

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

February 2015

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 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 postulate that the lack of regional variation in GSE mortgage rates is likely driven by political pressure. We quantify the economic impact of the national 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 national interest rate policy has significant ex-post redistributive consequences across regions.

*First draft: April 2014. We thank Sumit Agarwal, Heitor Almeida, John Leahy, Arvind Krishnamurthy, Stijn Van Nieuwerburgh, Monika Piazzesi, David Scharfstein, Adi Sunderam and seminar participants at the Bank of Italy, HEC, Indian School of Business, Kellogg, MIT, National University of Singapore, Ohio State, Rutgers, Stanford, Toronto, UBC, University of Chicago Booth, University of Chicago Harris, University of Illinois, University of Michigan, Wharton, the FRIC 2014 conference on financial frictions, and the NBER conference on Financing Housing Capital 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 United States 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 United States, the vast majority of such borrowing occurs through the mortgage market. In this paper we empirically document the extent to which local mortgage rates vary with local economic conditions. Government-sponsored 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, then individuals in those regions may face lower borrowing costs than if this default risk were priced into interest rates. Lower borrowing costs may in turn help to offset the negative local economic shock that increased local default probabilities. Thus, the constant interest rate “policy” followed by the GSEs results in resources being transferred across regions in state-contingent ways. Our objective is to quantify the size and welfare consequences of these implicit transfers. The extent of such redistribution can then be compared with the costs of providing such insurance through implicit subsidies to the GSEs, including too-big-to-fail subsidies, and can inform the debate on the costs and benefits of the GSEs.²

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). These findings hold throughout the entire 2000s and are not limited to the period surrounding the 2008 recession. Despite this finding, we further document that interest rates on loans securitized by the GSEs do not vary at all with predictable loan default risk. 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

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.

²Importantly, although we argue that the GSE national interest rate policy has important redistributive effects, alternative tax-and-transfer policies might be able to implement similar ex-post transfers without some of the costs induced by GSEs. Our paper is a positive analysis of the consequences of existing interest rate variation rather than a normative analysis of optimal policy.

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, given the large spatial variation in default risk. To do this, we exploit loan-level data containing 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 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 loan-to-value [LTV] ratio). Unlike the interest rate on GSE loans, we document that the interest rate on prime jumbo loans rises dramatically with ex-ante local predicted default risk. Thus, although there is no regional risk-based pricing in the government-backed GSE market, the private market does set interest rates based in part on regional risk factors.

Employing a variety of econometric techniques, including a regression discontinuity approach around the conforming limit threshold, we construct counterfactual estimates of the extent to which 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. These results are robust to controlling for many potentially confounding factors, including the possibility that prepayment propensities or points and fees vary spatially. We also document that there is no differential response in loan amounts between the GSE and prime jumbo borrowers to ex-ante predictable default. This suggests that the GSE market is not compensating for the lack of spatial variation in mortgage rates by reducing the amount of credit extended relative to the prime jumbo market.

In the second part of the paper, we explore a number of explanations for why the relationship between mortgage rates and predictable default differs in 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 by 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. This 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.

In the final part of the paper, we quantify the economic impact of the transfers induced by the GSEs’ constant interest rate policy. We begin with a back-of-the-envelope calculation that simply takes the existing portfolio of mortgages in the U.S. as given and computes how households’ payments would change if GSEs priced risk like the private market. However, this back-of-the-envelope calculation is likely to overstate the true effects of the policy because it is inherently static in nature and does not account for household reoptimization 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. Furthermore, some regions that 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. Capturing these dynamic considerations and accurately accounting for the endogenous response of households to changes in mortgage policy requires the use of a quantitative structural model that can more rigorously assess the economic impact of the GSEs' constant interest rate policy.

To address these issues, we build a 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 nondurable consumption. The model's consumption equivalents account for both reoptimization on the part of individuals and the persistence of 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 that accounts for household reoptimization, the GSE constant interest rate policy generates transfers across regions that are substantially smaller than the back-of-the-envelope calculation but still nontrivial. In our benchmark calibration, designed to match the regional variation observed during the Great Recession, the GSE pricing policy generates a one-time \$1,000 per-household tax on a region with a two-standard-deviation increase in regional activity (decline in predicted mortgage default) and generates a one-time subsidy of \$800 for a region with a two-standard-deviation decrease in regional activity (increase in predicted mortgage default). According to our model, about \$20.7 billion was transferred via the mortgage market from regions receiving better than average economic shocks to regions receiving worse than average economic shocks. For comparison, the Department of Labor forecasts that total unemployment insurance benefits paid in 2014 would equal \$49 billion. As an additional comparison, the one-time \$1,800 per-household transfer from regions with two-standard-deviation positive shocks to those with two-standard-deviation negative shocks is larger than the per-household tax rebate checks paid by the U.S. government during the 2001 and 2008 recessions. Thus, our results suggest that the magnitude of redistribution induced by the GSEs through the mortgage market is economically meaningful.

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 housing rental markets, then they are affected quite dramatically by the GSEs' constant interest rate policy because they are the most likely to

want to borrow during regional downturns.

Our work relates to a number of existing literatures. First, there is a small body of work that studies 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. 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 against local shocks. Lustig and Van Nieuwerburgh find that the extent of regional risk sharing varies with the state of the aggregate housing market. Our paper complements these findings by highlighting a direct mechanism by which the credit market serves to insure regional shocks. Our results suggest that GSEs have a large effect on the extent to which credit markets insure against regional risk, since GSE 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.³

Although not the primary focus of our paper, our work also 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. These critics point to the many ways that the GSEs distort the allocation of capital within the economy.⁴ Proponents, on the other hand, argue that the GSEs serve parts of the 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. Although Fannie Mae and Freddie Mac have a variety of well-documented impacts on housing and mortgage markets, our paper shows that their common national interest rate policy may have additional important and understudied consequences. However, our paper is silent on the fact that the implicit subsidy to the GSEs may distort the allocation of capital toward the housing market and away from other productive resources. The goal of our paper is to shed light on the empirical extent to which interest rate variation alters regional risk, not to provide a full evaluation of the costs and benefits of GSEs. With this goal in mind, we take the GSEs as given and explore the consequences of their policies for regional risk sharing rather than pursuing a normative analysis of optimal policy. 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.

³More broadly, 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 Nieuwerburgh (2005), Mian, Sufi, and Trebbi (2014), Mian, Rao, and Sufi (2014), Mayer et al. (2009), Piazzesi et al (2007), and Scharfstein and Sunderam (2013).

⁴See, for example, Calomiris (2001), Lucas and McDonald (2010), Lucas and Torregrosa (2010), and Acharya et al. (2011).

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 to 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).⁵ 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 where the guarantee fee could potentially vary with features of the borrower (FICO score) or loan (loan-to-value ratio). The interest rate charged on mortgages sold to the GSEs thus reflects the guarantee fee, additional guidelines imposed by the GSEs, and any other charges that could potentially vary with regional risk.

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 or interest on the loans underlying the securities. That is, 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 interest rate charged on mortgages sold through the private market thus reflects the guidelines imposed by investors, as well as other charges that could potentially vary with regional risk.

Figure 1 shows the share of mortgages in the secondary market that were securitized or directly held by the GSEs relative to the private market during the 2000s. Prior to 2004, roughly 80 percent of the securitized mortgage market was securitized by the GSEs (Fannie Mae, Freddie Mac, and Ginnie Mae).⁶ The private market securitized all other loans. The private market includes jumbo mortgages (loans that exceed the conventional mortgage size limits), subprime mortgages (loans for borrowers with poor credit histories), and Alt-A mortgages (loans for borrowers who provide less than full documentation). 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 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 reduces the GSEs' cost of funds relative to the private market. Estimates suggest that mortgage rates for conventional mortgages are between 20 to 40 basis points lower than mortgage rates for otherwise similar jumbo mortgages (see, for example, Sherlund 2008). This difference is attributed

⁵Specifically, 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.

⁶The data in Figure 1 come from data published by the Securities Industry and Financial Markets Association (SIFMA). With respect to the GSEs, Fannie Mae and Freddie Mac securitize conventional mortgages while Ginnie Mae securitizes mortgages issued by the FHA.

to both the implicit guarantee and the scale of the GSE market.⁷ This cost differential makes it difficult for the private market to undo pricing mistakes made by the GSEs. If political constraints prevent the GSEs from 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. However, this cost differential does provide a bound on the potential mispricing of local risk.

As discussed in the introduction, we argue that the difference in regional risk-based pricing between the GSE and private market is driven by political constraints on GSEs. It is natural to ask if alternative mechanisms could also explain our result. For instance, the GSE market is much bigger than the private market, which could contribute to its cost advantage. However, this cost advantage occurs in all regions and so cannot explain the lack of regional variation in interest rates. In addition, one might also worry that the GSE constant interest rate policy occurs because GSE loans are securitized, which allows for better diversification of idiosyncratic and regional risk. However, note that our comparison will be with loans in the private market that are also pooled together and securitized; Hence 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 GSE market, loans in the private 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 mispricing. Although institutional investors may hold both GSE-backed and private mortgage-backed securities, only the private securities 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 mispricing. From the investors' perspective, they only face the risk of early prepayment in GSE-backed mortgage securities. When the GSEs were publicly traded, their 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. Due to issues related to data coverage and comparability, we do not analyze loans securitized by Ginnie Mae. The second includes a sample of jumbo loans securitized by the private market.

⁷For a recent discussion of this literature, see Lehnert et al. (2008).

III.A Fannie Mae/Freddie Mac Sample

Our primary data sources are 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 2012.⁸ The data include 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 are provided monthly and include 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 is pre-paid. 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 either GSE. Finally, within our analysis sample, we include both loans associated with a new purchase mortgage or a refinancing.⁹ In total, our sample includes roughly 13 million loans that were originated during the 2001-2006 period and another roughly 5 million loans that were originated during the 2007-2009 period.

III.B Prime Jumbo Sample

Our second primary data source is 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 focus only on what we term fixed-rate “prime jumbo” mortgages. As noted above, loans securitized by the private market include both subprime 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 as similar as possible to the mortgages in the Fannie/Freddie pool. To do that, our “prime jumbo” mortgages: (1) have an origination value that is between the conforming mortgage limit and two times the conforming mortgage limit in the year of origination, (2) have a fixed interest rate, (3) have a LTV ratio at origination of less than 100 percent, (4) have a FICO score at origination of 620 or higher, (5) provide full documentation at the time of origination, and (6) were

⁸The 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 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.

