaults," *Journal of Economic Perspectives*, 2009.
- Mian, Atif, Kamelesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013.
- Mian, Atif, and Amir Sufi, "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, forthcoming.
- Mian, Atif, Amir Sufi, and Francesco Trebbi, "Foreclosures, House Prices, and the Real Economy," *Journal of Finance*, forthcoming.
- Midrigan, Virgiliu, and Thomas Philippon, "Household Leverage and the Great Recession," NBER Working Paper 16965, April 2011.

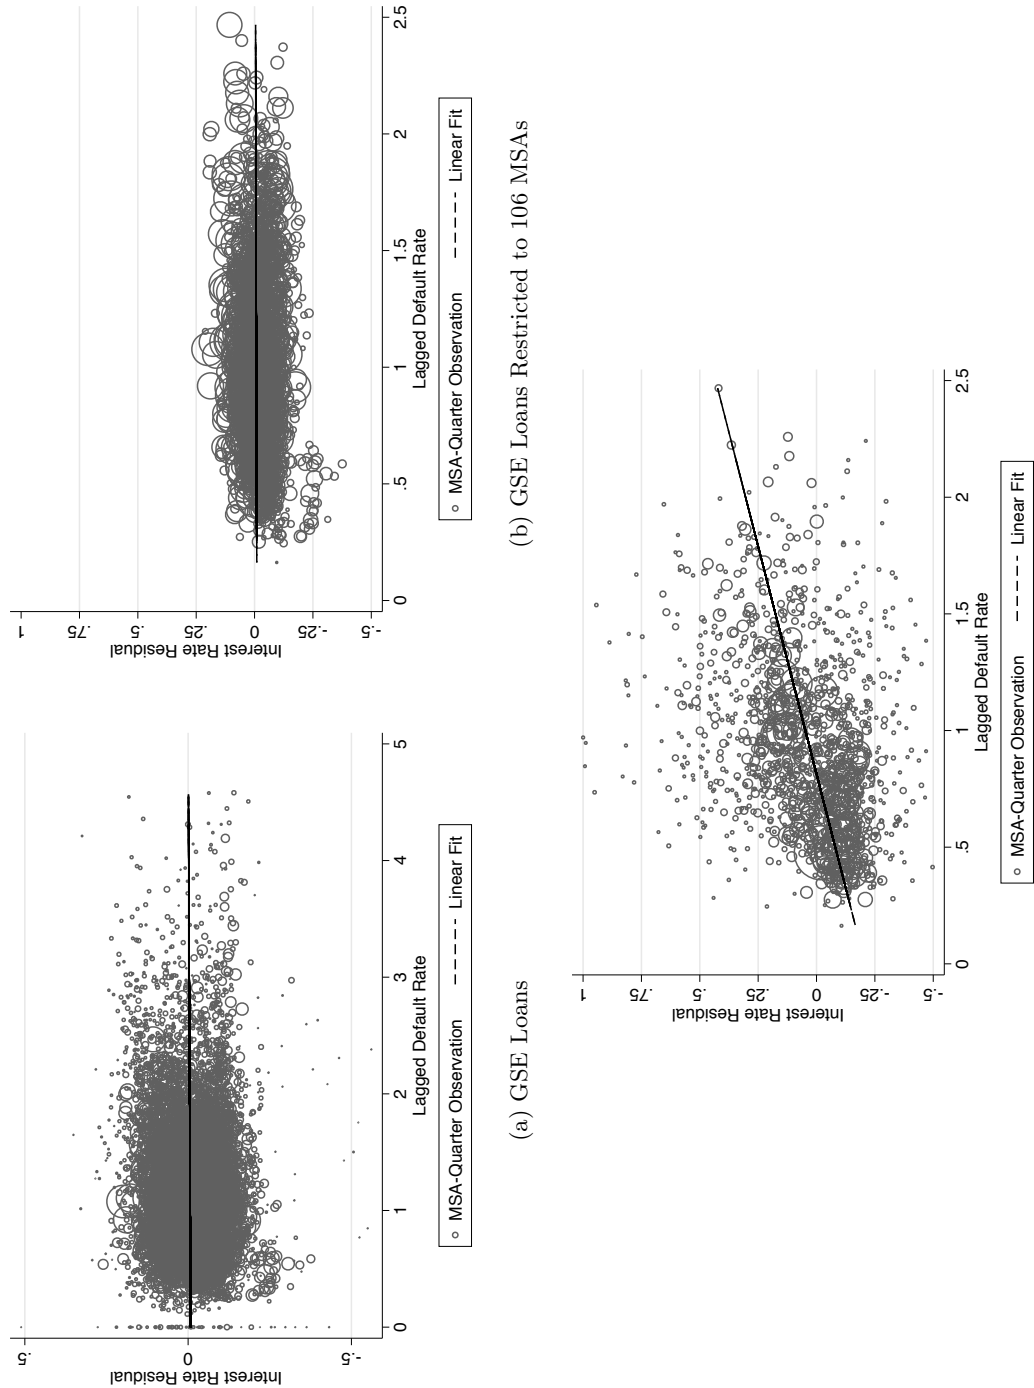
- Nakamura, Emi, and Jon Steinsson, “Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions,” *American Economic Review*, 104(3), March 2014.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel, “Housing, Consumption, and Asset Pricing,” *Journal of Financial Economics*, 83, March 2007.
- Sala-I-Martin, Xavier, and Jeffrey Sachs, “Fiscal Federalism and Optimum Currency Areas: Evidence for Europe from the United States,” NBER Working Paper 3855, October 1991.
- Scharfstein, David, and Adi Sunderam, “Concentration in Mortgage Lending, Refinancing Activity, and Mortgage Rates,” NBER working paper 19156, 2013.
- Sherlund, Shane M., “The Jumbo-Conforming Spread: A Semiparametric Approach,” Federal Reserve Board Finance and Economics Discussion Series 2008-01, January 2008.
- Shoag, Daniel, “The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plan Returns,” Harvard University Working Paper, 2014.
- Stroebel, Johannes, and Joseph Vavra, “Housing Wealth Effects and Retail Prices,” University of Chicago Working Paper, 2014.
- Suarez-Serrato, Juan Carlos, and Philippe Wingender, “Estimating Local Fiscal Multipliers,” University of California at Berkeley mimeo, 2010.
- Vargas-Silva, Carlos, “Monetary Policy and the U.S. Housing Market: A VAR Analysis Imposing Sign Restrictions,” *Journal of Macroeconomics*, 30(3), September 2008.
- Yagan, Danny, “Moving to Opportunity? Migratory Insurance Over the Great Recession,” Harvard University mimeo, January 2014.

