Regional Heterogeneity and Monetary Policy*

Martin Beraja Andreas Fuster Erik Hurst Joseph Vavra†

March 29, 2017

Abstract

We argue that the time-varying regional distribution of housing equity influences the aggregate consequences of monetary policy through its effects on mortgage refinancing. Using detailed loan-level data, we show that regional differences in housing equity affect refinancing and spending responses to interest rate cuts but that these effects vary over time with changes in the regional distribution of house price growth and unemployment. We then build a heterogeneous household model of refinancing and use it to explore the aggregate implications for monetary policy arising from our regional evidence. We find that the 2008 equity distribution made spending in depressed regions less responsive to interest rate cuts, thus dampening aggregate stimulus and increasing regional consumption inequality, whereas the opposite occurred in some earlier recessions. Taken together, our results strongly suggest that monetary policy makers should track the regional distribution of equity over time.

*First draft: May 2015. We thank Caitlin Gorback, Karen Shen and Eilidh Geddes for excellent research assistance. We would also like to thank Adrien Auclert, Wouter Den Haan, Daniel Greenwald, Arlene Wong, Mark Zandi and Junyi Zhu, as well as seminar participants at Chicago Booth, the University of Minnesota, NYU, MIT Sloan, Berkeley Haas, IIES Stockholm, Central Bank of Ireland, ECB Annual Research Conference, ASSA Chicago, Hutchins Center at Brookings and the CEPR–University of St. Gallen workshop on Household Finance and Economic Stability for helpful comments. The views expressed in this paper are solely those of the authors and not necessarily those of the Federal Reserve Bank of New York or the Federal Reserve System.

†Beraja: Princeton University and MIT. Fuster: Federal Reserve Bank of New York. Hurst and Vavra: University of Chicago Booth School of Business and NBER.

1 Introduction

Collateralized borrowing in the housing market can potentially play an important role in the monetary transmission mechanism, as interest rate declines encourage mortgage refinancing and home equity extraction to fund current consumption.[1] In this paper, we argue that the ability of declining interest rates to stimulate aggregate spending through this *refinancing channel* is time-varying because it depends crucially on the distribution of housing equity in the economy, which varies substantially across different recessions. Furthermore, we argue that interest rate cuts amplify regional consumption inequality through this channel when the hardest hit regions also experience the largest house price declines and therefore have little home equity.

Our analysis is motivated by striking differences in the regional distribution of house price growth over time. During the Great Recession, global house price declines left many households in the US and Europe with little home equity. Within these monetary unions, house price declines varied substantially across space. In particular, house prices fell the most during the late 2000s in the regions with the largest declines in economic activity (e.g., Nevada or Spain). In contrast, some earlier recessions exhibited very different patterns. For example, aggregate US house prices grew throughout the 2001 recession with little spatial variation.

Our paper proceeds in two parts. First, we use detailed microdata to explore the relationship between regional economic activity and refinancing in the US. We show that monetary policy during the Great Recession stimulated the most those regions with the smallest increases in unemployment and the smallest declines in house prices. In contrast, during the 2001 recession, when regional house price growth and unemployment were largely uncorrelated, refinancing activity was stronger in high-unemployment regions. Second, we build a heterogeneous household model of refinancing and use it to derive the aggregate implications of monetary policy from this regional evidence. When we calibrate our model to match economic conditions in 2008, we find that consumption responds less to interest rate cuts in depressed regions. This then dampens aggregate consumption stimulus and leads rate cuts to increase regional consumption inequality. However, under alternative distributions of house prices, such as in the 2001 recession, we show that interest rate cuts can both better stimulate aggregate consumption and reduce (rather than amplify) regional consumption inequality. Since the distribution of equity both varies across time and changes the consequences of monetary policy, we conclude that it is important for policy makers to track this variation when making decisions.

In more detail, the first half of our paper provides empirical evidence about the effects of regional heterogeneity in housing equity on the refinancing channel of monetary policy. We present two complementary sets of facts that exploit different sources of variation. The first set explores the regional response in the US to interest rate declines following the first round of the Federal Reserve's large-scale asset purchase program—commonly known as quantitative easing (QE1). QE1 provides a unique opportunity to study the refinancing channel and its interaction with the distribution of housing equity in the economy because of both its magnitude—30-year fixed-rate mortgage rates fell by around 1% in the month after announcement—and the large variation in regional housing market conditions at

[1] See e.g. https://www.federalreserve.gov/boarddocs/hh/2004/february/testimony.htm and https://www.newyorkfed.org/newsevents/speeches/2012/dud120106.html for recent policy discussion of this channel.

the time of the announcement. Because the US is a monetary union where interest rates are identical across regions and interest rate policy is independent of regional variation in economic activity, we can also explore how refinancing and regional economic activity respond to rate cuts in other periods.[2] In our second set of facts, we indeed move away from this particular QE1 episode and compare regional house prices, economic activity, and refinancing patterns across different recessions in the US.

Using a variety of loan-level data sources, we document three facts regarding the regional response to QE1. First, there was a boom in household mortgage refinancing right after the QE1 announcement. Second, refinancing activity and the amount of equity extracted increased more in metropolitan statistical areas (MSAs) that had lower unemployment and where homeowners had more housing equity on the eve of QE1. Specifically, very little refinancing occurred in places like Las Vegas and Phoenix, where most homeowners were underwater at the time QE1 was implemented. Third, MSAs where homeowners refinanced the most right after QE1 also experienced the largest resulting increase in consumption, as measured by car purchases. Moreover, individual households that refinanced increased their spending sharply, an increase that is even larger for households that removed equity by cashing out when refinancing. Overall, these facts show that the refinancing channel was weakest during the Great Recession in regions with the worst housing and labor market conditions, reducing monetary policy effects in these locations and thus increasing regional consumption inequality.

Our second set of empirical results moves beyond QE1 to provide evidence that the consequences of monetary policy vary substantially over time. We begin by showing that there is large variation in the distribution of regional house price growth, and therefore regional home equity, in the US over time. For example, in the recessions in 1990 and 2008, there were large aggregate declines in house prices, that varied substantially across regions. Regional declines in house prices were also highly correlated with increases in regional unemployment in these recessions. In contrast, aggregate house prices grew throughout the 2001 recession. There was also little variation in this house price growth across regions and little correlation between regional house price growth and unemployment. We also show that refinancing patterns differ across time. For example, aggregate refinancing propensities were much higher during the 2001 recession than during the 2008 recession. These recessions also exhibit different relationships between unemployment and refinancing: propensities increased more in low-unemployment than in high-unemployment MSAs during the Great Recession, whereas the opposite was true during the 2001 recession. This time-series finding complements the evidence from our QE1 event-study that regional refinancing responses depend crucially on regional economic conditions.

Next, we ask: what does this regional evidence imply for the aggregate consequences of monetary policy? Answering this question without a theoretical model is challenging. First, many features of the regional equity distribution move over time. With only a small number of recessions, it is essentially impossible to determine directly from the data which particular features of this distribution determine the strength of the refinancing channel of monetary policy. Second, echoing ideas in Beraja, Hurst, and Ospina (2016), making inferences about aggregates from regional evidence requires accounting for the offsetting effects of refinancing activity on the behavior of lenders in the economy, which cannot be directly measured in our data. Analyzing these issues thus requires a formal model in order to conduct counterfactuals that cannot be computed in our micro-data alone.

[2] See Hurst et al. (2016) for empirical evidence of regional rate equilization for mortgages.

Hence, in the second part of the paper we build an equilibrium, incomplete-markets, heterogenous agents model with both mortgage borrowers and lenders. The key model features are that mortgage borrowers face both idiosyncratic and regional income and house price risk, and they can refinance their mortgage and extract housing equity by paying a fixed cost. The model is rich enough to capture key aspects of the data developed in the first part of the paper. In addition, the model clarifies how the transmission from interest rate cuts to aggregate spending through the refinancing channel depends crucially on the distribution of housing equity. Finally, the model highlights the conditions under which monetary policy makers are likely to face a trade-off between stimulating aggregate spending and increasing regional consumption inequality. Through a series of counterfactual exercises, we show that the effects of interest rate changes on both aggregate consumption and regional consumption inequality change dramatically with the distribution of housing equity in the economy.

Our first theoretical results focus on the consequences of interest rate cuts in a benchmark economy that matches the joint distribution of housing equity and income observed in 2008. We pick the baseline parameters in our quantitative model to match the cross-region effects of QE1 documented in the first part of the paper and then compute the aggregate effects of this policy. We find that a decline in interest rates of the magnitude observed after QE1 modestly raises aggregate spending. This implies that the spending offset coming from lenders in equilibrium is not one-for-one. This is because our model features an important role for cash-out activity in determining spending. Households accumulate equity over time and periodically pay a refinancing cost to access this equity. Furthermore, since borrowers are more liquidity constrained than lenders, equity extraction increases spending on net. When interest rates decline, refinancing and equity extraction are accelerated and aggregate spending rises. However, under 2008 economic conditions, this aggregate spending effect is quantitatively small. After documenting the aggregate effects of rate cuts, we next show that, under 2008 economic conditions, they lead to a large increase in consumption inequality across regions. High-equity homeowners have stronger refinancing and spending responses to declines in interest rates than low-equity homeowners. Since income and equity were highly correlated in 2008, the variance of consumption across regions then rises. Thus, our model implies that monetary policy in 2008 modestly boosted aggregate spending but did so at the cost of much larger consumption inequality.

While it is interesting to model the economy's response to monetary policy during the Great Recession, the bulk of our analysis moves beyond this particular episode. In particular, we look at the consequences of interest rate cuts in counterfactual economies with alternative distributions of equity and income across regions. We find that as the average level of housing equity in the economy rises, both aggregate spending and regional consumption inequality respond more strongly to rate cuts. However, the spending response grows more rapidly with house prices and equity than the inequality response. This means that the trade-off between stimulus and inequality is not as bad in recessions with house price increases, such as 2001, as it is in recessions with house price declines, such as 2008. Conversely, a reduction in the variance of house price growth across regions leads to declines in both spending and inequality effects of monetary policy. However, reducing the variance of the equity distribution reduces inequality effects more than it reduces spending effects. This means that the trade-off between stimulus and inequality is not as bad in recessions with small regional house price variance, such as 2001, as it is in recessions with large variance such as 2008. Finally,

we show that when the correlation between house prices and income is low, interest rate declines no longer increase inequality. However, this correlation has almost no effect on the aggregate spending consequences of monetary policy. This means that in recessions such as 2001, with little correlation between house prices and income, monetary policy may face no trade-off between stimulating the economy and increasing inequality.