⁹The results are unchanged if we analyze Fannie Mae and Freddie Mac loans separately, or if we exclude refinance loans. The data appendix discusses other sample restrictions as well. In particular, we include only mortgages that have a FICO score at origination of at least 620 (the bulk of GSE data), were originated between January 2001 and December of 2009, and were originated within one of our included MSAs.

originated between 2001 and 2006.¹⁰ The 2006 end date is necessitated by the fact that the private market essentially dried up in 2007 (as seen in Figure 1). 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 that very expensive loans may differ along many dimensions of loan and borrower characteristics compared with 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 include only observations where there are at least five 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 ensures that there will be a minimum amount of loans for each MSA-quarter cell. We explore the robustness of this restriction in our work below. In total, our “prime jumbo” sample includes 70,327 loans originated during the 2001-2006 period.

III.C 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 (where at least five loans that meet our definition were originated during a quarter). 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 include only loans for the 106 MSAs where we have at least five observations of prime jumbo data. This ensures that the MSA-quarter coverage between the two samples is identical. The restriction 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

¹⁰The conforming limit was raised from \$275,000 to \$417,000 between 2001 and 2006. This period pre-dates the FHFA policy to vary loan limits regionally based on “high cost” areas, which began in 2008.

sample along two additional dimensions. First, we restrict the sample so that the sample sizes match exactly. 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. Second, we restrict the GSE sample so that it replicates the FICO and LTV distributions of the prime jumbo sample. As a result, the distribution of borrower quality as measured by FICO scores and LTV ratios will not differ between the two samples.¹¹ We refer to this sample as the “matched” GSE sample where the matching occurs on MSA-quarter, FICO score, LTV ratio, and sample size. For each prime jumbo loan we “draw” a similar loan from the GSE sample. 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, median LTV variation, and the MSA coverage match exactly with the prime jumbo sample. 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 33 basis points lower than the rate on prime jumbo loans (6.33% 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, 3.0% 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.D Controlling for Borrower and Loan Characteristics

Throughout the paper, we want to examine the spatial variation in mortgage rates and show how the variation correlates with the spatial variation in predicted future mortgage default rates. One reason interest rates and 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 spatial variation in mortgage rates and default rates. Of course, matching the two samples on FICO scores and LTV ratios 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.

¹¹All 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/LTV distributions. The results are nearly identical across the specifications. See the appendix that accompanies the paper for details of the exact selection criteria for our main sample to facilitate replication of our results. In Appendix Table A1, we show that the matching criteria resulted in both the mean and distribution of FICO and LTV being similar between the GSE and prime jumbo sample.

To formally illustrate these patterns, we purge the variation in mortgage rates and subsequent 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:

$$r_{ikt}^j = \alpha_0^j + \alpha_1^j X_{it} + \alpha_2^j D_t + \alpha_3^j D_t \cdot X_{it} + \eta_{ikt}^j$$

$$y_{ikt}^j = \varphi_0^j + \varphi_1^j X_{it} + \varphi_2^j D_t + \varphi_3^j D_t \cdot X_{it} + \nu_{ikt}^j$$

where r_{ikt}^j is the loan-level mortgage rate for a loan made to borrower i , in MSA k , during period t and y_{ikt}^j is an indicator variable for whether the loan made by borrower i , in MSA k , 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 is fully interacted with quarter of origination dummies. The goal of these specifications is to recover η_{ikt}^j and ν_{ikt}^j , the residual mortgage rate and residual ex-post delinquency rate, respectively, for borrower i in MSA k 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_{kt}^j , and location specific average ex-post default rates, Y_{kt}^j . We do this separately for each time period and for each sample. Specifically,

$$R_{kt}^j = \frac{1}{N_{kt}^j} \sum_{i=1}^{N_{kt}^j} \eta_{ikt}^j$$

$$Y_{kt}^j = \frac{1}{N_{kt}^j} \sum_{i=1}^{N_{kt}^j} \nu_{ikt}^j$$

where N_{kt}^j is the number of loans in the MSA k during quarter*year t within each sample. Formally, R_{kt}^j (Y_{kt}^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_{ikt} 's) across MSAs and the standard deviation of conditional mortgage rates (R_{kt}^j 's) across the MSAs for our matched GSE sample and our prime jumbo sample originated during 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 we condition 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 as in the matched GSE sample, while the conditional

cross-MSA standard deviation of delinquency rates is similar in the two samples. This shows that there is more cross-MSA variation in mortgage rates in privately securitized loans than in GSE loans.

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 prime jumbo rates do vary with this risk.

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 k 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 $t - 1$. Because our time unit of analysis is a quarter, our lagged delinquency measure is the fraction of all loans originated between 9 quarters prior and 1 quarter prior that became 60 days delinquent by the current quarter.¹² We refer to this measure as $E_{k,t-1}^j$, where $E_{k,t-1}$ denotes lagged economic activity in location k 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.¹³

To present the data, Figure 2a shows a simple scatter plot of local mortgage rates residuals for the GSE loans, R_{kt}^{GSE} , in the full GSE sample against lagged local GSE default rates, $E_{k,t-1}^{GSE}$, during the 2001-2006 period. Figure 2b presents the same result for the GSE sample matched on the distribution of FICO scores and LTV ratios. The matched GSE sample, as discussed above, only includes 106 MSAs, while the full sample includes 374 MSAs. Figure 2c analogously shows the scatter plot of local mortgage rates residuals for the prime jumbo loans, R_{kt}^{jumbo} , against lagged local GSE default rates, $E_{k,t-1}^{GSE}$, during the same time period. Each observation in the figures is an MSA-quarter pair.

Figures 2a and 2b show that there is no relationship between lagged local GSE default rates and average local mortgage rates in either the full GSE sample or in the matched GSE sample. Columns (1) and (3) of Table 2

¹²We have experimented with 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.

¹³We 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.

summarize the regression line of the scatter plots in Figures 2a and 2b, respectively. Focusing on the results from column (3) of Table 2, a one-percentage-point increase in lagged GSE default is associated with a (statistically insignificant) increase in local GSE mortgage rates of only 3.5 basis points (i.e., from 6.000 to 6.035).¹⁴ The standard deviation of lagged GSE default across MSAs is 0.7 percentage points, which implies that a one-standard-deviation increase in lagged default is associated with only a 2.5 basis point increase in local GSE mortgage rates. Even adjusting for the standard error of the estimate, this estimate is essentially a precise zero. As seen from comparing the first three columns of Table 2, there is no economically meaningful or statistically significant relationship between lagged GSE default and GSE mortgage rates regardless of the sample used for the GSE data. 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 economically meaningful relationship between lagged local mortgage default and local mortgage rates.

The pattern in Figure 2c is in stark contrast to those in Figures 2a and 2b. Figure 2c shows that there is a strong positive correlation between lagged GSE default rates and local interest rates for prime jumbo loans. MSAs that 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 coefficient is 10 times larger than the effect on GSE mortgage rates and is highly statistically significant.

IV.B Relationship Between Predicted Default and Mortgage Rates

The previous subsection showed the relationship between lagged economic conditions and current mortgage rates. What lenders are presumably interested in is how past economic conditions translate into future default risk. In this subsection, we assess the extent to which lagged local economic conditions predict subsequent actual default. We then assess the cross-region relationship between predicted default and mortgage rates for both the GSE and prime jumbo samples.

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 relationship between future default and lagged default conditional on borrower and loan characteristics. In particular, we run the following regression on both the GSE and prime jumbo samples using data from 2001-2006:

$$y_{ikt}^j = \theta_0^j + \theta_1^j X_{it} + \theta_2^j D_t + \theta_3^j D_t \cdot X_{it} + \lambda^j E_{k,t-1}^{GSE} + \nu_{ikt}^j$$

where y_{ikt}^j , X_{it} , D_t and $E_{k,t-1}^{GSE}$ are defined above. The goal of this regression is to use the underlying micro data to see whether lagged GSE default rates predict subsequent mortgage default (conditional on loan and borrower

¹⁴When fitting a line through the scatter plot or running regressions, 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.

observables). We use the lagged GSE default rate for both samples so that we capture the response of actual default rates in the two samples to the same underlying economic conditions. The primary coefficient of interest is λ^j , which we can use to define our first measure of predicted local mortgage default:

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

For both samples, λ^j is large and statistically significant, showing that lagged GSE default rates have significant predictive power for future default rates in both the GSE and prime jumbo samples. In particular, for the GSE market, the coefficient is 1.71 (SE=0.24, F-stat=50.5), while for the non-GSE market, the coefficient is 2.55 (SE=0.31, F-stat=68.1).¹⁵ For robustness, we also explore two additional measures of predicted local default. The first we refer to as our “random walk” forecast such that:

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

This specification implies 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. As with the results, lagged prime jumbo default rates were highly predictive of future prime jumbo default rates.

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

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

This “perfect foresight” specification implies that lenders’ best prediction of future default in a given sample in a given location (conditional on observables) is the actual future default rate (which we label $Y_{k,t}^j$ in the above specification).

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:

$$r_{ikt}^j = \omega_0^j + \omega_1^j X_{it} + \omega_2^j D_t + \omega_3^j D_t \cdot X_{it} + \beta^j \hat{Y}_{kt}^j + \eta_{ist}^j$$

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} (estimated from separate regressions on the GSE data and prime jumbo data, respectively). To address concerns related to statistical inference with generated

¹⁵One 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 various subperiods of our data. For example, within the GSE sample, λ^j was large, statistically significant, and of similar order of magnitude during the 2001-2003 period, the 2004-2006 period, and the 2007-2009 period. In all three subperiods, lagged GSE default positively and significantly predicted future default rates within each loan type.

regressors, every estimate reported in the paper that relies on predicted defaults uses bootstrapped standard errors (500 repetitions, clustered at the MSA level). 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.