Figure 1: MBS Issuance by GSE / Non-GSE Issuer, 1996-2013



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


Figure 2: Relationship between Interest Rates and Lagged Local Default, 2001-2006



(c) Non-GSE Loans


Note: Figure shows the relationship between residualized interest rates and lagged MSA-level default of loans originated within the last two years for three samples. 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 and the Loan Performance Database.

Figure 3: Average FICO Score, by loan amount, 2001-2006



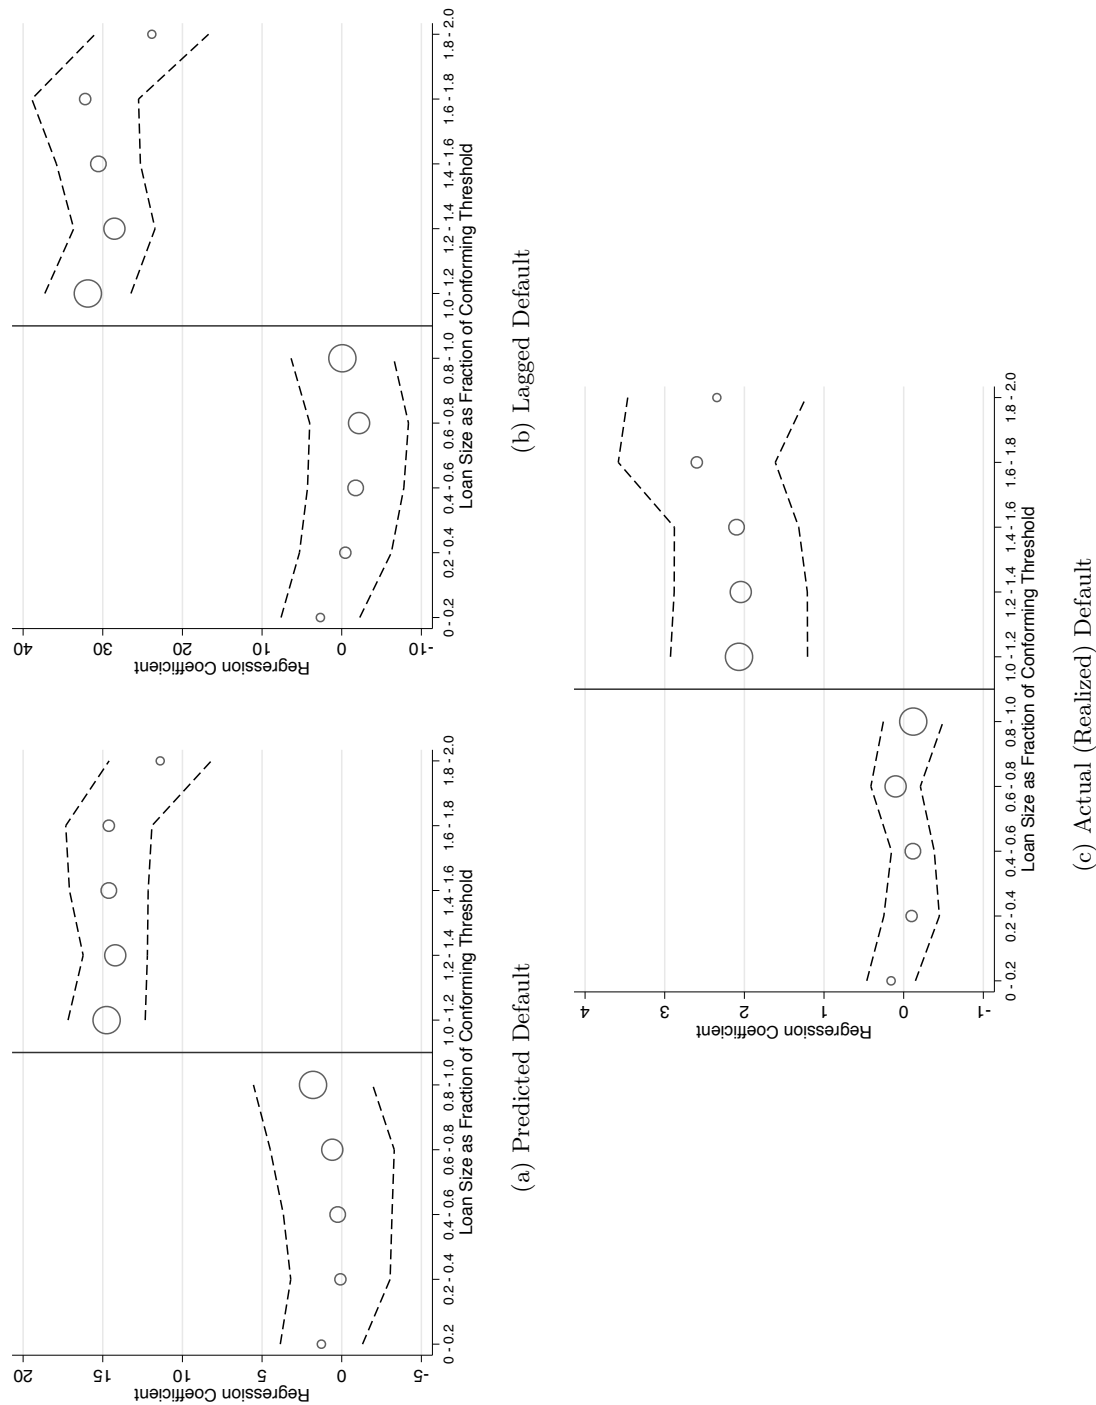
Note: Figure shows the average FICO credit score in each loan amount bin around the conforming loan limit. To the left of the limit (values ≤ 1), loans are insured and securitized by the GSEs. To the right of the limit (values ≥ 1), loans are securitized by the private non-GSE market. The GSE sample is restricted to the MSAs where non-GSE loans are present, and matched based on the FICO distribution of non-GSE loans for comparability. See text for details. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and the Loan Performance Database.

Figure 4: Default rates, by Loan Amount




Note: Figure shows the average residualized default rate in each loan amount bin around the conforming loan limit. Adjusted residual removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. To the left of the limit (values ≤ 1), loans are insured and securitized by the GSEs. To the right of the limit (values ≥ 1), loans are securitized by the private non-GSE market. The GSE sample is restricted to the MSAs where non-GSE loans are present, and matched based on the FICO distribution of non-GSE loans for comparability. See text for details. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and the Loan Performance Database.