It is notable that each of these effects makes the trade-off between stimulus and inequality worse during the Great Recession, since this period was characterized by large declines in house prices, large regional variance in house prices, and a high correlation between house prices and income. However, this trade-off between stimulus and inequality does not hold in general and varies across time. During other periods, such as 2001, there was likely no trade-off at all. While we think these time-varying trade-offs are interesting, it is also important to note that even under a policy mandate that places no weight on reducing inequality, one must still pay attention to the regional distribution of collateral, since it affects the stimulative power of monetary policy.

Why do changes in the equity distribution alter the economy's response to interest rate cuts? These effects arise from the non-linear interaction between household equity and refinancing decisions. Because households must satisfy a collateral requirement in order to refinance, when interest rates fall, those with substantial equity can reduce their interest payments while also extracting equity whereas those currently underwater must put up additional cash. Hence, when interest rates fall, many households with positive equity refinance and increase their consumption commensurately with extracted equity, whereas almost no households with negative equity do. This leads to a consumption response to interest rate cuts that is highly convex in equity because households that are mildly underwater exhibit the same zero response as those substantially underwater, whereas households with substantial positive equity exhibit much stronger consumption changes than those with mildly positive equity. This convexity explains why changing the distribution of equity affects the economy's response to interest rate declines. For example, an aggregate decline in house prices that reduces all individuals' equity proportionately reduces the consumption response of households with initially high equity but leaves the response of those with low equity unchanged.

After arguing that the equity distribution in 2008 hampered monetary policy's ability to stimulate the economy through the refinancing channel, we show that various policies can complement interest rate cuts to increase monetary policy's power. In particular, we show that both targeted debt reduction and relaxation of collateral constraints for refinancing can amplify the stimulative effects of monetary policy and also reduce the trade-off with inequality. Policies along these lines were implemented during the Great Recession through the Home Affordable Modification and Refinance Programs (HAMP and HARP), and our results show that such mortgage market interventions can interact importantly with monetary policy. In addition, we explore the role of macroprudential policy and show that time-varying countercyclical leverage requirements have the potential to both dampen the depth of house-price-induced recessions and strengthen the stimulative effects of monetary policy.

Finally, we show that our model implications continue to hold under many alternative assumptions. In these model extensions we account for the presence of adjustable-rate mortgages, calibrated to match the observed regional share in the data; allow for the fact that in this recession large busts were preceded by large booms; allow for cash-out and non-cash-out refinancing; extend our baseline

environment with unanticipated one-time interest rate shocks to environments with stochastic, transitory rate changes; show that our results are insensitive to assumptions about short-long interest rate spreads; and explore alternative assumptions on the lender side of the economy about the share of mortgage debt ultimately held by domestic consumers.

While we focus on the role of the distribution of equity across US regions for the refinancing channel of monetary policy, we believe our results apply more broadly. We concentrate on the regional dimension for housing equity because shocks to housing collateral values have a very large regional component. However, variation in the distribution of other types of collateral at even more disaggregated levels can potentially generate many of the same implications for monetary policy. For example, industry-level shocks may change the distribution of collateral across firms and affect the response of investment to monetary policy. Our conclusions also extend beyond the US. The last decade has given rise to persistent variation in economic activity across countries within Europe, and this activity has been strongly correlated with national house price growth. While institutional features of mortgage markets differ across countries, our results suggest that the distribution of house price growth in Europe may have produced similar challenges for monetary policy during this time period.

2 Related Literature

Our work is related to much existing research. First, a vast New Keynesian literature emphasizes intertemporal substitution as the main reason why interest rate changes affect household spending.[3] We also emphasize spending responses to interest rate changes, but depart from the standard New Keynesian assumption of frictionless household capital markets with one-period borrowing. In reality, the bulk of household borrowing occurs through the mortgage market, which features collateral requirements and long-term fixed nominal payments that can only be refinanced at some cost. Together, these features give rise to what we call the "refinancing channel" of monetary policy.

Importantly, the strength of this channel depends on the distribution of housing equity in the economy, which exhibits substantial time variation in the data. This means that policy makers must pay attention to this distribution when determining the rate necessary to achieve a given level of stimulus. We thus contribute a new channel to the growing literature arguing that the economy exhibits time-varying responses to aggregate shocks, which depend on the microeconomic distribution of agents in the economy.[4] Most closely related is Berger et al. (2015) who argue that changes in the distribution of household leverage during the housing boom contributed to the large decline in spending when house prices subsequently crashed. Interestingly, we show here that these same leverage patterns then hampered monetary policy's ability to stimulate the economy in response.

We are not the first to study the transmission of monetary policy through redistribution in the mortgage market. For example, on the theoretical side, Rubio (2011), Garriga, Kydland, and Sustek (2013) and Greenwald (2016) also study this channel. However, they assume a representative borrower and generally abstract from the costs of refinancing, in contrast to our environment, which accounts for

[3] See Woodford (2003) and Galí (2009) for canonical expositions.

[4] For example, Vavra (2014) and Berger and Vavra (2016) argue that the time-varying distribution of price changes matters for monetary policy, Caballero and Engel (1999) and Winberry (2016) argue that the distribution of capital matters for aggregate investment, and Berger and Vavra (2015) arrive at similar implications for durable spending.

heterogeneity, incomplete markets, and fixed costs of refinancing. This means that their models have no role for the distribution of housing equity across borrowers, which is the focus of our paper. Wong (2016) uses a model closer to our own, which includes borrower heterogeneity and costly refinancing but in partial equilibrium, and she focuses on how aging affects monetary policy. Since the age distribution changes slowly across time, age effects are more relevant for cross-country comparisons and long-run trends than for shorter-run changes in the refinancing channel of monetary policy.

Our focus on the implications of realistic modeling of household borrowing and how it interacts with heterogeneity in the economy also parallels many of the themes in Auclert (2015), who argues that the covariance of the marginal propensity to consume with interest rate exposure across agents matters for aggregate consumption responses to interest rate changes. His analysis abstracts from refinancing. We show that refinancing frictions combined with cash-out refinancing lead to an important role for the time-varying distribution of housing equity.

On the empirical front, Fuster and Willen (2010) measure the effects of QE1 on the primary US mortgage market. They emphasize differential effects on borrowers with different creditworthiness, while we emphasize regional disparities.[5] Chen, Michaux, and Roussanov (2013) investigate the link between macroeconomic uncertainty and cash-out refinancing while Bhutta and Keys (2016) show that low interest rates increase the likelihood and magnitude of home equity extraction. Calza, Monacelli, and Stracca (2013) document the importance of variation in mortgage structure for monetary policy across countries. Di Maggio, Kermani, and Ramcharan (2014) and Keys et al. (2014) study the effects of ARM resets on durable consumption, following work by Fuster and Willen (2016) and Tracy and Wright (2016) that studies the effects of resets on mortgage defaults. Di Maggio, Kermani, and Palmer (2016) also study the response of refinancing to quantitative easing efforts but focus on the distinction between conforming and non-conforming loans and on changes across time in the composition of large-scale asset purchases. Agarwal et al. (2015) use data from the Home Affordable Refinancing Program (HARP) to provide evidence that refinancing spurs spending and that this channel was strengthened by the program's reduction of collateral frictions. Fratantoni and Schuh (2003) use a heterogeneous-agent VAR with regional heterogeneity in housing markets to study time variation in monetary policy passthrough. Our empirical patterns in the QE1 episode are similar to those documented by Caplin, Freeman, and Tracy (1997) for the 1990 recession based on mortgage data from a single bank. We use more representative data over a longer time period and present a model that allows us to analyze which features of the regional equity distribution matter for aggregate policy, and to conduct counterfactuals. Our results show that the refinancing channel of monetary policy varies substantially across different recessions.

Finally, a large literature studies a "credit channel" of monetary policy, where changes in collateral values amplify output responses to rate changes.[6] This channel is complementary but distinct from ours, as it arises from monetary policy changing collateral values which, in turn, affect economic

[5] There is a growing literature specifically studying the effects of QE, but focused primarily on financial market reactions; see, for instance, Gagnon et al. (2011); Hancock and Passmore (2011); Krishnamurthy and Vissing-Jorgensen (2011, 2013); Stroebel and Taylor (2012). Chen, Cúrdia, and Ferrero (2012) study the effects of QE through the lens of a DSGE model.

[6] For example, Iacoviello (2005) shows that adding collateral constraints on housing to a financial accelerator model like that in Bernanke, Gertler, and Gilchrist (1999) amplifies the effects of rate changes. See also the related literature on the "balance-sheet" channel of monetary policy, e.g., Gertler and Karadi (2011).

activity. In contrast, we take the distribution of collateral at a point in time as given and show that it affects the transmission from interest rates to spending. We think both channels are important and exploring their interactions is an interesting area for future work.

3 Data

We use a variety of different data sources in our empirical work, which we briefly describe here. The Online Appendix provides additional details.

Our main measures of local refinancing activity come from Equifax's Credit Risk Insight Servicing McDash (CRISM) data set. This data set merges McDash mortgage servicing records (from Black Knight Financial Services) with credit bureau data (from Equifax) and is available beginning in 2005. The structure of the data set makes it possible to link multiple loans by the same borrower together, something that is not possible with mortgage servicing data alone. This allows us to measure refinancing activity much more accurately than what can be achieved with previously available data. Since we know both the outstanding amount of the old loan (as well as any second liens that get paid off around the same time) and the principal amount of the new loan, we can measure the dollar amount of equity that is removed (or "cashed out") from the home during refinancing. CRISM covers roughly two-thirds of the US mortgage market during the period we study.

We also use CRISM data to estimate borrowers' home equity. We define home equity as one minus the household's combined loan-to-value (CLTV) ratio, which we estimate for each household by adding balances of first mortgages and potential second liens and dividing by the estimated property value. To create estimates of the property value, we use the property's appraisal value at the time of mortgage origination and update it using location-specific house price indices from CoreLogic.[7] We use estimated equity as a measure of the ability and incentive of local homeowners to refinance. Our preferred summary statistic of equity conditions in a location is the equity of the median borrower in that location where we weight borrowers by their outstanding mortgage balances. We henceforth denote this statistic as $E^{med}_{j,t}$ which varies across MSAs j and time t.[8] We put particular emphasis on $E^{med}_{j,Nov2008}$, the median equity as of November 2008, just prior to the onset of QE1.

We supplement our analysis of refinancing activity using data from the Home Mortgage Disclosure Act (HMDA). For each application, HMDA reports the geographic location of the property, the desired loan amount, the loan purpose (purchase or refinance), and whether the loan application led to an origination, was rejected by the lender, or was withdrawn by the borrower. HMDA data has the benefit of broader coverage than CRISM data; it also covers a longer time period, which allows us to extend our analysis prior to the 2008 recession. However, it does not have information on equity removed during the refinancing process. In the Online Appendix, we show that regional patterns in

[7] We use zip code indices if available, and MSA-level indices if not. Additional details are provided in the Data Appendix.