In all cases, mortgage rates in the prime jumbo market respond much more to predicted default than do mortgage rates in the GSE sample. That is, these regressions show that the greater response of jumbo mortgage rates to lagged economic conditions is not driven by greater sensitivity of actual default to these conditions. Furthermore, it is not just that jumbo rates are more responsive than GSE rates: our regression shows that GSE interest rates do not respond in *any* meaningful way to predicted default. A one-percentage-point increase in local predicted default only raises local GSE mortgage rates by 2 basis points, an effect that is statistically indistinguishable from zero.¹⁶

We can also explore the differential responsiveness of local mortgage rates to measures of local predicted default using a regression discontinuity approach to estimate ($\beta^{jumbo} - \beta^{GSE}$) around the conforming loan threshold. Specifically, we estimate:

$$r_{ikt}^j = \delta_0 + \delta_1 X_{it} + \delta_2 D_t + \delta_3 D_t \cdot X_{it} + (\tilde{\delta}_1 X_{it} + \tilde{\delta}_2 D_t + \tilde{\delta}_3 D_t \cdot X_{it}) D_{it}^{jumbo} + \delta_4 Bin_{it} + \beta Bin_{it} \cdot \hat{Y}_{kt}^j + \eta_{ist}^j$$

For this regression, we pool together the prime jumbo sample and the matched GSE sample for the years 2001-2006. D^{jumbo} is a dummy variable indicating that the loan is from the prime jumbo sample, and our specification allows the responsiveness of mortgage rates to observables (FICO, LTV) and time effects to differ across the two samples.

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 (given the prime jumbo sample includes only loans that were originated up to two times the conforming limit). These loans will all be from the prime jumbo sample. Loans below the conforming threshold will have a metric between 0 to 1. 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. As noted above, we 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 this threshold. This ensures that there are similar numbers of loans in each symmetric bin to the left and right of the threshold.

¹⁶It is important to note that because the measures of predicted default are in different units, the coefficients cannot be directly compared across rows within a given column. 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.

Selection is a potential concern for any such regression discontinuity approach, and we address it in a number of ways. More specifically, the concern is that loans just above the threshold may be similar on observables but might differ on unobservables that affect their propensity to default. This type of selection would not be surprising given the large financial benefit in terms of lower average interest rates 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. Figures 3a and 3b show that there is no discrete change in FICO scores or LTV ratios so that observable characteristics do not change across the conforming threshold. This is not surprising given that the samples were matched on exactly these measures.

Figure 3c 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 (on unobservables) put up more cash so that they secure a loan lower than the conforming threshold. Figure 3c shows that there is a very slight increase in default probabilities for prime 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). Although 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 default probability relative to the GSE loans just below the threshold. The differential default probability between GSE loans close to the 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. Thus, although 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.¹⁷

Figure 4 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 3c, 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.

¹⁷Given that the second bin has a loan value that is, on average, between \$40,000 and \$80,000 above the threshold, it is challenging for most households buying a \$500,000 home to substantially reduce the loan balance so that it could be securitized by the GSEs.

IV.C How Much Should GSE Loan Rates Have Varied with Predictable Default?

In this subsection, 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. Table 4 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 sample and our prime jumbo sample during the 2001-2006 period. The last column examines predicted default measures for a sample of GSE loans restricted to the same MSAs as the prime jumbo loans, but during the 2007-2009 instead of the 2001-2006 period.

Table 5 is our key counterfactual table. Given the standard deviation of predicted default rates (Table 4), Table 5 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 (5) of Table 3). Table 5, 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 5, which uses the regression measure of predicted default) suggest 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 difference between the two periods results from the fact that the variation in predicted default across regions was much higher during the 2007-2009 period.

The other specifications of lagged default give roughly similar estimates. In our modeling section below, we are particularly interested in measuring the extent of resource transfers due to the GSEs' constant interest rate policy during the Great Recession because regional risk was particularly important during this time period. With this goal in mind, we choose parameters so that a two-standard-deviation shock to local economic activity across regions would generate a 25 basis point movement in mortgage rates across regions if the GSEs abandoned its constant interest rate policy and allowed mortgage rates to adjust to local default risk as in the private market. Given our counterfactual estimates for the other predicted default measures shown in Table 5, we examine the robustness of our model results when a two-standard-deviation shock causes a 15 basis point or a 35 basis point movement in mortgage rates across regions.

IV.D Robustness and Extensions

It is useful to explore the robustness and extensions of our results to alternate specifications and controls before turning to quantifying the economic importance of the constant interest rate policy for regional redistribution. We describe the specific robustness tests in more detail in the accompanying appendix.

IV.D.1 Exploring Regional Variation in Mortgage Quantities Instead of Interest Rates

Our analysis has only explored the adjustment of mortgage prices in response to spatial variation in regional risk. One may also expect some adjustment to occur on the quantity side, i.e., both on the extensive (loan approval) and intensive (loan amount, conditional on approval) margins.¹⁸ Unfortunately, we are not able to explore variation on the extensive margin, because the only available data on the extensive margin (HMDA database) does not have borrower-level variables, which are crucial for differentiating borrower-level risk from location specific risk. We are, however, able to explore quantity movements on the intensive margin using our data. Appendix Figure A1 shows the relationship between lagged default rates and LTV residuals for both the GSE sample (top panel) and the prime jumbo sample (bottom panel). These figures are similar to Figure 2. We residualize LTV controlling for FICO score and time effects in a way similar to our residualization of interest rates. As seen from Appendix Figure A1, there is little LTV adjustment across MSAs in response to differences in lagged default rates in either sample. If anything, borrowers in riskier places are slightly more leveraged on average. Moreover, there are no statistical differences in the response rates between the two samples. Given this finding, we focus our analysis only on the price margin and not the quantity margin.

IV.D.2 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 prime jumbo loans in a way that is correlated with local default risk, the lack of variation in GSE mortgage rates with local default risk may not be surprising.

In our data, we can track prepayments and thus create a measure of predicted local prepayment risk. In particular, we follow the same procedure as above for local default risk and create three different measures of predicted prepayment risk: the regression-based approach for both samples using lagged local GSE prepayment rates, a perfect foresight model in which the predicted prepayment rate is the actual local prepayment rate for each sample, and a random walk model in which the predicted prepayment rate is the actual lagged local prepayment rate for each sample. Our first finding is that predicted prepayment rates, conditional on loan and borrower observables, are very similar for GSE and prime jumbo loans. For example, using our RD approach, 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 (38% vs. 40%).

What matters is whether predicted prepayments are differentially correlated with predicted default rates across the two samples in a way that undoes the results documented above. To explore this, we added predicted prepayment rates as an additional control to all our main empirical specifications. Table 6 shows one such specification. We

¹⁸Note that most models would imply that when lenders reduce loan quantities they would also raise loan prices, so the fact that there is no regional variation in GSE loan prices strongly suggests there is no quantity variation.

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

IV.D.3 MSA Fixed Effects

Another potential concern with the interpretation of our previous results is that identification could be driven by across-MSA differences in the composition of GSE versus private loans rather than from differential responses of these loans to common local conditions. To address this concern, we reestimated all our specifications including MSA fixed effects. This allows us to compare GSE loans within an MSA to prime jumbo loans within the same MSA. Column (3) of Table 6 controls for MSA fixed effects in our RD specification, while column (4) controls for both MSA fixed effects and local prepayment risk. As can be seen from the table, the estimated difference in interest rate responsiveness to local default risk, $(\beta^{jumbo} - \beta^{GSE})$, is essentially unchanged in all the specifications.

IV.D.4 Points, Fees, and Private Mortgage Insurance

Up until this point, we have not examined regional variation in points paid or other loan fees because points and fees are not recorded in our data. It may be the case that mortgage rates do not vary across MSAs in the GSE sample, but that points and other fees do vary with local default risk. To address this concern, we obtained additional data from one of the GSEs to directly estimate the relationship between effective interest rates and regional risk. The measure of effective interest rates in this data nets off any points and fees (including closing costs) that are charged to the borrower. As shown in Appendix Table 4, we find no significant relationship between effective interest rates in the *universe* of GSE loans that meet our sample criteria and regional risk, as measured by lagged GSE default. The effect of a two-standard-deviation increase in regional risk is an insignificant 5 basis points increase in effective interest rate. In results not shown, no component of the effective interest rate (either points or fees) were found to be statistically associated with regional risk for these loans.¹⁹

One additional concern with our analysis is that the securitized lenders may require borrowers with higher LTVs

¹⁹For additional robustness, we also secured data from LoanSifter, a company that collects loan quotes from various lenders about the interest rate they charge for a given loan type. Loan type is defined as a function of FICO score, fixed interest rate, points charged, and initial LTV. 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 advantage of this data is that points are held fixed across all loan quotes. Given the loan size, all quotes were for loans eligible for securitization by the GSEs. Within this data, we find no relationship between quoted mortgage rates for a given contract at the local level and local measures of default risk. These results can be seen in Appendix Figure A2.

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

IV.D.5 Exploiting Time Variation in the Conforming Limit

Finally, we conduct a robustness specification that exploits time variation in the conforming limit. In particular, we focus on specific set of loans of a given dollar amount that changed between being above the conforming limit to being below the conforming limit during the early 2000s. For instance, a loan amount of \$350,000 was above the limit (and thus excluded from the GSE market) in the period prior to 2005, but in 2005 the conforming loan limit was raised such that a loan of this size could be purchased and insured by Fannie Mae or Freddie Mac. We then ask if interest rates for loans of \$350,000 respond to local predicted default risk prior to 2005 but not respond to local predicted default risk after 2005. To explore this, we reestimate our main specifications using only the set of loans in the range of \$276,000 and \$417,000 that change conforming status between 2001 and 2006. For these loans, prior to the conforming limit change, they responded significantly to local predicted default risk. However, for loans in this same range, interest rates no longer responded to local default risk once they became conforming loans. The magnitudes of the differences are consistent with our estimates in Table 3. We view this as further evidence that GSEs do not vary interest rates in response to local predictable default risk.

V Why Do GSE Rates Not Vary with 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?²⁰ There is some evidence that the quasi-public nature of the GSEs may impose

²⁰There is nothing in the GSE charters that prevents charging differential interest rates across localities. However, the current charter of the Federal Home Loan Mortgage Corporation states that the GSEs are to “promote access to mortgage credit throughout the Nation...by increasing the liquidity of mortgage investments and improving the distribution of investment capital available for residential mortgage financing.” See Lucas and Torregrosa (2010) for a discussion of the origins of the GSEs being driven in part by the volatility in mortgage access across U.S. subregions in the periods surrounding the Great Depression.

political economy constraints on the extent to which they can vary mortgage rates across space. In early 2008, the GSEs attempted to implement a “declining market” policy that restricted credit differentially across U.S. locations. The policy required more equity at the time of origination in markets for which house prices were declining. In non-declining markets, the GSEs would purchase mortgages that had an initial LTV lower than 95%. However, in declining markets, the GSEs would only purchase mortgages where the initial LTV was lower than 90%.²¹ 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.²² 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.”²³ The *Washington Post* reported, “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.”²⁴ 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 GSEs, proposed a new 25 basis point fee at the time of origination that would differ across locations. The fee was tied to states that had long judicial delays in foreclosures. The rationale was that these states’ institutional features increased the length of the foreclosure process and the associated GSE losses. At the time of its original announcement, the fee would only have 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. For example, 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 states (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 had been delayed indefinitely.