Figure 5: Relationship between Interest Rates and Three Measures of Default, 2001-2006




Note: Figure shows the relationship between residualized interest rates and default rates in each loan amount bin around the conforming loan limit. Adjusted residual removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. To the left of the limit (values ≤ 1), loans are insured and securitized by the GSEs. To the right of the limit (values ≥ 1), loans are securitized by the private non-GSE market. The GSE sample is restricted to the MSAs where non-GSE loans are present, and matched based on the FICO distribution of non-GSE loans for comparability. See text for details. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and the Loan Performance Database.

Figure 6: Relationship between Interest Rates and Lagged Default Rates, Conditioning on Prepayment Risk, 2001-2006



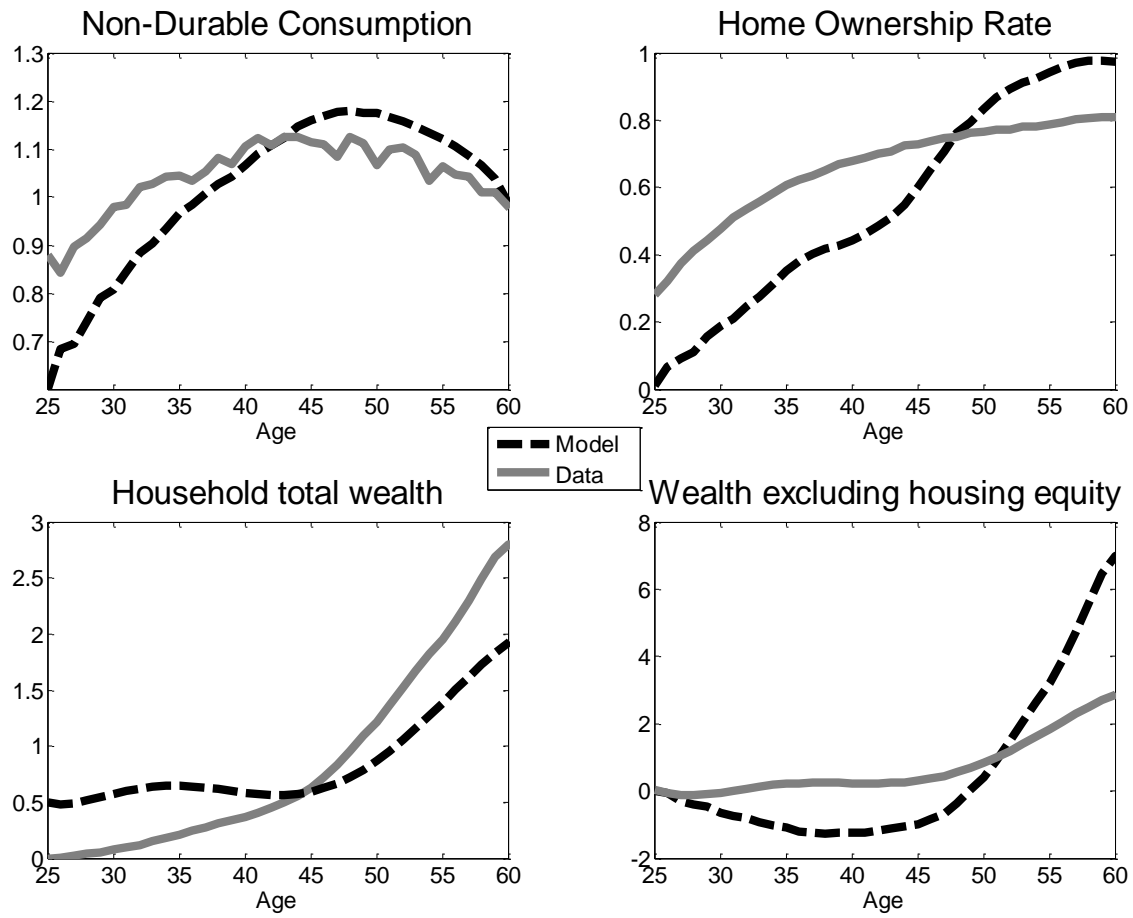
(a) GSE Loans



(b) Non-GSE Loans

Note: Figure shows the relationship between residualized interest rates and lagged default rates conditional on lagged prepayment risk for loans originated between 2001 and 2006. The GSE sample is restricted to the MSAs where non-GSE loans are present. Source: Authors' calculations using Fannie Mae and Freddie Mac Single-Family Loan Databases and the Loan Performance Database.