[8] We have also conducted our analysis instead using measures of the fractions of borrowers with a CLTV above 0.8 or above 1, thresholds above which it becomes expensive and difficult (or even impossible) to refinance. As we show in Appendix Figure A-1 , $E^{med}_{j,t}$ is strongly is correlated with the fraction of mortgage holders in an MSA above these thresholds. As a result, performing our empirical work with these alternative measures of local housing equity produced very similar results. We prefer the median equity measure, since it is slightly easier to interpret and corresponds more closely to the quantities emphasized in our theoretical model later in the paper.

refinancing following QE1 are nearly identical when using HMDA data instead of CRISM.

Finally, we supplement our borrower-level data with a variety of additional data sources. First, we show aggregate refinancing trends within the US using published statistics from the Mortgage Bankers Association's (MBA) Refinance Index. Second, as alluded to above, we use house price data from CoreLogic at the zip code or MSA level to measure local house price appreciation. Third, we measure unemployment rates for each MSA using data from the BLS's Local Area Unemployment Statistics. Finally, as our measure of local spending we use data from R.L. Polk on new car purchases at the MSA level (aggregated from new auto registrations at the zip code level).

4 The Refinancing Channel and Regional Heterogeneity: Evidence from QE1

This section documents several facts relating regional heterogeneity in housing equity to the refinancing channel. In this section, we use an event-study methodology that exploits regional variation in responses to the interest rate decline that followed QE1. We begin with a brief description of aggregate mortgage activity around QE1 in order to provide some background. Then, we proceed to document our main findings, which can be summarized as follows:

1. Mortgage originations increased substantially after QE1 because households refinanced their existing mortgages.

2. Refinancing activity and equity extraction were higher in MSAs where homeowners had more equity (which were also locations where unemployment was lower) prior to QE1.

3. Car purchases increased the most after QE1 for individuals who removed equity when refinancing and in MSAs with the largest refinancing response.

4.1 Aggregate Trends in Mortgage Activity Around QE1

Figure 1 shows time-series patterns in the monthly MBA Refinance Index over 2000 to 2012 (solid line). The figure also includes the difference between the average 30-year fixed-rate mortgage (FRM) rate (also from MBA) in month t and the average of the 30-year mortgage rate over the prior five years (dashed line). Negative values of this metric mean that mortgage rates in a given month are low relative to what they had been over the previous years, giving many mortgage borrowers an incentive to refinance. A few things are noticeable from Figure 1. First, there is a very strong relationship between refinancing activity and mortgage rates. The simple correlation between the two series is -0.77. When mortgage rates fall relative to the average over the prior few years, refinancing activity increases. Second, mortgage rates fell sharply and refinancing activity expanded sharply when QE1 was announced in late November 2008, marked as a vertical line in the figure. The refinancing boom from December 2008 through April 2009 was larger than in any period since mid-2003. Finally, we note that since refinancing activity increases essentially whenever mortgage rates fall, the applicability of our results extends to any period in which Federal Reserve policy moves mortgage rates. We primarily

Figure 1: Mortgage Refinancing Activity in the US over 2000-2012

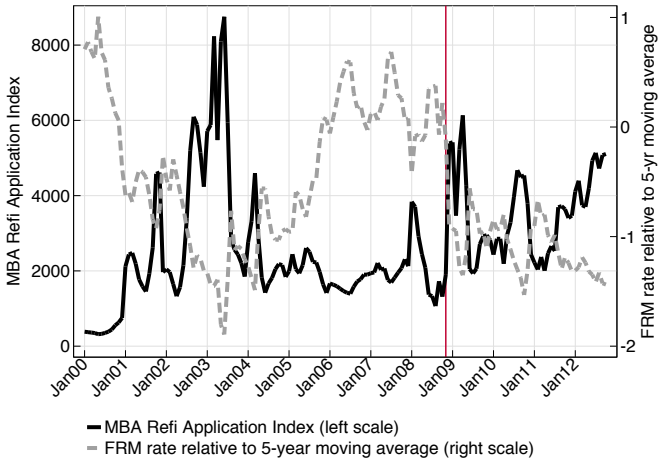


Figure shows monthly average of Mortgage Bankers Association (MBA) Refinancing Index (seasonally adjusted; March 1990=100) and the 30-year fixed-rate mortgage rate (relative to 5-year moving average), also from MBA.

focus on the QE1 announcement simply because it was largely unexpected and led to such a sharp drop in mortgage rates.[9]

4.2 Regional Variation in Equity Distributions Prior to QE1

To explore regional variation in loan origination around QE1, we need to define our notion of regions. Throughout our paper, we use metropolitan statistical areas (MSAs). We begin by showing that equity distributions evolved very differently across MSAs between 2007 and 2008.

Figure 2 shows the distribution of household housing equity in two different time periods for five MSAs: Chicago, Las Vegas, Miami, Philadelphia, and Seattle. We pick these MSAs to show examples of MSAs that had housing price declines between 2007 and 2008 that were large (Miami and Las Vegas), medium (Chicago), and small (Philadelphia and Seattle).[10] Panel (a) shows the housing equity distribution for these five MSAs in January 2007. We choose January 2007 because it is a period prior to the nationwide house price decline. For all five MSAs, housing equity distributions are quite similar. As noted above, we often summarize the distribution in each MSA using the median household equity within that MSA during the period, $E^{med}_{j,t}$. In January 2007, the median household in Chicago, Miami, Philadelphia and Seattle had housing equity worth between 30 and 40 percent of its house value. The median household in Las Vegas had equity that was roughly 23 percent of its house value. Very few households in any of these MSAs had negative equity as of January 2007.

By November 2008, when QE1 was announced, there was large variation in the equity distribution across MSAs. This can be seen in panel (b) of Figure 2. The median household in Las Vegas had negative equity of roughly 17 percent of their house value, while the median household in Miami had zero equity. Conversely, the median households in Philadelphia and Seattle still had home equity

[9] In Appendix Figure A-2, we show that monthly mortgage origination activity in HMDA data over the same period displays very similar patterns. HMDA data also shows that from December 2008 through early 2009 the vast majority of originations were due to refinancings. There were very few mortgage originations for new home purchases during this time period, partly because of seasonality. For this reason, we focus our analysis in this paper on refinancings.

[10] Table A-1 in the Online Appendix shows descriptive statistics for all 381 MSAs in our analysis.

Figure 2: Distributions of Borrowers' Equity in their Homes across 5 MSAs

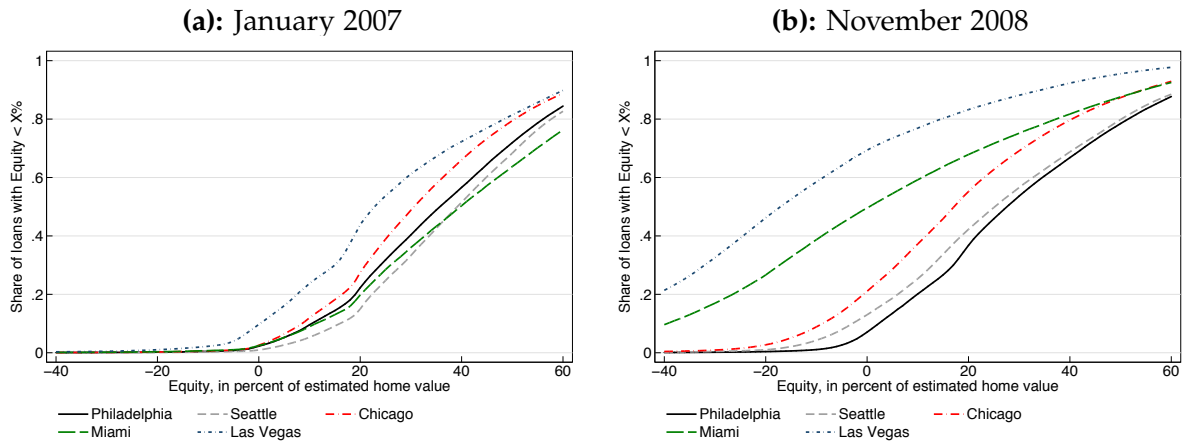


Figure shows the cumulative distribution of borrower equity in five illustrative MSAs in January 2007 (panel a) and November 2008 (panel b). Equity is measured for each household in an MSA using CRISM data on the estimated current house value minus total current mortgage debt, divided by an estimate of their current house value (i.e., equity = 1−CLTV). Distributions are weighted by mortgage balance.

that was roughly 25 to 30 percent of their house value. Between early 2007 and late 2008, the equity distribution in places like Las Vegas and Miami shifted dramatically to the left relative to places like Philadelphia and Seattle. The equity of the median household correlates strongly with other measures of the equity distribution (see also Figure A-1 in the Online Appendix). For example, roughly 50 percent of households in Miami and 70 percent of households in Las Vegas had negative housing equity in November 2008, while only 6 to 10 percent had negative equity in Philadelphia and Seattle.

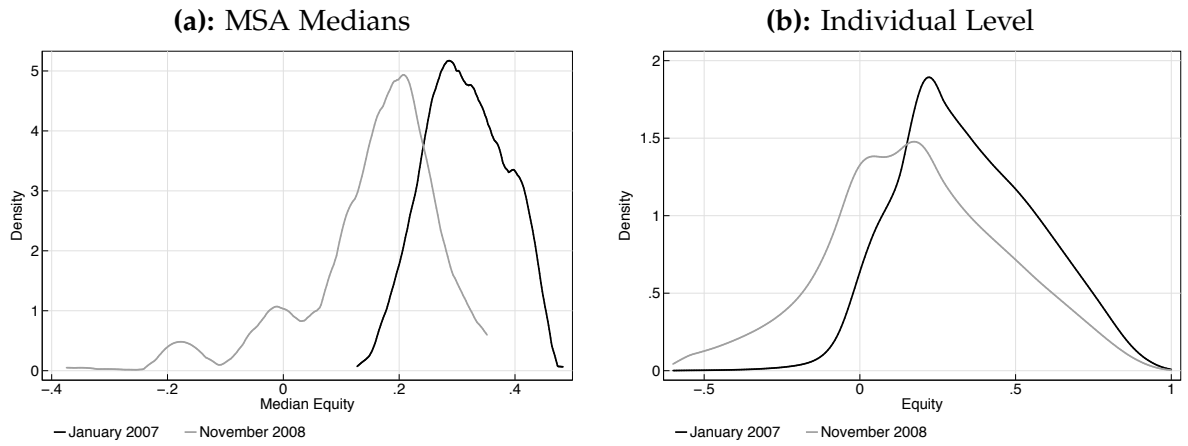
Panel (a) of Figure 3 shows the distribution of $E^{med}_{j,t}$ across MSAs (weighted by MSA population) in January 2007 and November 2008 to highlight that the MSAs we included in Figure 2 are representative of the full distribution. On average, equity declined sharply between early 2007 and late 2008 but these declines were not uniform across MSAs. Thus, the distribution of median household equity in an MSA shifted sharply left and fanned out over this period. Panel (b) shows that patterns are similar for the distribution of individual equity rather than $E^{med}_{j,t}$ to illustrate that focusing on median equity is not essential for our conclusions.[11]

Appendix Figure A-3 shows the relationship between equity, unemployment changes, and house price growth during the period January 2007 through November 2008.[12] Over this period, differential house price declines across MSAs were the main driver of differences in $E^{med}_{j,t}$. On average, a 10 percent decrease in house prices from January 2007 to November 2008 is associated with an 8.3 percentage point lower $E^{med}_{j,Nov2008}$. This relationship will prove useful when we extend our analysis to pre-2005 periods for which we do not have measures of median housing equity at local levels. When examining longer time series patterns, we will use regional variation in house price growth as a proxy for regional variation in housing equity. Additionally, Appendix Figure A-3 documents that MSAs with the largest

[11] This individual distribution is also more easily mapped to our model results. Individual equity is much more dispersed than median equity at the MSA level because individual variation has life-cycle and other idiosyncratic components that are large relative to cross-region variation.