Even though these policies focused on imposing either spatial variation in down payments or spatial variation in loan guarantee fees, they can shed light on reasons why the GSEs may not raise interest rates in riskier markets

²¹Fannie 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.

²²See “Fannie Mae Sets New Loan Boundaries,” *NPR’s Marketplace*, April 18, 2008.

²³See “Fannie Is Poised to Scrap Policy Over Down Payments,” *Wall Street Journal*, May 16, 2008.

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

during a 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 toward markets with weaker economic activity and greater default risk.

VI Consequences of Constant Interest Rates: A Structural Model

We now turn to a more formal quantitative assessment of the welfare effects of the GSEs' constant interest rate policy. We begin by discussing a naïve back-of-the-envelope calculation using the above estimate and aggregate mortgage market statistics. We then discuss why this back-of-the-envelope calculation likely overstates the extent to which resources get transferred across regions through the constant interest rate policy. We then set out the details of our structural model.

VI.A A Back-of-the-Envelope Calculation for Regional Redistribution

To compute our simple back-of-the-envelope calculation of the extent to which resources get transferred across regions through the constant interest rate policy, we need three pieces of information. First, we need a measure of how much interest rates should have responded if GSE mortgage rates varied as much as the private market. Our RD results above (column 5 of Table 3) show that a one-percentage-point increase in predicted default results in a 13.48 basis point additional increase in mortgage rates in the prime jumbo sample relative to the GSE sample. Second, we need a measure of the variance of predicted default. Table 5 provides measures of predicted default for different time periods and different samples. Finally, we need a measure of average loan size to which the interest rates apply. For our entire sample period, the median loan size in the GSE sample was \$240,000.

We can compute a rough estimate of the transfers which are induced by the GSE constant interest rate from regions with predicted default that is one-standard-deviation below the average predicted default to regions with predicted default that is one-standard-deviation above the average predicted default. We find that in 2007-2009, a typical homeowner in a region with predicted default one standard deviation below average transfers roughly \$350 to other regions (by paying higher mortgage rates than they otherwise would ($\$240,000 * 0.011 * 0.1348$)). This calculation applies only to homeowners, but we can scale it by homeownership rates in the U.S. economy (two-thirds) to arrive at a transfer of about \$230 per household. The model is symmetric in that households in regions with predicted default one standard deviation above average default will receive a net transfer of \$230 per household.

Although this back-of-the-envelope calculation is straightforward to implement, it has three deficiencies. First, the total amount of transfers from good regions (below average predicted default) to bad regions (above average

predicted default) depends on the persistence of the regional shocks. If the regional predicted default probabilities are purely transitory, households in regions one-standard-deviation below the average will receive a transfer of \$230 for one year and then (in expectation) zero thereafter. On the other hand, if the regional predicted default probabilities are perfectly persistent, then the households in regions one-standard-deviation below the average will receive a transfer of \$230 per year forever. Thus, the total amount of transfer depends on the persistence of the local predicted default process that needs to be specified. 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. Such re-optimization necessarily means that for a given level of persistence, the back of the envelope calculation will be biased upward. Finally, the interest rate policy may have general equilibrium feedback effects on house prices and local economic activity. If the constant interest rate policy props up local income and house prices, the redistribution effects can be larger than the back-of-the-envelope calculation.

VI.B Structural Model

In this section, we lay out a quantitative spatial model that captures various salient factors of the U.S. housing market. This structural model is designed to overcome the deficiencies of the simple back-of-the-envelope calculation. We develop a multi-region life-cycle consumption model 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 nondurable consumption and liquid savings. Owner-occupied housing is subject to fixed adjustment costs but serves as collateral against which individuals can borrow to smooth nondurable 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 for households of different ages. Because housing portfolios differ dramatically by age, the GSE policy should differentially affect households in different parts of the life cycle. Moreover, as our micro data does not include information on household age, calculating these distributional effects requires our structural model.²⁵

The previous sections have shown that GSE mortgage rates do not respond to regional shocks. Accordingly, 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. We will use the empirical work in the prior part of the paper to discipline the counterfactual interest rate policy.

²⁵In our base model specification, we do not allow for the interest rate policy to have feedback effects on either local house prices or on local economic activity. However, after presenting our base specification, we explore how our results change by allowing such potential feedback effects. As we show below, allowing for such feedback effects only strengthens the magnitudes from our base model specification.

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. The second channel further lowers welfare as higher interest rates mean that households in deteriorating regions will delay purchasing housing and reduce the sizes of their eventual purchases. A constant interest rate policy eliminates these effects, and it is in this sense that the policy transfers resources toward regions experiencing deteriorating local economic conditions.

VI.B.1 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 60.²⁶ After retirement, households face stochastic mortality risk with probability of death d_j . Households live to a maximum age of 85, so $d_{85} = 1$.

Households live in specific regions indexed by k , and we assume that households never move. In our above empirical results, we showed that local economic conditions such as lagged mortgage default predict current mortgage rates in the prime jumbo sample but that there is no such relationship in the GSE sample. While we would want to include various dimensions of regional economic activity in our model, it is intractable to include all the separate dimensions as separate stochastic processes. Instead, within our model, we capture various measures of local economic conditions in a parsimonious manner by collapsing them into a single stochastic process, γ . We assume that this measure of economic activity (γ) in region k and period t follows the following process:

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

In turn, we assume that $\gamma_{k,t}$ affects other local variables such as income and house prices. The effect of this regional shock $\gamma_{k,t}$ on these other aspects of the model will be made concrete below as we describe the evolution of income and house prices.

VI.B.2 Preferences and Household Choices

Household i receives flow utility

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

²⁶Since we are interested in how interest rate policies interact with household self-insurance, we focus on households prior to retirement who are subject to labor market shocks. In the data, retirement starts to rise substantially around age 60. Since we do not model this endogenous decision, we choose a low threshold for exogenous retirement in the model following the approach in Kaplan and Violante (2010).

from nondurable consumption c_{ijk} and housing services h_{ijk} .²⁷ Households discount expected flow utility over their remaining lifetimes with discount factor β .

VI.B.3 Income Shocks

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

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

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

where χ_j is a deterministic age profile common to all households, $z_{i,t}$ is a purely idiosyncratic persistence income shock, and $\phi^y \gamma_{k,t}$ is a region-specific shock to income. ϕ^y is a parameter that governs the sensitivity of household income to the underlying latent local economic conditions. One symptom of a depressed local economy will be declines in household income, and this is captured by that parameter. It is important to include this channel because changes in household income will directly affect the borrowing decisions of households and, as a result, will affect their response to interest rate variation.

When retired, households receive Social Security benefits. These benefits are based on lifetime earnings prior to retirement, and they are deterministic until household death. We describe the computation of these benefits in the calibration section of the model, but they mirror payments under the actual US Social Security system.

VI.B.4 Housing Markets and Interest Rates

Housing services can be obtained from owner-occupied housing or through a rental market. Housing can be purchased at price $p_{k,t} = (\gamma_{k,t})^{\phi^h}$ or rented at price $p_{k,t} r^f$. We assume that house prices move exogenously with local economic activity, and ϕ^h governs the strength of this correlation. 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 that is proportional to the current value of the house. That is, the fixed fraction lost for household i when the owners buy or sell their home takes the following form:

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

²⁷This specification of utility between housing and nondurable consumption is common in the literature. See, for example, Piazzesi et al. (2007) and Davis and Van Nieuwerburgh (2014). Using data from the consumer expenditure survey, Aguiar and Hurst (2013) find that the share of expenditure households allocate to housing out of total expenditure is roughly invariant with either the level of household income or the level of household expenditure.

Offsetting the disadvantage that it is costly to adjust one's housing services, 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 $\delta^f > \delta^h$. This is a standard assumption that provides a reason that individuals prefer to own rather than rent. 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:

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

Thus, $\delta^f > \delta^h$ implies that the imputed rental price of owner-occupied housing is lower than that of renting.

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

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

where the risk adjustment is declining in regional economic activity:

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

$\bar{\Psi}$ is a fixed risk adjustment associated with mortgage lending that is constant across locations. ϕ^r represents the sensitivity of local mortgage rates to local economic conditions. In our base specification, we set $\phi^r = 0$, consistent with the patterns documented for the GSE loans described above. Our main counterfactuals will be based on changing ϕ^r so that it matches the regional variation in response to predicted local default risk found among the prime jumbo loans.

Finally, 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 . We assume that households are otherwise liquidity constrained in that they can only borrow against the value of their home.

VI.B.5 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 of age j in region k . b_j, m_j, h_j reflect start of the period values before households make their current period 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. Because there are no aggregate shocks, separately tracking t does not change the solution.

Within each period, households choose whether to adjust their housing stock, continue owning their same home, or continue to rent. The adjusters include those homeowners that remain homeowners but change the size of their house, those homeowners that become renters, and those renters that become homeowners. Conditional on their adjustment decision, the households choose the level of their consumption, their savings in bonds, and their mortgage debt.

Formally, 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

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

For brevity, we leave the value functions for non-adjusting homeowners ($V_j^{noadjust}$) and non-adjusting renters (V_j^{rent}) to the computational appendix. 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, as described above. A computational appendix discusses the numerical solution of the model.

VII Calibration

Our benchmark calibration strategy proceeds in two parts. First, we calibrate parameters that do not depend on regional economic activity to standard values from the literature together with standard moments from wealth data. Second, 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 earnings in the PSID. Their calculation conditions on education and age and so captures residual earnings risk. Since households in our model are ex-ante identical, this is the relevant empirical object. To calibrate the life-cycle profile of earnings, χ , we use the age-earnings profile in PSID data estimated by Kaplan and Violante (2010).