Figure 7: Average Life Cycle Profiles: Model Simulation vs. Data



Data sources: Non-durable consumption comes from Aguiar and Hurst (2013). Home ownership rates are calculated from the March CPS, and wealth statistics are calculated from PSID data. See text for additional details.

Table 1: Descriptive Statistics

	Time Period					
	2001-2006			2007-2009		
	GSE All	GSE Restricted MSAs	GSE Matched Sample	Prime Jumbo	GSE All	GSE Restricted MSAs
Number of Loans	13,110,212	8,052,967	70,560	70,680	4,861,259	3,677,984
Median FICO	728	727	657	656	756	757
Median LTV	0.78	0.75	0.77	0.80	0.76	0.75
MSAs covered	374	106	106	106	374	106
Mean Interest Rate (%)	6.25	6.22	6.30	6.66	5.65	5.63
Mean 2-Yr Delinquency Rate (%)	1.6	1.4	2.9	2.1	3.8	4.0
<u>Cross MSA SD of Interest Rates</u>						
Unconditional (percentage point)	0.544	0.557	0.568	0.657	0.627	0.623
Conditional (percentage point)	0.076	0.072	0.086	0.165	0.070	0.064
<u>Cross MSA SD of Delinquency Rates</u>						
Unconditional (percentage points)	1.5	1.2	3.1	2.7	4.0	4.3
Conditional (percentage points)	1.3	1.1	2.7	2.5	2.9	2.9

Note: The table provides summary statistics for the samples of GSE and non-GSE loans. The different columns refer to different samples and different time periods. The “Restricted MSA” sample uses only those MSA with prime jumbo loans present. The “GSE Matched Sample” restricts to these 106 MSAs and matches the distribution of FICO scores in the non-GSE sample. See text for details. 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 and the LoanPerformance Database.

Table 2: Responsiveness of Conditional MSA Interest Rates to Lagged GSE Default Rates

	Time Period					
	2001-2006			2007-2009		
	(1)	(2)	(3)	(4)	(5)	(6)
	GSE All	GSE Restricted MSAs	GSE Matched Sample	Prime Jumbo	GSE All	GSE Restricted MSAs
Coefficient on Lagged GSE Default Rate	0.16 (0.29)	2.40 (2.84)	3.38 (2.72)	30.55 (2.49)	1.12 (0.23)	1.09 (0.27)
Implied Basis Point Change in Mortgage Rate to a Two Standard Deviation Change in Lagged GSE Default	0.28	1.78	2.43	22.61	3.18	3.27
Sample Size	13,109,968	8,052,967	70,560	70,327	4,861,218	3,677,984

Note: The table shows the coefficient from a regression of conditional MSA interest rates during a given quarter on lagged GSE default rates. The different columns refer to different samples and different time periods for which the conditional MSA interest rates and lagged default rates are based. The different sample definitions are discussed in the notes to Table 1. The implied change in interest rate to a one standard deviation change in lagged GSE default is simply the coefficient times the standard deviation of lagged GSE default across the MSAs in the relevant sample. Standard errors in parentheses clustered at the MSA level.

Table 3: Relationship Between Conditional MSA Interest Rates and MSA Predicted Defaults, 2001-2006

	Base Specification			Regression Discontinuity Specification	
	(1) GSE Matched Sample	(2) Prime Jumbo Sample	(3) Difference in Coefficients	(4) p-value of difference	(5) RD Coefficient
Predictive Default Measure					
Predicted Default Using Lagged Local GSE Default	1.94 (1.56)	11.94 (0.97)	10.0	<0.001	12.96 (2.53)
Lagged Default (Random Walk)	3.38 (2.72)	12.60 (3.16)	9.22	<0.001	12.46 (4.56)
Actual Default (Perfect Foresight)	0.15 (0.13)	2.12 (0.40)	1.97	<0.001	2.19 (0.46)
Underlying Sample Size of Loans	70,560	70,327			
Time, FICO, and LTV Controls Included	Yes	Yes			

Note: This table presents coefficients from regressions of conditional MSA interest rates on three measures of predictive default: lagged default rates, actual default rates, and predicted default rates. Lagged default is measured within-sample depending on GSE or non-GSE loans. The sample of GSE loans is restricted to the 106 MSAs where non-GSE loans are present during the time period 2001-2006 and matches the distribution of FICO scores in the non-GSE sample. The different sample definitions are discussed in the notes to Table 1. Standard errors in parentheses clustered at the MSA level.