[12] Appendix Table A-2 further documents the correlation of median equity with other characteristics of the stock of outstanding mortgages across locations.

Figure 3: Distributions of Borrowers' Equity in their Homes — MSA Medians and Individual Level



Panel (a) shows kernel density of $E^{med}_{j,t}$ across 381 MSAs in January 2007 and November 2008; MSAs are weighted by their 2008 population. Panel (b) shows kernel density of individual borrower equity in January 2007 and November 2008; borrowers are weighted by loan amount.

increases in local unemployment rates also had the lowest $E^{med}_{j,Nov2008}$. This is unsurprising since the literature has shown that house price declines were associated with weakening labor markets during this period (e.g., Charles, Hurst, and Notowidigdo, 2013; Mian and Sufi, 2014), but it is important for interpreting the inequality effects of monetary policy, since we will now show that refinancing activity responded least to QE1 in the locations with the least home equity.

4.3 Regional Variation in Mortgage Activity Around QE1

We now show that in the months after QE1, refinancing activity was much higher in regions where individuals had sufficient equity in their home and the unemployment rate was relatively low. To facilitate the exposition of our results, we divide all MSAs into quartiles based on $E^{med}_{j,Nov2008}$.[13] Figure 4 shows refinancing activity over time for MSAs in the top and bottom quartiles of $E^{med}_{j,Nov2008}$. The bottom quartile of $E^{med}_{j,Nov2008}$ includes MSAs like Las Vegas where the median mortgage borrower was underwater. The top $E^{med}_{j,Nov2008}$ quartile includes MSAs like Seattle where most borrowers had sufficient equity to refinance.

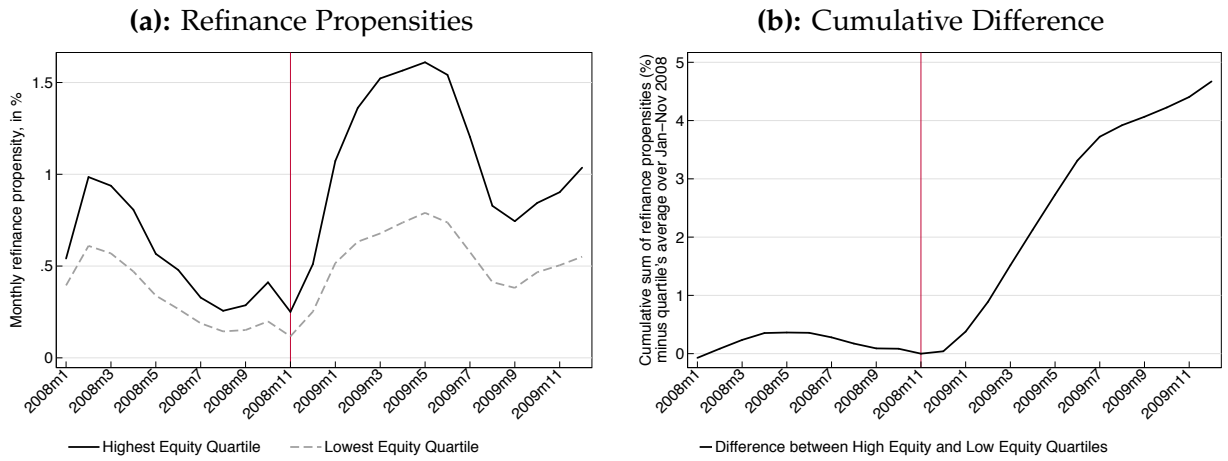
Panel (a) shows monthly refinancing propensities from January 2008 through December 2009. Refinancing propensities are higher throughout in the high equity quartile, but they evolve similarly between high and low equity MSAs up to November 2008. After QE1, refinancing activity jumped—but it jumped much more in the high equity MSAs relative to the low equity MSAs.[14]

Panel (b) shows the cumulative difference between the two groups, after subtracting each group's average refinancing propensity from January to November 2008 to remove the initial level difference.

[13] Quartiles are population-weighted using 2008 numbers from the Census. This ensures that there are the same number of people within each quartile. Appendix A.1 lists the specific MSAs within each of the $E^{med}_{j,Nov2008}$ quartiles.

[14] The jump happens in January/February (rather than December) because CRISM measures originations, not applications, and there is a delay of 1-3 months between when a mortgage application is initially made and when the actual mortgage origination takes place. In Appendix Figure A-4, we use HMDA data with exact application dates to show that applications jumped immediately after the announcement of QE1, and more so in high equity MSAs.

Figure 4: Mortgage Refinance Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by Median Borrower Equity in November 2008



Panel (a) shows monthly refinance propensities in CRISM, defined as the dollar amount of refinance mortgage originations divided by outstanding mortgage amounts in the prior month. Calculations are done over MSA quartile groups for the highest and lowest $E^{med}_{Nov2008}$ quartiles. Panel (b) shows the cumulative difference between the two groups, after subtracting each group's average refinancing propensity from January to November 2008. Vertical lines show the month of the QE1 announcement (November 2008).

Prior to QE1, the cumulative difference is essentially flat at zero, reflecting the parallel pre-trend in panel (a). After QE1, a sharp difference emerges, eventually leading to a total increase in refinancing propensity about 4 percentage points larger in the high equity MSA group than in the low equity MSA group. This is a substantial difference, since the cumulative refinancing propensity in the low equity group is only 7 percent over the entire year 2009.

We complement these figures with a difference-in-differences style regression which allows us to control for additional local factors. Specifically, we estimate the following:

$$Refi_{j,t} = \alpha_j + \alpha_t + \beta(E^{med}_{j,Nov2008} \times postQE) + \Gamma(X_{j,Nov2008} \times postQE) + \varepsilon_{j,t}, \quad (1)$$

where $Refi_{j,t}$ is the monthly refinancing propensity in each MSA over the six months prior to QE1 and the six months after QE1, α_j and α_t are MSA and time fixed effects, respectively, and $postQE$ is an indicator variable that equals one for the six months after QE1. We use February 2009 as the start of the post-QE-announcement period given the lag between mortgage applications and originations discussed above. $X_{j,Nov2008}$ is a vector of potential local controls, including the change in the unemployment rate between January 2007 and November 2008 and measures of loan characteristics (e.g., average FICO score, ARM share, and GSE share) in November 2008. To conserve space, we show results in Appendix Table A-3. A few key takeaways are worth mentioning. First, in all specifications β is positive and highly statistically significant. This reinforces the findings in Figure 4. Additionally, Appendix Table A-3 shows that our findings are robust to the inclusion of controls for local economic conditions and loan characteristics.[15] Adding the average FICO score of mortgage borrowers (inter-

[15]One might also worry that mortgage rates may vary across regions with different equity, and that these rate differences might explain different refinancing patterns. However Hurst et al. (2016) shows that rates on new conforming loans, which

acted with *postQE*) to the regression reduces the coefficient on equity by almost half. But it is likely that average FICO scores are themselves affected by the local equity distribution (since underwater borrowers are more likely to default); therefore, we view the fact that equity by itself remains strongly significant as underscoring its importance in explaining differences in refinancing. The same is true for the last column in the table, where we add all additional controls at once. A linear combination of these variables is very highly correlated with median equity (see the last column of Table A-2); nevertheless, equity remains individually significant. One additional result worth mentioning is that, even conditional on local housing equity levels, MSAs with larger increases in unemployment between January 2007 and November 2008 saw smaller increases in refinancing. These results show that both local labor market conditions and the local equity distribution determine local refinancing propensities.

Collectively, the results from Figure 4 and Table A-3 show there were large regional differences in refinancing activity in response to QE1. Regions with the largest house price declines and thus the least equity were the least responsive to QE1 in terms of subsequent mortgage refinancing activity.

4.4 Regional Variation in Equity Extraction and Spending Around QE1

To what extent do these spatial differences in refinancing activity lead to differences in spending? Unfortunately, data on local spending are extremely limited, but we provide evidence on this front in two ways. First, we explore the extent to which households removed equity from their home at the time of refinancing. Prior research has shown that households, on average, spend a large amount of the equity they remove during the refinancing process on current consumption and home improvements.[16] Second, as described above, we have data on new car purchases at the MSA level from R.L. Polk. This data has been used recently to measure spending at the local level.[17]

Figure 5 shows the amount of equity removed during refinancing for the top and bottom quartile MSAs by $E^{med}_{j,Nov2008}$. Panel (a) shows estimated dollar amounts per month, while panel (b) shows the cumulative difference between the two groups, after subtracting each group's average cash-out amounts from January to November 2008.[18] The total amount of equity removed during the refinancing process sums over people who removed no equity, those who put equity into their home, and those who extracted equity. On net, borrowers remove equity during the refinancing process in both high and low equity locations. At all points in time there is more equity removal in high equity locations, but trends evolve similarly prior to the QE1 announcement. After QE1, equity removal increases substantially in the high $E^{med}_{j,Nov2008}$ locations relative to the low $E^{med}_{j,Nov2008}$ locations.