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 lifetime earnings, but this would substantially complicate the solution of the model because these lifetime 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 lifetime income given income in the final period of working and then apply the actual benefit 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 $1/2$. As stated above, our model period is annual, and we accordingly 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-adjustment ($\Psi_{k,t}$) of 0.01 to match the average real mortgage rate in our data. We calibrate $\delta_h = 0.03$ to roughly match the average ratio of residential investment to the residential stock in Bureau of Economic Analysis (BEA) data. We set $\theta = 0.20$ so that households are required to have a minimum 20% down payment. We pick $F = 0.05$ so that there is a 5% transaction cost from adjusting housing. This encompasses costs of real estate broker fees, closing costs, and other costs associated with buying/selling a home.

We jointly pick β, r^f and α to match various wealth and home ownership targets. We do so under the assumption that $\phi^r = 0$, which corresponds to the data-generating process under current policy. To estimate these parameters, we target a home-ownership rate of 69% as in the Survey of Consumer Finances (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 $\alpha = 0.88$.

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 (γ). Using annual employment data from the U.S.'s Bureau of Labor Statistics (BLS) from 1991-2013, we estimate an annual AR(1) process for log MSA employment, which yields $\rho_\gamma = 0.947$ and $\sigma_\varepsilon = 0.018$.²⁸ These findings suggest that shocks to local employment are somewhat persistent. 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 in our robustness analysis.

To estimate ϕ^h , we regress log MSA house prices on log MSA employment during the same time period, which yields $\phi^h = 0.48$.²⁹ We think of this elasticity as a short-to medium-run adjustment in house prices. In our base analysis, we are abstracting from housing supply adjustments, which limit the relationship between local employment growth and house price growth in the long run. However, to account for differential long-run supply effects, we show the robustness of our results to values of ϕ^h between zero and 0.48. A value of $\phi^h = 0$ implies a perfectly elastic housing supply curve even in the short run in response to local shocks.

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 by the prime jumbo data described above. In particular, we pick ϕ^r so that a two-standard-deviation in γ increases local mortgage rates by 25 basis points. This is the variation in local mortgage rates to a two-standard-deviation change in predicted default during the 2007-2009 period for prime jumbo loans (see Table 6).³⁰ 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 5 shows the life-cycle profiles in our model compared with 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 nondurable consumption as well as the increasing homeownership rate as households age. Like many models, we underpredict household wealth holdings in the model relative to the data for older households. 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

²⁸When estimating the AR(1) process for MSA employment, we remove permanent differences in employment across MSAs by including MSA fixed effects. Likewise, 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.

²⁹For the local house price series, we use the MSA level repeat sales price indices published by the Federal Housing Finance Authority (FHFA). The loans underlying the Fannie and Freddie mortgage pools are used to make these price indices.

³⁰The implied elasticity of the total borrowing rate ($r + \bar{\Psi}\gamma$) to γ is 0.54.

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 to provide a broad estimate of the impact of the GSE constant interest rate policy on the economy. Household re-optimization in response to changing policy is a first-order effect that must be modeled in order to get the impact of this GSE policy roughly correct. We are less concerned that modest departures between our model and data along the above dimensions will dramatically affect our broad policy conclusions.

VIII Model Results

To examine the effect of a constant interest rate policy on household well-being, we simulate household consumption under both the constant interest rate and the variable interest rate policy. For ease of discussion we label regions with low economic activity and high predicted default “bad” (low γ) regions and regions with high economic activity and low predicted default “good” (high γ) 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.³¹

VIII.A Baseline Results

To assess the quantitative size of this “transfer” between good and bad regions under the constant interest rate policy, 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.³²

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 $\widetilde{c}_{jk}$ and $\widetilde{h}_{jk}$ be the choice for nondurable consumption and housing services, respectively, that obtain this maximal value. Finally, let $E_{\gamma,z,j}$ denote the expectation of these value functions over values of the

³¹Note that the impact of households defaulting are built into the model through its effect on the risk adjustment factor (Ψ). However, for tractability, we have chosen not to explicitly model households’ default decisions. Endogenizing the default decision will only matter if the elasticity of default with respect to changes in the mortgage rate is large. It is only the interest rate that is changing across our policy experiment. Given that the variation in interest rates across the two policies is small, we do not expect our welfare calculations to change much if default was explicitly modeled.

³²The Computational Appendix which accompanies the paper discusses exactly how we solve the model. With respect to initial conditions, we assume households start their working lives with no wealth. We also assume that the initial distribution of the idiosyncratic and regional income variation matches the observed cross sectional distributions for those variables.

idiosyncratic shock and age, conditional on living in a region with economic activity γ .³³ 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 today 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. We discretize the distribution of γ and focus on regions that had shocks that were one and two standard deviations above and below the mean region. The first row expresses the utility gain/loss (λ) from the constant interest rate policy. This is the lifetime consumption gain as a fraction of today's consumption. The second row turns the lifetime consumption equivalent into a dollar amount. To do this, we first estimate average consumption in dollars per household by dividing consumption from the BEA by the number of households in the U.S.³⁴ This calculation gives that average household consumption is roughly \$82,000. Although 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 as well by using the model's implied consumption difference across regions to scale average consumption accordingly.

The worst regions (on the left side of the table) are made substantially better off by the constant interest rate policy, while the best regions (on the right side of the table) are made substantially worse off. The calibration suggests that the lifetime gain from moving to the constant interest rate policy for the regions hit with a negative two-standard-deviation shock is roughly 1 percent of today's consumption. The results are symmetric in that the regions that received a positive two-standard-deviation local economic shock are made strictly worse off. They would be willing to pay roughly 1 percent of today's consumption to avoid the constant interest rate policy. In terms of dollar values, the constant interest rate policy redistributes roughly \$900 per person from the best regions to the worst regions in Table 7 during a period of increased regional dispersion similar to the Great Recession. This is a net transfer of \$1,800. Although we will talk more about magnitudes below, this \$900 transfer to relatively depressed regions (those with a two-standard-deviation negative shock) is of similar size to the tax rebate checks authorized by the U.S. Congress during the 2001 and 2008 recessions (which ranged from \$500 to \$1,000 per qualifying household). The difference is that the transfer provided via the constant interest rate mortgage policy is funded by "taxing" the regions that are doing relatively well by roughly the same amount. The constant interest rate policy provides lower

³³In 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 behavior in this final period is essentially irrelevant for working age households.

³⁴Total consumption in 2013 was \$9,398 billion = (\$11,484 billion personal consumption expenditures minus \$2,086 billion 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 \$81,722 per household.

interest rates to the regions hit with relatively negative economic shocks by charging higher mortgage rates in the regions hit with the relatively positive economic shocks. As discussed in Section II, the GSEs can maintain this pricing scheme given their cost of funds advantage relative to private lenders.

Rather than focusing on implicit transfers from regions with particular employment shocks, we can also calculate the total resources transferred from *all* regions with positive employment shocks to *all* regions with negative employment shocks.³⁵ The average household in a location with a negative regional employment shock values the constant interest rate policy at \$205 (a one-time consumption equivalent). At the same time, the average household in a location with a positive employment shock would pay \$154 in lifetime consumption units to move to a variable interest rate policy. Thus, the total implicit net transfer implied by the constant interest rate policy is \$360 per household. This is roughly one-third to two-thirds of the transfers received from the U.S. tax rebates paid during the past two recessions. These payments were also one-time payments made to households during a recession. The total present value of transfers through the constant interest rate policy from regions receiving any positive regional shocks to any regions receiving negative regional shocks for a period where the shocks are similar to the Great Recession is \$20.7 billion.³⁶ For comparison, the Department of Labor forecasts that total unemployment insurance benefits paid in 2014 will equal \$49 billion.³⁷ In summary, our results imply large transfer payments across regions so that the constant interest rate policy has significant redistributive consequences across regions.

Although the estimated annual transfer per household estimated in our base case of the model is quite large, it is important to note that, as predicted, it is substantially smaller than the annual net transfer per household implied by the back-of-the-envelope calculation discussed above.³⁸ This highlights that ignoring household responses to changing interest rate policy can lead to an overstatement of the welfare consequences of a constant interest rate policy. To understand the differences between the model results and the back of the envelope calculation, it is helpful to explore the factors within the model driving the attenuation. 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 and (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.

³⁵To compute this average transfer, we compute $\int_{\gamma < 0} f(\gamma) (1 + \lambda(\gamma)) d\gamma$ and $\int_{\gamma > 0} f(\gamma) (1 + \lambda(\gamma)) d\gamma$, where $f(\gamma)$ is the probability of experiencing a given γ shock and $\lambda(\gamma)$ is the welfare evaluation of a constant interest rate policy for a household living in region with that γ . To provide a better approximation to $f(\gamma)$ we expand the γ grid from 5 to 15 points for this calculation, but we find nearly identical results when using 5 points.

³⁶This is equal to $\$360 \times 115$ million households divided by two. We divide by 2 because half the households in our model live in regions that get positive shocks while the other half live in regions that get negative shocks.

³⁷See http://www.dol.gov/_sec/media/reports/annual2013/2013annualreport.pdf, page 123.

³⁸In order to compare the back-of-the-envelope calculation directly with the model results, we need to adjust the back-of-the-envelope calculation so the persistence of the shock is similar to the model. With a persistence of about 0.95 and a household discount rate of 8%, as in our baseline model, the present value gain from a constant interest rate policy in the back-of-the-envelope calculation to a one standard deviation regional shock during a period like the Great Recession of about \$1,800. This number is about 4 times bigger than the transfer our model estimates for a similar shock.

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 who own a home already and have accumulated equity. Table 8 uses our model to explore the age-specific effects of the constant interest rate policy more formally. This table recomputes the calculation in Table 7 but conditions the results on age. That is, we look at how much young households (those aged 25-34) and middle-aged households (those aged 35-60) in different regions are willing to pay to move to a constant interest rate policy. Again, within the model (and data), those aged 25-34 are most likely to be purchasing their first home.

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 for middle-aged households, the first channel (borrowing to smooth consumption) is more important than the second (the effect on home purchases). The welfare gains are dampened for younger households because when there are negative regional income shocks, households in bad regions delay home purchases. 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 of shutting down the rental market and forcing all households to purchase houses. When there is no rental market, the welfare effects of the constant interest rate policy are amplified because households lose access to an important margin of adjustment. Unsurprisingly, this moves the model results closer to the back-of-the-envelope results highlighted in Section V, since that calculation assumes that households do not endogenously respond to changing mortgage rate policies. Additionally, the results by age in Table 9 are the reverse of those in Table 8. Without a rental market, the young in regions experiencing a negative shock are worse off with a variable interest rate than are the middle-aged. 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.