Table 4: Robustness of Regression Discontinuity Estimates

Predictive Default Measure	Specification		
	(1)	(2)	(3)
Predicted Default Using Lagged Local GSE Default	12.96 (2.53)	12.50 (2.83)	11.63 (2.70)
Time, FICO, and LTV Controls Included	Yes	Yes	Yes
Predicted Payment Controls Included	No	Yes	No
MSA Fixed Effects Included	No	No	Yes

Note: This table presents regression-discontinuity style estimates of the difference in the relationship between interest rates and predicted defaults around the conforming loan limit. Lagged default and lagged prepayment measures are constructed within-sample depending on GSE or non-GSE loans. See text for details. The GSE sample is restricted to the MSAs where non-GSE loans are present and matched on the FICO distribution of the non-GSE sample for better comparability. Each coefficient represents a separate regression. Standard errors in parentheses clustered at the MSA level.

Table 5: Standard Deviations of Predicted Default

Predicted Default Measure	Time Period		
	2001-2006		2007-2009
	GSE Matched Sample	Prime Jumbo Sample	GSE Restricted MSAs
Predicted Default Using Lagged Local GSE Default	0.006	0.009	0.011
Lagged Default (Random Walk)	0.004	0.005	0.015
Actual Default (Perfect Foresight)	0.031	0.027	0.043

Note: This table presents the standard deviation of each measure of predicted default for each sample used in the analysis. See text for details of sample construction.

Table 6: Predicted Counterfactual Two Standard Deviation Cross MSA Variation in GSE in Interest Rates

Predicted Default Measure	Time Period	
	2001-2006	2007-2009
Predicted Default Using Lagged Local GSE Default	0.156	0.285
Lagged Default (Random Walk)	0.100	0.374
Actual Default (Perfect Foresight)	0.136	0.188

This table presents the interest rate response to a two standard deviation change in each predicted default measure for two time periods. These values are obtained by multiplying the values in Table 4, column 1 with the standard deviations found in table 5 for GSE loans.

Table 7: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates

	Regional Employment				
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviations
Percent consumption gain ($\lambda * 100$)	1.14%	0.60%	0.08%	-0.42%	-0.96%
Dollar per household effect $(\lambda \times 63790 \times \frac{\bar{C}_{region}}{C_{overall}})$	\$670	\$363	\$51	-\$281	-\$681

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 employment rate. See text for details.

Table 8: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates, by Age

	Regional Employment				
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviations
Percent consumption gain ($\lambda * 100$): Overall	1.14%	0.60%	0.08%	-0.42%	-0.96%
Percent consumption gain ($\lambda * 100$): Young	0.62%	0.30%	0.00%	-0.38%	-0.90%
Percent consumption gain ($\lambda * 100$): Middle Aged	1.38%	0.76%	0.14%	-0.44%	-1.00%

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

Table 9: One-Time Consumption Equivalent Necessary to Accept Region-Specific Rates, No Rental Market, by Age

	Regional Employment			
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation
Percent consumption gain ($\lambda * 100$): Overall	2.02%	1.10%	0.08%	-1.46%
Percent consumption gain ($\lambda * 100$): Young	3.14%	1.70%	0.12%	-2.72%
Percent consumption gain ($\lambda * 100$): Middle Age	1.44%	0.72%	0.06%	-0.80%

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

Table 10: Sensitivity to Different Values of ϕ_r

	Regional Employment			
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation +2 Standard Deviations
Percent consumption gain ($\lambda \cdot 100$): Benchmark (25 bp)	1.14%	0.60%	0.08%	-0.42% -0.96%
Larger Variation (35 bp)	1.56%	0.86%	0.14%	-0.58% -1.44%
Smaller Variation (15 bp)	0.60%	0.32%	0.04%	-0.24% -0.44%

Table 11: Sensitivity to Other Elasticities

	Regional Employment				
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviation
Percent consumption gain (λ * 100): Benchmark	1.14%	0.60%	0.08%	-0.42%	-0.96%
Reduce Regional Income Variation in Half	0.98%	0.52%	0.04%	-0.44%	-0.86%
Reduce Regional House Price Variation in Half	1.16%	0.64%	0.10%	-0.30%	-0.80%
Allow for Endogenous Feedback	2.20%	1.16%	0.08%	-0.94%	-1.94%