Summing across all MSAs in the top equity quartile, about $23.7 billion of equity was cashed out during the refinancing that took place in the six months after QE1 (January-June). Conversely, for the MSAs in the bottom equity quartile, only $10.9 billion of equity was cashed out. As panel (b) shows, the cumulative difference in cash-out amounts over 2009 between the two MSA groups—after

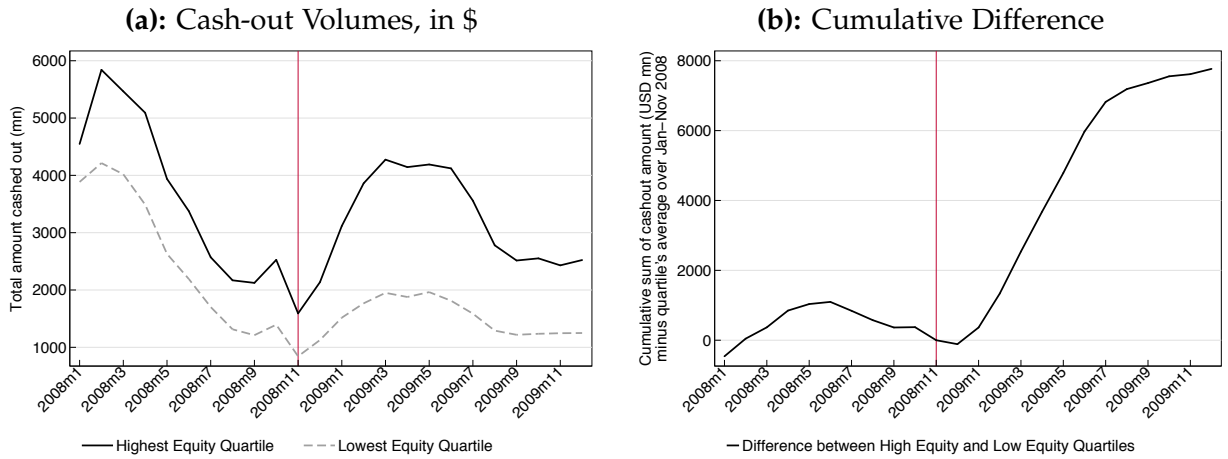
make up almost the entire market in this time period, are equalized across space after controlling for borrower observables. Table A-3 also controls directly for previous outstanding rates.

[16] See Brady, Canner, and Maki (2000), Canner, Dynan, and Passmore (2002), Hurst and Stafford (2004) and Bhutta and Keys (2016).

[17] See, for example, Mian, Rao, and Sufi (2013).

[18] Since the CRISM data does not cover the whole mortgage market, we scale up dollar amounts in CRISM for this figure; see Appendix A.2.3 for details.

Figure 5: Cash-Out Refinancing in Top and Bottom Quartile of MSAs by Median Borrower Equity in November 2008



Panel (a) shows total cash extracted during refinancing in the top and bottom MSA quartiles by $E^{med}_{j,Nov2008}$. Panel (b) shows the cumulative difference between the two groups, after subtracting each group's average cash-out amounts from January to November 2008. Vertical lines show the month of the QE1 announcement (November 2008).

subtracting each group's pre-QE averages—amounted to about $8 billion. Again, this cumulative difference shows the sharp break around QE1. Using BEA data, this differential cash-out adds up to approximately 9% of the total differential change in spending across these groups between 2008 and 2009.[19] This effect is both large and similar to the size of effects in our model below, in which monetary policy explains 10% of observed spending differences between high and low equity regions.

In Appendix Table A-4, we show results from a regression similar to equation (1) above but with monthly equity removed (relative to outstanding balance) as the dependent variable. We refer to monthly equity removed relative to outstanding balance as the cash-out fraction. Echoing the results in Figure 5, we find that the cash-out fraction is positively related to $E^{med}_{j,Nov2008}$ after QE1. This relationship is highly statistically significant and is robust to a variety of additional local economic and loan level controls. Importantly, we also document that high equity places extract more equity even conditional on the frequency of refinancing. That is, the patterns in Figure 5 are not driven purely by the differential refinancing propensities shown in Figure 4. To show this, we add monthly refinancing propensities over the same period as an explanatory variable in our regression. We find that the coefficients on both the refinancing propensity and on $E^{med}_{j,Nov2008}$ interacted with the post-QE indicator are positive and strongly significant. Hence, low $E^{med}_{j,Nov2008}$ MSAs both refinanced less and removed less equity, conditional on refinancing. This is intuitive, since these places indeed have less equity to remove when refinancing.

Since prior research has shown tight links between equity removal and spending, these results suggest that locations with different $E^{med}_{j,Nov2008}$ had different spending responses to QE1. We now turn to more direct measures of auto spending from Polk to indeed show this is the case. Panel (a) of Figure 6 shows total monthly auto sales in the top and bottom $E^{med}_{j,Nov2008}$ groups. A few things stand

[19] Total GDP in low equity MSAs fell by $132.8 billion more than total GDP in high equity MSAs between 2008 and 2009. Scaling these differences by the aggregate share of consumption in GDP of 68% delivers .09=8/(.68*132.8).

out. First, the trend in auto sales was very similar in the high and low equity quartiles prior to QE1. In both groups, new auto sales fell rapidly throughout 2008. Second, the trajectory of new auto sales remained parallel through February 2009. This is not surprising since refinancing applications that took place in December 2008 would not result in new mortgage originations until January or February 2009. Third, and most important, after February 2009, auto sales diverge sharply between the low and high equity groups. On average, in the high equity MSAs, sales increased by 31 percent in March-May 2009 relative to November 2008, while they increased only 15 percent in the low equity MSAs. This difference is highly economically and statistically significant, with a p-value <0.01, and the timing lines up perfectly with expected spending responses to QE1. Panel (b) shows that cumulative differences between the groups through 2009 amount to over 200,000 cars, or roughly one month of sales.[20] It is also worth noting that these differential car purchases with housing equity all occur prior to the start of the Cash-for-Clunkers program in July 2009 and so are not driven by this program.[21]

Figure 6: New Auto Sales in Top and Bottom Quartile of MSAs Defined by Median Borrower Equity in November 2008

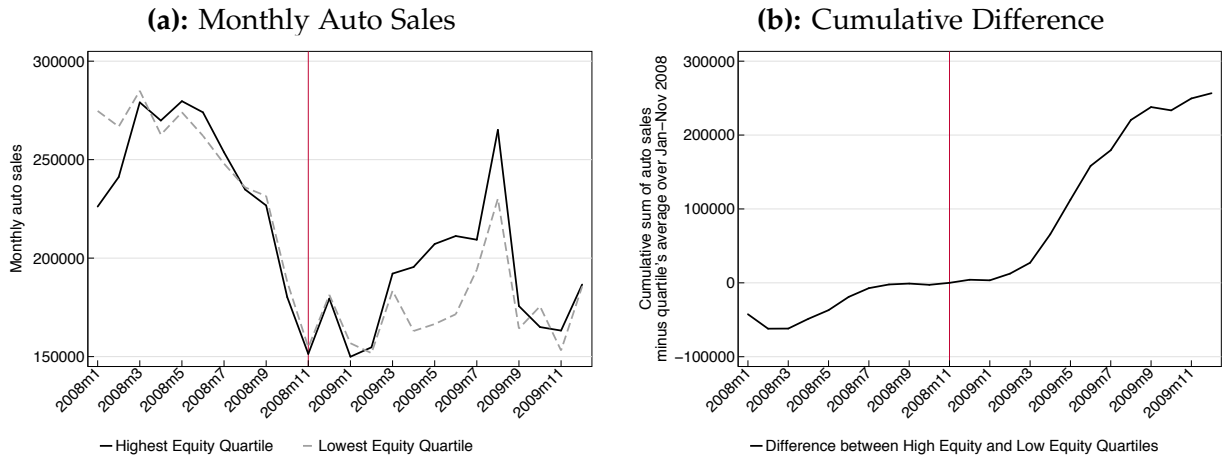


Figure shows the total volume of new car purchases in the top and bottom MSA quartiles by $E^{med}_{j,Nov2008}$. Panel (a) shows total monthly car purchase volumes, summed across MSAs in each quartile. Panel (b) shows the cumulative difference between the two groups, after subtracting from each group's series the group's average car purchases over January to November 2008. Car purchase volumes come from R.L. Polk. MSA quartile groups are the same as in Figure 4. In both panels, the vertical line indicates the month of the QE1 announcement (November 2008).

[20] Appendix Table A-5 shows MSA-level regressions, similar to equation (1), but using the log of monthly new car sales as our dependent variable. β is positive and statistically significant across a variety of specifications with various controls. Furthermore, to directly test the link between refinancing, equity extraction, and auto sales, we add both refinancing and cash-out propensities as independent variables and find they are both correlated with stronger auto sales over this period.

[21] The start of the program can be seen as a large spike in panel (a). Although irrelevant for our regressions with MSA fixed effects, clunker shares are actually mildly negative correlated with median equity which should imply *larger* responses to cash-for-clunkers in low equity regions. (A one standard deviation increase in median equity is associated with a decrease in clunker shares by one-fifth of a standard deviation.) This correlation, together with the fact that the program pulled forward many purchases from the future, may explain why spending differences with equity disappear after July 2009.

4.4.1 Household-level Analysis of Spending after Refinancing

Of course, high and low equity MSAs differ in many ways, so that observed differences in spending following QE1 could be due to factors other than mortgage refinancing and equity withdrawal. Our regressions include both MSA fixed effects and long-run MSA trends to control for many differences across MSAs, so any alternative explanation would need to interact at high frequencies with our QE1 event study. However, there are spending channels that might satisfy this requirement even if they interact in no way with refinancing. For instance, mortgage defaults in low $E^{med}_{j,2008}$ MSAs were substantially higher, which, in turn, decreased credit scores and would lead to increased difficulty in benefiting from lower interest rates on car loans. To provide direct evidence that spending effects indeed arise from refinancing, we move to CRISM data at the household level (rather than the MSA level) and study spending responses of borrowers who refinanced at some point during 2009. The credit record component of the data includes information on each borrower's monthly auto loan balances, and we follow the literature (e.g. Agarwal et al., 2015; Di Maggio, Kermani, and Palmer, 2016) in using large balance increases from one month to the next as a proxy for a new car purchase.[22] We then conduct an "event study" where we look at borrowers' propensity to purchase a car in the months around refinancing. Given our earlier analysis, we are particularly interested in how spending responses differ between cash-out and non-cash-out refinances.[23]