All of our results so far (in Tables 7, 8, and 9) focus on the ex-post redistributive consequences of the GSE pricing rule, because 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 on utility. 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 in consumption, it would necessarily lower ex-ante welfare. For our benchmark results, we find a very small overall consumption equivalent response of a lifetime gain equal to 0.04

percent of today’s consumption from moving from the variable interest rate to the constant interest rate policy. The ex-ante welfare effects we find are quite small and are also sensitive to some of the robustness exercises we perform below.³⁹ Part of the reason for the small welfare effects is that in general the costs of business cycles are quite small (a point made in Lucas 1987). The welfare costs of the regional business cycles in our model are slightly higher because the regional shocks we estimate are more persistent than the shocks typically estimated for aggregate business cycles. But, despite the potentially small welfare costs of business cycles, our ex-ante risk sharing results are also small given that in expectation individuals are equally likely to be a net recipient of a regional transfer as they are to be a net provider of a regional transfer. A region that receives a bad shock today will likely receive a good shock in the future. When a region receives a good shock, they are worse off under the constant interest rate policy and this lowers the ex-ante welfare gains of moving to a constant interest rate policy.

VIII.B Robustness

Despite having relatively small ex-ante welfare gains, our model predicts large ex-post transfers across regions. These results are robust across many of our parameter specifications.

Table 10 shows how our base results change as ϕ^r changes. In our benchmark calibration, we pick ϕ^r to match the variation in interest rates observed in the jumbo market during the Great Recession. However, if the jumbo market piggybacks off of the GSE policy in picking interest rates, the true sensitivity might be larger than what we find in the data. Conversely, our empirical section controlled for various observables around the conforming threshold and argued that observable default varies smoothly across the threshold. However, in presence of sorting – with more risky borrowers being unable to supplement resources that would allow them to borrow below the conforming limit – our benchmark estimates of ϕ^r might be overstated. Table 10 shows implied transfers under various values of ϕ^r . In particular, we targeted a 35 basis point variation in response to a two-standard-deviation regional shock (referred to as “larger variation”) and then separately a 15 basis point variation in response to a two-standard-deviation regional shock (referred to as “smaller variation”). Unsurprisingly, the level of implied transfers is increasing in ϕ^r . Increasing the variation by 10 basis points in response to a two-standard-deviation change in local economic activity increased the total transfers from good to bad regions by about 50 percent relative to the base specification. Cutting the variation by 10 basis points in response to a two-standard-deviation change in local economic activity reduced the total transfer from good to bad regions by about 50 percent. The key point is that even with a much smaller value of ϕ^r (compared with what we estimate from the prime jumbo data), the transfers from good regions to bad regions through the constant interest rate policy remain quite large.

Table 11 explores how the overall level of regional risk interacts with the GSE constant interest rate policy. The

³⁹Note that this number is not the sum of the numbers in row 1 of Table 7 for two reasons: 1) There are not equal numbers 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 given that we measure “good” and “bad” regions as being ones that receive one or two standard deviation regional shocks. 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.

benchmark results are redisplayed at the top of the table, and the next rows show the sensitivity of our results to different values of ϕ^y . In our base specification, we set $\phi^y = 1$. As a robustness exercise, we show results for $\phi^y = 0.5$, and $\phi^y = 0.0$. Basically, the different estimates of ϕ^y determine the extent of regional income risk. In the extreme case, $\phi^y = 0$, there is no regional income risk at all. The next rows of the table show the sensitivity of our results to reducing ϕ^h , where ϕ^h is the responsiveness of local house prices to changes in local economic activity. If housing supply is perfectly elastic in all periods, $\phi^h = 0$. To explore the sensitivity of our results, we recompute results setting $\phi^h = 0.25$ and $\phi^h = 0$. The main take-away from Table 11 is that varying sensitivity of regional income and house prices has essentially no effect on how households evaluate interest rate variation. In other words, although regional income and house price variation matter for the level of risk faced by households, they have little effect on the way that households are affected by changes in the interest rate across the two policies.

More specifically, the counterfactual interest rate policy asks how much a household would pay to avoid having the interest rate rise today. What the non-interaction with ϕ^y and ϕ^h means is that the answer to this question depends little on whether a household is currently in a high-house-price area, a low-house-price area, a high-income area, or a low-income area. This can be seen most clearly in the second to last row, where both $\phi^h = \phi^y = 0$ so there is no regional variation in income or in house prices. This means that under the constant interest rate policy, regions are identical. The only way that regions vary when $\phi^h = \phi^y = 0$ is that in the variable interest rate counterfactual, some regions will have high interest rates and some regions will have low interest rates. Comparing that row with the benchmark shows that the welfare effects of interest rate variation in a world where house prices and income vary spatially is very similar to one where they do not. This occurs because there are offsetting interactions between regional risk and interest rate variation. For example, higher income makes households more likely to want to buy a big house, so that they dislike high interest rates more than low-income households do. But at the same time, higher income also makes households less likely to borrow, so they care less about high interest rates. The net effect is that these two effects roughly cancel each other out so the level of regional risk has little effect on implicit transfers from interest rate variation.⁴⁰

VIII.C Model Extensions and Discussion

In our benchmark results, for simplicity, we assume that regional economic conditions do not depend on interest rate policy. However, it is possible that changing interest rates will have a feedback effect on local employment and local house prices. For example, when interest rates are lower, the demand for housing will increase and this will put upward pressure on house prices. Additionally, lower interests, through their increase in house prices, could result in higher local employment and earnings. This latter effect is emphasized in the recent work of Charles et al. (2013). If the constant interest rate policy is abandoned, interest rates will rise in locations with poor economic conditions.

⁴⁰It is important to note that this non-interaction does not imply that ϕ^y and ϕ^h do not matter for welfare: households are much worse off in the world with regional income variation. But the relevant question is to what extent households' willingness to tolerate interest rate variation interacts with these other regional shocks. Table 11 shows that there is no substantive interaction.

This increase in interest rates will then exacerbate the already poor economic conditions by further depressing house prices and income in that location. We excluded this feedback mechanism from our baseline model because (1) we think such feedback will only amplify our estimated cross-region transfers and (2) it is harder to select empirical estimates to discipline this feedback mechanism.

However, we feel it is important to at least attempt to quantify this endogenous feedback mechanism. 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 federal funds rate 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. In this case, unlike in the previous sensitivity exercises, there is now a direct interaction between regional interest rate variation and the level of regional risk: When interest rates are more variable, the level of regional risk rises. The final row of Table 11 shows that allowing for this form of feedback from interest rates to local economic conditions amplifies the degree of transfers implied by the constant interest rate policy. In particular, using the estimated feedback from Christiano, Eichenbaum, and Evans (1999) and Vargas-Silva (2008) roughly doubles the amount of transfers implied by the constant interest rate policy. Given this finding, we view the results in our base specification as a lower bound on the actual regional transfers made through the constant interest rate policy of the GSEs.

One potential criticism of our model is that it does not endogenize housing prices. Although we think endogenous house prices would be a good addition to the model, doing so would significantly complicate the analysis. To reassure ourselves that this abstraction is not dramatically altering our results, we explored the sensitivity of our results to alternate values of ϕ^h . Assuming a value of $\phi^h = 0$ implies that housing supply is perfectly elastic. Given that our results were essentially invariant to alternate values of ϕ^h (reported in Table 11), we feel the abstraction of exogenous housing prices is not biasing our results in any meaningful way.

A more serious assumption we make is that labor is immobile across regions. One way to support regional equilibrium is allow factors to move across regions. Again, adding costly labor mobility to our model would significantly complicate the analysis. Two facts make us believe that allowing for such mobility would not alter our results in any meaningful way. First, we estimate our regional income processes on actual data. The underlying data takes into account both the true underlying process driving the regional shocks to income as well as any endogenous response of factors across regions. Our approach cannot distinguish between large regional shocks that are mitigated in part through factor mobility and slightly smaller shocks in a world with less factor mobility. Given that we are using these processes to calibrate our model, any migration that actually occurs will be captured in our estimates. Second, there is a large literature showing that permanent regional shocks lead to sizable migration responses (see, for example, Blanchard and Katz 1992). There is less evidence that regional migration is important in response

to the kinds of temporary regional shocks captured by our model. With costly migration, individuals may choose to ride out the regional business cycle as opposed to paying the migration cost and moving to another region. In fact, there was very little net migration from regions hit hard during the most recent recession to regions that were hit less hard. For these reasons, we think that abstracting from migration will not significantly change the model results.

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 some 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 \$12 trillion U.S. 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 that is comparable in size to recent fiscal stimulus programs such as tax rebates and tax holidays. The extent of such redistribution can be compared with the costs of providing such insurance through implicit subsidies to the GSEs, including too-big-to-fail subsidies, and can inform the debate on costs and benefits of the GSEs.

Although a range of consequences to the housing and mortgage markets 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 countercyclical 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 two important caveats of our result. 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, LTV ratio, documentation type, fixed-rate, and 30-year term. However, we realize that political economy considerations may also limit the extent to which interest rates can vary spatially even 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 their pricing models' transparency for secondary market investors. Both of these factors may lead us to likely understate the true spatial variation we would 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.

Additionally, we want to stress that we are not saying the GSE policy is the optimal way to transfer resources across regions in state-contingent ways. Many have argued the potential welfare-reducing role of the GSEs in distorting the allocation of capital. Moreover, policy makers have many other tools to transfer resources across regions if they so desire. Our goal in this paper is to study the impact of a current policy as opposed to providing either a full welfare analysis of the existence of the GSEs or discussing the optimal way to transfer resources 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.
- Aguiar, Mark, and Erik Hurst, “Deconstructing Lifecycle Expenditure,” *Journal of Political Economy*, 212(1), 2013.
- 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.
- Beraja, Martin, Erik Hurst, and Juan Ospina, “The Regional Evolution of Prices and Wages During the Great Recession,” University of Chicago Working Paper, 2014.
- Berger, David and Joseph Vavra, “Consumption Dynamics During Recessions,” *Econometrica*, forthcoming.
- Blanchard, Olivier J., and Lawrence F. Katz, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992:1, pp. 1-75, 1992.
- Calomiris, Charles W., “An Economist’s Case for GSE Reform,” in P. Wallison, ed., “Serving Two Masters Yet Out of Control,” Washington, D.C.: AEI Press, 2001.
- 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.
- Davis, Morris A., and Stijn Van Nieuwerburgh, “Housing, Finance, and the Macroeconomy,” NBER Working Paper No. 20287, July 2014.
- 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.
- Green, Richard K., and Susan Wachter, “The American Mortgage in Historical and International Context,” *Journal*

- of Economic Perspectives, 19(4), 2005.
- Guvenen, Fatih, and Anthony Smith, “Inferring Labor Income Risk and Partial Insurance from Economic Choices,” Federal Reserve Bank of Minneapolis Research Department Staff Report No. 485, June 2013.
- Kaplan, Greg and Giovanni L. Violante, “How Much Consumption Insurance Beyond Self-Insurance?” *American Economic Journal: Macroeconomics*, 2(4), October 2010.
- Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 125(1), February 2010.
- Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig, “Mortgage Financing in the Housing Boom and Bust,” in E. Glaeser and T. Sinai, eds., “Housing and the Financial Crisis,” Chicago, IL: University of Chicago Press, 2013.
- 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.
- Lucas, Deborah, and Robert McDonald, “Valuing Government Guarantees: Fannie and Freddie Revisited,” in D. Lucas, ed., “Measuring and Managing Federal Financial Risk,” Cambridge, MA: National Bureau of Economic Research, 2010.
- Lucas, Deborah, and David Torregrosa, “Fannie Mae, Freddie Mac, and the Role of the Secondary Mortgage Market,” Congressional Budget Office Study, 2010.
- Lucas, Robert E., “Models of Business Cycles,” Jahnsson Lectures, Oxford, UK.
- 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 Issuer, 1996-2013




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

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




Note: This figure shows the relationship between residualized interest rates and residualized lagged MSA-level default of loans originated within the last two years for three samples. Panel a presents the relationship in the GSE market for all 374 available MSAs. Panel b restricts the GSE loans to the 106 MSAs where non-GSE loans are present, and matched based on the FICO and LTV distributions of non-GSE loans for comparability. Panel c shows the relationship in the non-GSE loan market. The adjusted residual removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects.

Figure 3: Average FICO Score, LTV Ratio, and Default Rate, by Loan Amount, 2001-2006




Note: This figure shows (a) the average FICO credit score, (b) the average LTV ratio, and (c) the average residualized default rate in each loan amount bin around the conforming loan limit. Residualized default rate 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 and LTV distributions of non-GSE loans for comparability. Each point in each figure is an average for a loan amount bin representing 10% of the loan amount distribution from \$0 to twice the conforming loan limit. 95 percent confidence intervals are represented by dashed lines. See text for details.

Figure 4: Relationship between Interest Rates and Three Measures of Default, 2001-2006



Note: This 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 and LTV distributions of non-GSE loans for comparability. Each point in each figure is a regression coefficient for a loan amount bin representing 10% of the loan amount distribution from \$0 to twice the conforming loan limit. 95 percent confidence intervals based on standard errors clustered at MSA level are represented by dashed lines. Standard errors for results relying on predicted default are bootstrapped (500 repetitions, clustered at MSA level). See text for details.

Figure 5: Average Life Cycle Profiles: Model Simulation vs. Data



Note: This figure shows average lifecycle profiles generated from the model compared with actual data. Data sources: Nondurable consumption comes from Aguiar and Hurst (2013). Home ownership rates are calculated from the March CPS, and wealth statistics are calculated from PSID data. All variables (except for the homeownership rate) are normalized to their life-cycle means. 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	70,327	70,327	4,861,259	3,677,984	
Median FICO	728	727	658	656	756	757	
Median LTV	0.78	0.75	0.79	0.80	0.76	0.75	
MSAs covered	374	106	106	106	374	106	
Mean Interest Rate (%)	6.25	6.22	6.33	6.66	5.65	5.63	
Mean 2-Yr Delinquency Rate (%)	1.6	1.4	3.0	2.1	3.8	4.0	
<u>Cross MSA SD of Interest Rates</u>							
Unconditional (percentage points)	0.544	0.557	0.578	0.657	0.627	0.623	
Conditional (percentage points)	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	3.2	2.7	4.0	4.3	
Conditional (percentage points)	1.3	1.1	2.8	2.5	2.9	2.9	

Note: This table provides summary statistics for the samples of GSE and non-GSE (“prime jumbo”) loans. The different columns refer to different samples and different time periods, with the first four columns referring to loans originated between 2001 and 2006, and the last two columns featuring loans originated between 2007 and 2009 (after the non-GSE market ceased large-scale operation). The first column uses all loans in our sample originated by the GSEs, the “Restricted MSA” sample uses only those MSAs with prime jumbo loans present (during 2001 to 2006), and the “GSE Matched Sample” restricts to these 106 MSAs and matches the distribution of FICO scores and LTV ratios in the non-GSE sample. Conditional measure of standard deviation removes year*quarter fixed effects and semi-parametric controls for FICO and LTV interacted with year*quarter fixed effects. See text for details.

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)	3.54 (2.75)	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	2.54	22.61	3.18	3.27
Sample Size	13,109,968	8,052,967	70,327	70,327	4,861,218	3,677,984

Note: This 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 two 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. See text for details.

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	
Predictive Default Measure					(5) RD Coefficient
Predicted Default Using Lagged Local GSE Default	2.10 (1.78)	12.04 (1.68)	9.94	<0.001	13.48 (4.56)
Lagged Default (Random Walk)	3.56 (2.76)	12.60 (3.16)	9.04	<0.001	13.04 (4.57)
Actual Default (Perfect Foresight)	0.26 (0.14)	2.12 (0.40)	1.86	<0.001	2.06 (0.44)
Underlying Sample Size of Loans	70,327	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. Predicted default rates are constructed using lagged GSE default rates. 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 and LTV ratios in the non-GSE sample. The different sample definitions are discussed in the notes to Table 1. The first two columns show the separate OLS estimates, columns 3 and 4 test for differences, while column 5 shows the “regression discontinuity” estimates shown in Figure 4, using bins that are each 20% of the loan amount distribution between \$0 and twice the conforming loan limit. Standard errors in parentheses clustered at the MSA level. Standard errors for results relying on predicted default are bootstrapped (500 repetitions, clustered at MSA level) to account for the generated regressor. See text for details.

Table 4: 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)	0.030	0.027	0.043

Note: This table presents the standard deviation of each measure of predicted default for each sample used in the analysis, GSE loans and non-GSE loans originated between 2001 and 2006, and GSE loans originated between 2007 and 2009. The GSE sample during the 2001-2006 period is restricted to the MSAs where non-GSE loans are present and matched on the FICO and LTV distributions of the non-GSE sample for better comparability. The GSE sample during the 2007-2009 period is restricted to the MSAs where non-GSE loans were present during the period 2001-2006. See text for details of sample construction.

Table 5: 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	0.162	0.297
Lagged Default (Random Walk)	0.104	0.391
Actual Default (Perfect Foresight)	0.124	0.177

Note: This table presents the interest rate response to a two standard deviation change in each predicted default measure for two time periods, 2001-2006 and 2007-2009. These values are obtained by multiplying the values in Table 3, column 5 with two times the standard deviations found in table 4 for GSE loans.

Table 6: Robustness of Regression Discontinuity Estimates

Predictive Default Measure	Specification			
	(1)	(2)	(3)	(5)
Predicted Default Using Lagged Local GSE Default	13.48 (4.56)	12.99 (5.04)	11.73 (4.74)	15.64 (4.56)
Time, FICO, and LTV Controls Included	Yes	Yes	Yes	Yes
Predicted Payment Controls Included	No	Yes	No	No
MSA Fixed Effects Included	No	No	Yes	No
Restrict to LTV ≤ 0.8	No	No	No	Yes

Note: This table presents regression-discontinuity estimates of the difference in the relationship between interest rates and predicted defaults around the conforming loan limit. The regression estimates here are estimated as in Figure 4, using bins that are each 20% of the loan amount distribution between \$0 and twice the conforming loan limit. Lagged default and lagged prepayment measures are constructed within-sample depending on GSE or non-GSE loans. The GSE sample is restricted to the MSAs where non-GSE loans are present and matched on the FICO and LTV distributions of the non-GSE sample for better comparability. Each coefficient represents a separate regression. Standard errors in parentheses clustered at the MSA level. Standard errors for results relying on predicted default are bootstrapped (500 repetitions, clustered at MSA level) to account for the generated regressor. See text for details.

Table 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	1.06%	0.56%	0.02%	-0.52%	-1.09%
Dollar per household effect	\$797	\$434	\$17	-\$446	-\$998

Note: This table shows the consumption gains estimated from our baseline model. See text for a description of baseline parameters and the policy experiment. The consumption gain in row 1 is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. To compute dollar equivalents in row 2, we use the formula $(\lambda \times 100)(\lambda \times 81722 \times \frac{c_{region}}{c_{overall}})$, where \$81,722 is average household consumption, adjusted for the fact that consumption varies with local economic activity. Calculations are restricted to working age households subject to labor market risk.

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
Consumption gain: Overall	1.06%	0.56%	0.02%	-0.52%	-1.09%
Consumption gain: Young	0.58%	0.32%	0.00%	-0.36%	-0.78%
Consumption gain: Middle Aged	1.30%	0.68%	0.04%	-0.60%	-1.24%

Note: This table shows the consumption gains estimated from our baseline model, by age. See text for a description of baseline parameters and the policy experiment. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. Calculations are restricted to working age households subject to labor market risk. We define “young” as households of ages 25-34 and “middle aged” as households of 35-60.

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	+2 Standard Deviations
Consumption gain: Overall	1.32%	0.56%	0.02%	-0.50%	-0.86%
Consumption gain: Young	2.18%	1.16%	0.08%	-1.00%	-2.02%
Consumption gain: Middle Aged	0.90%	0.44%	0.04%	-0.20%	--0.28%

Note: This table shows the consumption gains estimated from our baseline model shutting down the rental market. See text for a description of baseline parameters and the policy experiment. We keep all parameters at their baseline values in this experiment, but set a cost of renting such that all households choose to be homeowners. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. Calculations are restricted to working age households subject to labor market risk. We define “young” as households of ages 25-34 and “middle aged” as households of 35-60.