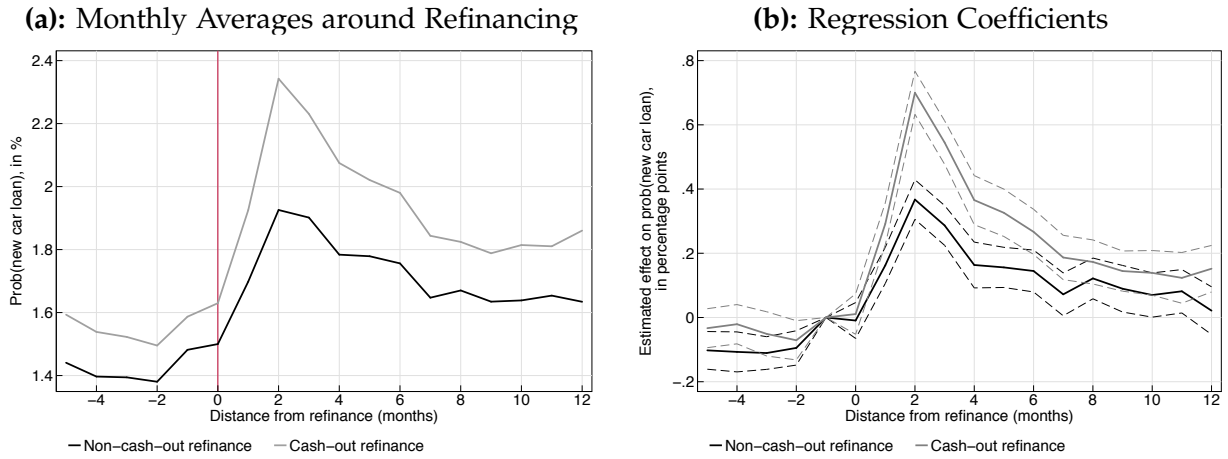
Panel (a) of Figure 7 simply shows the average car purchase propensities for borrowers completing a cash-out or non-cash-out refinance in month 0 over the surrounding months. The figure shows that this propensity spikes for both groups following the refinance, but more so for borrowers who remove equity during the process, supporting the view that this stimulates spending. Panel (b) instead shows the monthly coefficients (and 95 percent confidence intervals) from a regression of the car purchase indicator on calendar-month-by-MSA fixed effects, as well as indicators of the distance in months to the time of the refinance, with month -1 as the omitted month. One can alternatively implement this regression with borrower and month fixed effects, and the results are very similar. The results show a strong increase in car purchase propensities in the months following refinancing, with a peak after two months that is almost twice as large for cash-out as for non-cash-out refinances. For both types of refinancing the effect stays positive for several months after the refinancing, and there does not seem to be a significant pre-trend prior to the refinancing. While this link from refinancing to car spending is not necessarily causal (we do not have exogenous variation in refinancing propensity at the individual level), it nevertheless supports the view that refinancing in the wake of QE1 did stimulate consumer spending, and more so if the borrower also removed equity in the process.

In sum, the results in this section suggest that (1) QE1 increased spending in the aggregate (in part by inducing households to remove equity) and (2) the amount of spending by households differed across regions in a way that is correlated with the equity distribution at the onset of QE1. Collectively, these results show strong evidence for the "refinancing channel" of monetary policy. They also show that the distribution of housing equity is important for determining the aggregate stimulus power of

[22] Following Agarwal et al. (2015) we use a $2,000 increase as an indicator for a new car purchase, but higher cutoffs make little difference to the qualitative results.

[23] We call a refinancing a cash-out if, after subtracting 2 percent from the new loan to cover closing costs, the new mortgage is at least $5,000 above the old mortgage(s). This is close to the binary indicator used by GSEs, and conclusions are not sensitive to altering the cutoff. According to this metric, 44 percent of refinances in our 2009 sample were cash-outs.

Figure 7: Borrower-Level Evidence on Response of Car Purchase Propensity to Refinancing



Panel (a) shows average fractions of refinancing borrowers who obtain a new car loan in each month surrounding a mortgage refinancing (in month 0), distinguishing between cash-out and non-cash-out refinancing. We identify a borrower as obtaining a new car loan in month t if their car loan balance increases by \$2,000 or more between month $t-1$ and month t. Panel (b) shows coefficients on month indicators (relative to the refinancing month, with -1 as the omitted month) from a regression of the new car loan indicator on calendar-month-by-MSA fixed effects as well as months-from-refinancing indicators interacted with whether the refinancing involves cash-out. Dashed lines are 95 percent confidence intervals based on robust standard errors clustered at the MSA level. Both panels are based on a 50 percent sample of borrowers identified as refinancing in 2009 in the CRISM sample.

this channel. Finally, because monetary policy stimulated spending most in the locations where equity was high and unemployment was low, the refinancing channel of monetary policy amplified regional inequality during the Great Recession.

5 The Time-Varying Refinancing Channel in the US

In this section, we show that there is substantial variation across time in the distribution of house price changes and its correlation with changes in unemployment. We then show that these changes lead to different relationships between equity, unemployment and refinancing in other recessions. Hence, when combined with the evidence from QE1, these results suggest that the regional distribution of housing equity produces time-varying effects of monetary policy through the refinancing channel.

Since CRISM data begins in 2005, we cannot construct the distribution of housing equity necessary to simply replicate our previous analysis in earlier recessions. However, as noted above, there is a strong correlation between house price changes and housing equity during this recession. Motivated by this relationship, we now show that the regional distribution of house price changes varies substantially across time. Panel (a) of Figure 8 shows the mean of the annual real house price growth distribution across states calculated using CoreLogic data for each year from 1976-2013.[24]

While the large house price declines in the Great Recession stand out, there is substantial variation

[24]State-level house prices are deflated using the CPI-U and are population weighted. We concentrate on the state-level rather than MSA-level distribution of house price growth, since MSA-level labor market data is only available starting in 1990, but patterns at the MSA level are similar for the mean and standard deviation.

Figure 8: Changes in Distribution of House Price Changes Across Time

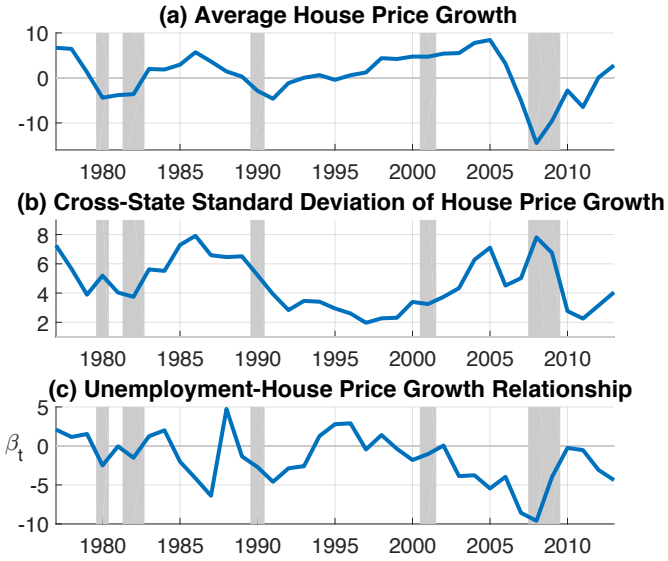


Figure shows time-series for the mean and standard deviation of state-level house price growth as well as the cross-state relationship between house price growth and unemployment rate changes. Calculations are population weighted by state.

across time. During the 2001 recession, real house prices were growing rather than falling, and house price declines during the previous three recessions were mild. Next, we calculate the standard deviation of house price growth across states in each year. Panel (b) plots the evolution of this cross-state standard deviation across time. Clearly, spatial variation in house price growth was unusually high during the late 80s and during the Great Recession. During the 2001 recession, spatial variation was unusually low, and in the earlier 80s recessions, there were intermediate levels of dispersion.

Finally, panel (c) explores relationships between state-level unemployment and house price growth, a moment that will be particularly important for determining monetary policy implications for inequality. In particular, we run a regression of one-year house price growth on state-level one-year changes in unemployment (in percentage points) interacted with annual dummies for each year from 1976-2013 and including year and MSA fixed effects so that the results are not driven by slow moving state-level trends or aggregate unemployment changes or house price growth:

$$\Delta \log HP_{i,t} = \alpha + \beta_t \Delta UR_{i,t} + \gamma_t + \zeta_i + \epsilon_{i,t} \tag{2}$$

Since identification is purely cross-sectional, β_t measures the comovement between house prices and unemployment in the cross-section in each year. Panel (c) shows that in the Great Recession, unemployment and house price growth were unusually negatively correlated: locations with large increases in unemployment exhibited large house price declines. Again, these patterns vary across time. In the 2001 recession and in the 80s recessions, there was essentially no relationship between unemployment and house price growth, while the 1990 recession exhibited intermediate behavior.

In the following section we build a theoretical model consistent with these time-series patterns and show that they imply striking variation across time in the consequences of monetary policy. Supporting these conclusions, Figure 9 illustrates that the pass-through of monetary policy through

the mortgage market was much stronger in the 2001 recession than in the 2008 recession, and that regional effects were very different as well. Panel (a) shows the times-series of monthly refinancing propensities in HMDA for both the lowest and highest unemployment quartile MSAs during the 2008 recession.[25] MSAs are sorted into quartiles based on the total increase in unemployment between November 2007 and October 2009; the top (bottom) quartile experienced unemployment increases of 6.3 percent or more (3.8 percent or less) over this period.[26] Given the high correlation between unemployment changes and house price changes, the unemployment results in panel (a) are very similar to the pattern shown for equity quartiles in Figure 4.

Figure 9: Refinancing Propensities for High and Low Unemployment Change Quartiles, 2008 Recession and 2001 Recession

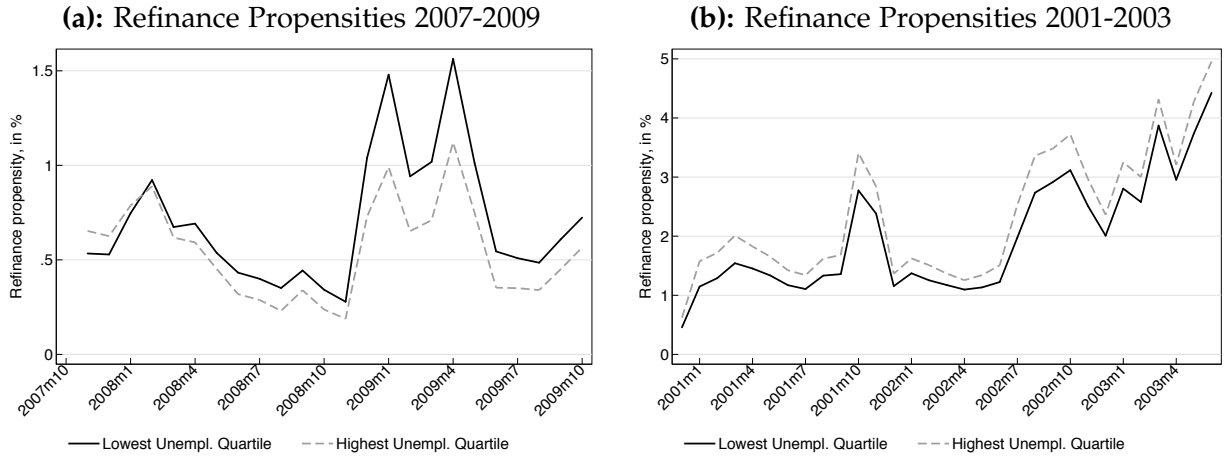


Figure shows monthly refinancing propensities for the top and bottom quartiles of MSAs by unemployment changes. Specifically, for each MSA we compute the change in unemployment between November 2007 and October 2009 (panel a) and December 2000 to June 2003 (panel b). Each MSA is placed into quartiles based on this measure. The quartiles are defined such that there is equal population in each quartile. Refinancing propensities are computed using HMDA data on origination volumes (by month of the loan application) divided by the total number of mortgages in MSA quartiles from the American Community Survey in 2008 (panel a) or 2000 (panel b).