Table 10: Sensitivity to Different Values of ϕ^r

	Regional Employment				
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviations
Consumption gain: Benchmark (25 bp)	1.06%	0.56%	0.02%	-0.50%	-0.86%
Larger variation (35 bp)	1.60%	0.82%	0.06%	-0.70%	-1.52%
Smaller variation (15 bp)	0.68%	0.32%	0.02%	-0.34%	-0.80%

Note: This table shows the robustness of our estimated consumption gains to different interest rate sensitivities to local economic conditions. See text for a description of baseline parameters and the policy experiment. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. We alter the value of ϕ^r such that the benchmark variation in regional employment produces alternative variability in mortgage rates.

Table 11: Sensitivity to Other Elasticities

	Regional Employment				
	-2 Standard Deviations	-1 Standard Deviation	0	+1 Standard Deviation	+2 Standard Deviations
Consumption Gain: Benchmark	1.06%	0.56%	0.02%	-0.52%	-1.10%
Reduce Regional Income Variation in Half	1.10%	0.56%	0.08%	-0.32%	-0.82%
Reduce Regional Income Variation to Zero	0.98%	0.54%	0.06%	-0.56%	-1.14%
Reduce Regional House Price Variation in Half	1.16%	0.64%	0.10%	-0.56%	-1.06%
Reduce Regional House Price Variation to Zero	1.06%	0.52%	0.06%	-0.38%	-0.88%
Reduce Regional House Price and Income Variation to Zero	1.06%	0.50%	-0.02%	-0.56%	-1.08%
Allow Endogenous Feedback	2.22%	1.18%	0.06%	-1.06%	-2.26%

Note: This table shows the robustness of our implied consumption response to alternative parameterizations. See text for a description of baseline parameters and the policy experiment. The consumption gain in each row is equal to $\lambda \times 100$, where λ is the percentage change in consumption that makes a household indifferent between a variable and constant interest rate. In rows 2 and 3, we reduce the value of ϕ^h , in rows 4 and 5 we reduce the value of ϕ^y , and in row 6 we reduce the value of both elasticities. In the endogenous feedback model (row 7), we allow interest rate movements to cause movements in regional income and house prices. We calibrate this feedback so that a 25 basis point increase in interest rates lowers house prices by 0.40% and income by 0.20%.

A-1 Empirical Appendix

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

A-1.A Matching Procedure

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

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

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

A-1.B Robustness

Table A-2 presents a few additional robustness checks (discussed in Section IV.D of the text) on our main empirical findings. These results are suppressed in the main text for brevity, but are consistently similar to our main reported effects. The first column of Table A-2 shows our baseline estimates of the regression-discontinuity style effect of the difference in the relationship between interest rates and regional risk across the GSE and non-GSE markets and uses the same sample and same specification as described in the main text. Column 2 restricts our baseline sample to only those markets with at least 10 non-GSE loans in a given quarter. Column 3 includes a set of lagged FICO and lagged LTV controls, thus conditioning on both contemporaneous and prior composition of loan quality in each market. Across all of these additional specifications, the results are consistent in magnitude.

During the period of our analysis, 2001 to 2006, the conforming loan limit increased in each year, rising from \$276,000 to \$417,000. As an additional source of variation, we explore the relationship between interest rates and regional risk over the range of loans that switched status. This is one more way to control for loan size and focus solely on the difference between the pricing in the GSE and non-GSE markets. As Table A-3 shows, the relationship between interest rates and regional risk is quite strong in this range of loan amounts (\$276,000 to \$417,000), with the difference in coefficients ($\beta^{jumbo} - \beta^{GSE}$) = 14.14 (p-value < 0.001). This is nearly identical to the estimate of our main specification in row 1 of Table 3 (column 3).

A further issue is whether the intensive margin of lending varied significantly with regional risk; That is, whether lenders offered restricted loan-to-value ratios in riskier markets. In Figure A-1, panels a and b, we present evidence that shows no statistically significant relationship between LTV ratios and predicted regional risk as measured by lagged MSA default rates. If anything, these conditional correlations are slightly positive, suggesting that lenders allowed greater leverage in riskier markets for loans originated between 2001 and 2006 in both the GSE market (panel a) and non-GSE market (panel b).

A-1.C Points and Fees

In the text (see Section IV.D.4) we mention two additional robustness checks to explore whether the pricing of GSE loans differs by region based not on the guarantee fee (or “G-Fee”) of Fannie or Freddie, but if this is instead priced into the up-front costs paid by borrowers in the form of points or fees. In this section, we examine two additional datasets to establish that up-front pricing is not driving our results. First, we obtained additional data from one of the GSEs to directly estimate the relationship between effective interest rate and regional risk as measured by lagged GSE default. The measure of effective interest rate nets off any points and fees (including closing costs) charged

to the borrower (and is converted into the units of interest rates). As shown in Table A-4, we find no significant relationship between either interest rates or effective interest rates in the *universe* of GSE loans that meet our sample criteria. The two-standard-deviation effect is an insignificant 5 basis points for effective interest rates. In results not shown, no component of the effective interest rate (either points or fees) were found to be statistically associated with regional risk for the loans insured by this GSE.

Second, we examine rate quotes from LoanSifter, a firm that collects mortgage contract quotes across a range of U.S. markets and contract types. The advantage of this data is that the prices quoted are intended to be holding points constant (in our case, at zero). We use rate quotes from September 2009 through November 2010 across 57 metro areas, for quoted prices for a 30-year fixed-rate loan with a 20 percent downpayment (that is, 80% LTV ratio) and a 750 FICO score. As shown in Figure A-2, we observe no relationship between either unemployment (panel a) or house price growth (panel b) and quoted mortgage rate prices.

A-2 Computational Appendix

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

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

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:

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

and a household that chooses to sell its current house⁴¹ and rent has value function

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

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.

In order to implement the solution to this model numerically, we proceed as follows. First, we note that since $r_{k,j}^m > r \forall k, j$ it is never optimal to simultaneously hold both positive m and positive b . Thus, we can replace m and b with a single variable a . This financial asset variable is positive if $b > 0$ and negative if $m > 0$, and households face $r_{k,j}^m$ when $a < 0$ and r when $a > 0$. Thus, the household state reduces to $\mathbf{s}_{jk} = (a_j, h_j, z_j; \gamma_{jk})$.

In order to rectangularize the choice set and simplify the computational problems imposed by the endogenous liquidity constraint, we follow Diaz and Luengo-Prado (2010) in reformulating our problem in terms of voluntary equity. In particular, define $q_j \equiv a_j + (1 - \theta) \gamma_{k,j}^{\phi_h} h_j$. After substituting the budget constraint into the utility function to eliminate non-durable consumption as a choice variable, the value function can then be rewritten in terms of two endogenous variables q_j and h_j , the choice of which is restricted to be strictly positive. Note that $q_{j+1} \equiv a_{j+1} + (1 - \theta) \gamma_{k,j+1}^{\phi_h} h_{j+1}$ but that a_{j+1} and h_{j+1} are chosen in period j . Thus, shocks to house prices mean that voluntary equity realized at the start of period $j+1$ may differ from that chosen at the end of period j . Define $\widetilde{q_{j+1}} \geq 0$ to be the choice of voluntary equity for period $j+1$ made in period j . The state of realized voluntary equity relevant for $j+1$ then evolves as $q_{j+1} = \widetilde{q_{j+1}} + (1 - \theta) h_{j+1} (\gamma_{k,j+1}^{\phi_h} - \gamma_{k,j}^{\phi_h})$. This implies that although households are constrained to always choose $\widetilde{q_{j+1}} \geq 0$, actual voluntary equity can become negative if house prices fall by a large enough amount. To account for this, we solve the model for states that include negative voluntary equity even though households are constrained to choose non-negative values for this variable.⁴²

We discretize the problem so it can be solved on the computer by first discretizing γ and z using the algorithm of Tauchen (1986). We use 13 grid points for z and 5 grid points for γ . We then approximate $V_j^{adjust}(q_j, h_j, z_j, \gamma_j)$, $V_j^{noadjust}(q_j, h_j, z_j, \gamma_j)$, and $V_j^{rent}(q_j, h_j, z_j, \gamma_j)$ as multilinear functions in the endogenous states. In our benchmark calculation, we use 50 knot points for q_j (we space these points more closely together near the constraint) and 36 knot points for h_j . The presence of fixed adjustment costs on housing together with the borrowing constraint make the household policy function highly non-linear. For this reason, we follow Berger and Vavra (2014) and compute optimal policies for a given state-vector using a Nelder-Mead algorithm initialized from 3 different starting values, to reduce the problem of finding local maxima. The value of adjusting, not adjusting and renting are then compared to generate the overall policy function. We proceed via backward induction from the final period of life.

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

⁴¹If 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.

⁴²Shocks to house prices in the model are not large enough to ever reach a situation where q is so negative that households would be unable to choose $\widetilde{q_{j+1}} \geq 0$ without having negative consumption.

Table A-1: Descriptive Statistics of Matched Samples

Percentile	Interest Rate		FICO Score		LTV Ratio	
	Non-GSE	Matched GSE	Non-GSE	Matched GSE	Non-GSE	Matched GSE
10th	5.625	5.5	627	628	60	57
25th	5.99	5.875	637	638	70	70
50th	6.5	6.25	656	658	80	79
75th	7.125	6.875	698	700	85	85
90th	7.95	7.25	745	759	90	95
Mean	6.66	6.32	672	677	76.7	76.1
N	70327	70327	70327	70327	70327	70327

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

Table A-2: Additional Robustness Results

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

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

Table A-3: Focus on Loans Transitioning Around the Conforming Loan Limit

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


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

Table A-4: Relationship between Effective Interest Rates and Regional Risk


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

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

Figure A-1: Relationship between LTV Ratio and Regional Risk




(a) GSE Loans




(b) Non-GSE Loans

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

Figure A-2: Relationship between Interest Rates and Regional Risk, LoanSifter data



(a) Unemployment



(b) House Price Growth

Note: The figure uses MSA-level monthly data from September 2009 through November 2010 from LoanSifter to explore quoted mortgage prices with points and fees held constant at a level of zero. The interest rate measure is residualized to remove time fixed effects. The estimated relationships between interest rates and drivers of regional risk, unemployment (panel a) and house price growth (panel b) are statistically and economically insignificant. See text for details.