Panel (b) shows that time-series patterns for monthly refinancing by unemployment quartiles are quite different during the 2001 recession. The sorting is again done by total increases in unemployment over this period; for the top (bottom) quartile the unemployment rate increased by 2.5 percent or more (1.6 percent or less). The results are very different from those shown in panel (a): during the entire period studied, refinancing volumes were higher in the MSAs more strongly affected by a surge in unemployment.[27] Even more important, overall refinancing propensities were dramatically higher in this earlier episode when house prices were growing rather than shrinking (especially between mid-2002 and mid-2003). While there was no single monetary policy "event" during this period, as

[25] Since CRISM data starts in 2005 we must use HMDA data for this analysis.

[26] The periods chosen for calculating unemployment differences are determined by turning points in the overall seasonally adjusted US civilian unemployment rate. In the earlier episode, national unemployment started increasing from a (local) low of 3.9 percent in December 2000 up to a high of 6.3 percent in June 2003. In the latter episode, unemployment increased from 4.7 percent in November 2007 to 10.0 percent in October 2009.

[27] Chen, Michaux, and Roussanov (2013) similarly find that over 1993-2009, refinancing activity was relatively higher in states with worse economic conditions (after controlling for aggregate conditions or including quarter fixed effects).

the federal funds rate instead declined gradually, it is evident from Figure 1 that overall mortgage rate declines relative to recent levels were similar in 2001-2003 and in 2007-2009.[28] The fact that the incentive to refinance in terms of lower mortgage rates was similar during both recessions but refinancing activity was much higher in 2001-2003 is consistent with the fact that fewer households were underwater during this earlier period.

6 A Model of Regional Heterogeneity and Monetary Policy

In this section, we move to a theoretical analysis of the interaction between monetary policy and regional heterogeneity using a quantitative equilibrium model of household refinancing. Our goal is both to clarify how the distribution of housing equity affects the transmission from interest rate cuts to aggregate spending through the refinancing channel and to highlight that because this distribution is time-varying, monetary policy makers may at times face a trade-off between stimulating aggregate spending and increasing regional consumption inequality. Our model includes many important and realistic features of mortgage markets, including equilibrium effects between borrowers and lenders, but it is intentionally stylized in many other dimensions, as the goal of the analysis is to carefully illustrate and clarify the ways in which the refinancing channel of monetary policy interacts with the regional distribution of equity. Embedding our framework into a richer DSGE structure would potentially produce more realistic numerical results but would complicate the analysis substantially in a way that obscures the intuition for the interaction between monetary policy and the regional distribution of housing equity. More importantly, as further discussed in Section 7.4, we argue that complicating the analysis in this direction would only amplify our conclusions.[29]

The starting point is a standard consumption-savings model with income shocks and borrowing constraints as in Huggett (1993). To this standard framework, we add housing and mortgages. Houses are subject to stochastic regional house price shocks, and individuals can borrow against the value of their homes using a fixed-rate mortgage that can be refinanced at any time by paying a fixed cost. We account for equilibrium interactions between borrowers and lenders by assuming that mortgage payments are received by lenders in the economy who will potentially adjust consumption when borrowers refinance. Throughout the analysis, we focus on an intentionally stylized environment in which we take the regional distribution of house prices and income as exogenous and explore the behavior of consumption, savings, and refinancing in response to exogenous interest rate changes. We look at these responses both in the aggregate as well as across regions and ask how ex-ante features of regional heterogeneity, such as the regional distribution of housing equity, affect both aggregate consumption stimulus and ex-post regional consumption inequality in response to monetary policy.

We focus on inequality across regions, as opposed to across households within regions, because the refinancing channel of monetary policy is less relevant for within-region household inequality: other

[28] Hurst et al. (2016) also show there is little spatial variation in rates in either period.

[29] For example, treating the regional income distribution as exogenous but consumption as endogenous amounts to assuming that consumption is fully tradable. This dramatically simplifies the model solution, and accounting for the feedback from local consumption to local income would only amplify our conclusions: if monetary policy increases consumption by more in regions with high initial income, inequality will be further amplified, as income then rises by more owing to the greater consumption response. Endogenizing house prices would also amplify our conclusions, as rate declines increase regional house price inequality. Similar arguments apply for aggregate time-series variation.

factors such as income and wealth are much more important determinants of household inequality within regions. However, house price shocks have a large regional component, which means that the refinancing channel of monetary policy interacts importantly with cross-region inequality even if it matters little for understanding within-region inequality. Furthermore, the regional dimension is relevant because labor markets and housing markets are geographically segmented, which induces a joint distribution of equity and income that is very different than at the individual level. For example, within a region, richer households could be buying more expensive houses. But for understanding the time-varying consequences of monetary policy for consumption inequality through the refinancing channel, it matters not that average income and house prices at the individual level are highly positively correlated. Instead, it matters that this correlation changes over time, and this occurs mainly at the regional level because of geographic segmentation. Finally, our previous empirical evidence on the heterogeneous response of refinancing and consumer spending to QE1 provides important discipline for our quantitative model. This evidence is inherently regional.

6.1 Model Description

Environment. The economy is populated by a continuum of infinitely-lived households, indexed by i and located in region $j = 1, 2, ..., J$, and a representative lender.

Idiosyncratic Earnings. In each period t, household's earnings are given by y_t^{ij}, which follows a random walk with drift,

$$log(y_t^{ij}) = \mu_y^j + log(y_{t-1}^{ij}) + \varepsilon_t^{ij}$$

with region-specific income drift μ_y^j and mean zero income shock ε_t^{ij}, which is i.i.d. over time but possibly correlated both across and within regions. In order for the process to have a stationary distribution, we exogenously bound it between $[\underline{y}, \bar{y}]$.

Assets and Liabilities. Households have access to a risk-free asset a_t^{ij} paying interest rate r_t, with a no-borrowing constraint $a_t^{ij} \geq 0$. They are endowed with one unit of housing with price q_t^j, which can be used as collateral for mortgage borrowing. House prices follow a random walk with drift,

$$log(q_t^j) = \mu_q^j + log(q_{t-1}^j) + \nu_t^j$$

where μ_q^j is region-specific trend house price growth and ν_t^j is a mean zero individual house price shock that is i.i.d. over time but is perfectly correlated within a region.

We assume that both earnings and house prices are random walks for two reasons. The first is computational: it allows us to reduce the state space, as we show in Claim 1. The second is because it simplifies aggregation: households can be collected into regions and aggregated in a straightforward manner without adding additional state variables because aggregate, regional, and idiosyncratic household shocks enter symmetrically in the problem.[30]

Since our empirical evidence focuses on refinancing for non-moving households, we assume for

[30]For example, whether current household earnings are high because regional earnings are high or because the household was hit by a positive idiosyncratic earning shock is equivalent from the household's point of view.

simplicity that agents cannot buy or sell houses and that mortgage debt is of infinite maturity.[31] Mortgage debt requires a constant mortgage payment $r^m_{\tau_0} m^{ij}_{\tau_0}$ every period, which is determined at the moment of debt issuance τ_0. However, households can refinance their mortgage at any time $\tau > \tau_0$ by paying a fixed monetary cost $F^{ij}_\tau q^j_\tau$.[32] We assume that F^{ij}_t is an i.i.d. stochastic process that is uncorrelated with house prices or income in order to introduce additional heterogeneity in refinancing decisions across households and so better fit the data.

When refinancing, households lock in current interest rate r^m_τ and future mortgage payments $r^m_\tau m^{ij}_\tau$. Furthermore, we assume that when refinancing, households always borrow up to the maximum amount allowed by the loan-to-value requirement γ. This simplifies the household decision problem and computations in our benchmark model, but we show in Section 7.2 that relaxing this assumption has little effect on our conclusions. This implies that the new mortgage balance is $m^{ij}_\tau = \gamma q^j_\tau$ and that cashed-out equity is $m^{ij}_\tau - m^{ij}_{\tau_0}$. These borrowing constraints capture the primary features of fixed-rate mortgages with options to refinance common in the US. Finally, we assume interest rates $\{r_t, r^m_t\}$ follow an exogenous Markov process, which we discuss in the calibration.

Household Problem. For notational clarity, we drop the agent and region indices ij when describing the individual household problem. We assume that households derive period-utility from consumption, $u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}$ with $\sigma \geq 1$. The value function of a household with assets a_t, earnings y_t, mortgage rate and balance $\{r^m_0, \gamma q_0\}$, facing current house price q_t, and current mortgage rate r^m_t and interest rate r_t can be written recursively as,

$$
\begin{aligned}
V(a, y, q_0, r^m_0, q, r^m, r, F) &= \max\{V^{norefi}(a, y, q_0, r^m_0, q, r^m, r, F), V^{refi}(a, y, q_0, r^m_0, q, r^m, r, F)\} \\
V^{norefi}(a, y, q_0, r^m_0, q, r^m, r, F) &= \max_{\{a', c\}} u(c) + \beta \mathbb{E}[V(a', y', q_0, r^m_0, q', r^{m'}, r', F')] \\
s.t. \quad c + a' &\leq a(1+r) + y - \gamma r^m_0 q_0 \\
a' &\geq 0, c \geq 0 \\
log(q') &= \mu_q + log(q) + \nu \\
log(y') &= \mu_y + log(y) + \varepsilon \\
V^{refi}(a, y, q_0, r^m_0, q, r^m, r, F) &= \max_{\{a', c\}} u(c) + \beta \mathbb{E}[V(a', y', q, r^m, q', r^{m'}, r', F')] \\
s.t. \quad c + a' &\leq a(1+r) + y - \gamma r^m_0 q_0 + \gamma(q - q_0) - Fq \\
a' &\geq 0, c \geq 0 \\
log(q') &= \mu_q + log(q) + \nu \\
log(y') &= \mu_y + log(y) + \varepsilon
\end{aligned}
$$

By inspecting this problem, we can see the household's incentives to refinance. When the current interest rate r^m is below the rate when the household last refinanced r^m_0, the household can secure a lower mortgage payment forever, even if house prices are unchanged. When the current house price

[31] Allowing households to move to extract equity would complicate the setup but would produce similar non-linear interactions between equity and consumption.

[32] We assume the fixed cost is proportional to current house prices so that these costs remain relevant in the presence of house price growth and because this is necessary to make the household value function homogeneous in house prices.

q is above the price when the household last refinanced q_0, the household can refinance to cash out accumulated equity even if rates have not changed, but will then face increased mortgage payments in the future. In order to characterize policy functions, it is useful to eliminate a state variable. We do so by showing that the value function is homogeneous in house prices:

Claim 1 *For fixed* $\{r, r^m, r^0\}$*, the value function is homogeneous of degree* $1-\sigma$*.*
This implies that $V(a, y, q_0, r_0^m, q, r^m, r, F) = J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F)q^{1-\sigma}$ *where* $\tilde{a} \equiv \frac{a}{q}, \tilde{y} \equiv \frac{y}{q}, \tilde{x} \equiv \frac{q_0}{q}$ *and* $J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F) \equiv V(\frac{a}{q}, \frac{y}{q}, \frac{q_0}{q}, r_0^m, 1, r^m, r, F)$.

Proof. See Appendix A.3. ■

The transformed problem in Claim 1, with the inverse of equity $\tilde{x}_t$ as a state variable, provides further intuition on the determinants of household refinancing. When $\mu_q > 0$, the household accumulates equity, on average, in the part of the state-space where refinancing is not optimal. Once sufficient time has passed since the last refinancing to acquire substantial equity, it is optimal to refinance and cash out this equity. This logic is typical of fixed adjustment cost models, where inaction is optimal until the state changes enough to justify paying the fixed cost. Often, these models are stylized enough that the state-space is one-dimensional and optimal policies are characterized by a single adjustment "threshold." This is not true in our setup since we have a richer state-space, with a state variable (assets a_t) evolving endogenously even when not refinancing. However, when solving the model numerically, we find that the refinancing decision is indeed a threshold level of equity, which depends on the level of assets, earnings, fixed cost, and interest rates.

Mortgage Lenders. We assume that mortgage debt is paid into a mutual fund. Share θ of this mutual fund is held by a representative US lender while share $1-\theta$ is held by foreign lenders (e.g., Chinese investors). This means that the representative lender ultimately holds a fixed fraction θ of the outstanding mortgage debt. Allowing for foreign investors captures the empirically relevant fact that much mortgage debt is ultimately held abroad.

We introduce mortgage lenders because we are interested in computing aggregate consumption responses to changes in interest rates. If we ignored lenders when computing aggregate consumption, we would potentially miss important offsetting general equilibrium effects. For example, when borrowers refinance after a decline in interest rates, the dividends accruing to lenders decrease, which should reduce their consumption (see Greenwald, 2016).

For simplicity, we assume the representative lender is a permanent income consumer and receives certain dividend payments d from the mutual fund (a consequence of a law of large numbers for households). Given a law of motion for dividends $d' = \tilde{G}(.)$, their value function is[33]

$$\begin{aligned} V_R(a_R, d, r) &= \max_{c_R, a_R'} u(c_R) + \beta_R V(a_R', d', r') \\ s.t. \quad c_R + a_R' &\leq a_R(1+r) + \theta d \\ c_R &\geq 0 \end{aligned}$$

[33]The law of motion is determined in equilibrium as a function of the refinancing decisions of all households. Since we do not endogenize interest rates, this is straightforward since we can calculate the present value response of dividends to a decline in interest rates entirely from the borrower side and then separately calculate lender consumption responses.

Recursive Equilibrium Definition. A Recursive Equilibrium is an initial distribution S for $\{\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, F, q\}$ across households i in regions j; initial lender assets a_R^0; a law of motion for d', r^m and r; value functions $J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F), V_R(a_R, d, r)$; and policy functions $\left[\tilde{a}', \tilde{c}, \tilde{x}, \mathbb{I}^{refi}\right](\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F)$ and $\left[a_R', c_R\right](a_R, d)$ such that

1. The policy functions solve households' and lender's problems.

2. For all realizations of $\{y_t(i,j), q_t(i,j), F_t(i,j), r_t^m, r_t\}_{t=0}^{\infty}$ across households and regions (i,j), the law of motion for dividends $G(S, r^m, r)$ implied by the policy functions is equal to the perceived law of motion by lenders $\tilde{G}(.)$ and satisfies

$$d_t = \int \left[\gamma r_{0,t}^m(i,j)\tilde{x}_t(i,j) - (\gamma(1-\tilde{x}_t(i,j)) - F_t(i,j))\,\mathbb{I}_t^{refi}(i,j)\right] q_t(i,j)didj.$$

This equilibrium definition does not impose market clearing on the asset and consumption market because we take interest rates, goods and house prices, and incomes as exogenously determined. However, we assume the mortgage market clears by introducing the representative lender.[34] Finally, a region in this economy is defined as a collection of households that experience common shocks to house prices and income. Aggregates are simply obtained by adding up all regions.

6.2 Calibration Strategy

The model is calibrated annually and most parameters of the model are calibrated at standard values. As is standard in the risk-sharing literature, we set $\sigma = 2$ to generate an intertemporal elasticity of substitution of $1/2$. We assume the risk-free rate is constant at $r = 0.03$ and set $\beta = 0.93$. For simplicity we set $\beta_R = \frac{1}{1+r}$ so that lenders perfectly smooth consumption. For our baseline results, we assume the mortgage rate is constant at $r^m = 0.06$ and then consider the response to a one-time, unanticipated permanent decline to $r^m = 0.05$. We begin with this specification as it illustrates effects transparently in the simplest possible environment and also facilitates numerical calculations, since it eliminates the calculation of r^m expectations from the value function. The initial value of 0.06 and decline to 0.05 roughly matches 30-year fixed rates before and after QE1. In robustness results we instead assume that r^m follows an AR process with annual persistence of 0.88 and standard deviation of 0.0055 to match the behavior of 30-year rates from 1990-2015 and show that this specification delivers extremely similar quantitative conclusions. We set $\gamma = 0.8$, assuming that lenders require a minimum of 20 percent equity. We set $\theta = 0.5$ to account for the fact that a substantial fraction of US mortgage debt is held abroad, either directly (about 20 percent in 2008, according to the Flow of Funds) or indirectly through ownership of other institutions (such as banks) that hold mortgage debt. We show robustness to this assumption in Section 7.

Following Kaplan and Violante (2010), permanent income shocks are drawn from a normal distribution with standard deviation of 0.1 to match earnings changes in PSID data. While this is the total standard deviation of income at the household level, we assume that some fraction of these income

[34] One can interpret this as a partial equilibrium or as a small open economy with a fixed exchange rate, where local households have endowments of perfectly tradable goods and interest rates and prices are determined internationally.

shocks is common to all households in a region. In our baseline calibration, we assume that the standard deviation of this common region component is 0.025, to match the differences across regions in our empirical work.[35] The drift in income and house prices is assumed to be identical to generate constant real house prices on average.[36]

House price shocks are calibrated to match the annual growth rate and standard deviation of MSA house price changes in CoreLogic data of 0.025 and 0.05. We assume that in the stochastic steady-state, shocks to house prices and shocks to income are independent. While these shocks are independent *on average*, and this independence of shocks is assumed when households make their policy decisions, one of the primary questions we explore in our model is how particular sequences of shock realizations affect the consequences of monetary policy. During the Great Recession, house prices fell substantially on average and there were substantial differences across regions. Furthermore, these region-specific changes were highly correlated with changes in income: the locations with the largest declines in house prices saw the largest declines in income. Since our empirical evidence is drawn from this period, we calibrate the remaining parameters of the model to match the distribution of house prices and income in 2008 and the refinancing responses across regions following QE1. We call this calibration our *baseline* economy.

More specifically, to construct our baseline economy, we initialize the model from the stochastic steady-state, but in period t, we hit the economy with the aggregate house price decline of 12.5% observed in 2008. Households are also hit with an additional regional house price shock which can take the value -12.5%, 0%, +12.5% so that 1/3 of regions experience a total house price decline of 25%, 1/3 experience a decline of 12.5% and 1/3 experience no decline. The 25% decline is picked to match the house price decline for the lowest housing equity quartile in our empirical analysis, while the 0% change matches that in the highest housing equity quartile. That is, we define regional house price shocks so that regions in our model can be mapped directly to regional differences in our empirical analysis. Similarly, we calibrate regions so that they differ by -2SD, 0 and +2SD of the regional income shock. If these income shocks were uncorrelated with house price shocks, then our simulated economy would be composed of 9 possible regions representing the 3x3 combinations of house price and income shocks. However, since income and house prices were highly correlated in the Great Recession, we instead assume that the regional shocks to house prices are perfectly correlated with those to income. This means that in our baseline economy, there are effectively three regions: relatively high house price and income, middle house price and income and low house price and income.

Finally, the fixed cost process is calibrated to match the level of refinancing just prior to QE1 as well as the regional responses to mortgage rate reductions under QE1. In particular, we assume that each household draws an i.i.d. fixed cost each period that can take on either a high or a low value and pick the levels of the high and low fixed costs and their relative probabilities to target a monthly refinancing rate just prior to QE1 of 0.0025, an increase in the refinancing rate of 0.0025 in the lowest

[35] Note that under the random walk assumption, shocks to idiosyncratic and regional income enter identically in the value function and so only total income needs to be tracked. The regional share of income shocks only matters for calculating measures of regional inequality and has no effect on any aggregate impulse responses.

[36] This implies that in the original problem, both house prices and income remain relevant in the long-run. If $\mu_y > \mu_q$, then in the long-run, housing becomes a vanishing share of household budgets so the refinancing decision is irrelevant for decisions, and if $\mu_y < \mu_q$ then eventually housing is the only thing that matters in the budget.

house price region, an increase of 0.0075 in the middle house price region, and an increase of 0.011 in the highest house price region.

Our primary counterfactual analysis explores whether the same change in interest rates would have had different consequences if it had occurred during a recession with a different distribution of house prices and income. For example, we will frequently compare impulse responses in our 2008 calibration to those in the stochastic steady-state. In the stochastic steady-state, aggregate house price growth is 2.5%, the standard deviation of house price growth across regions is 5% and there is no correlation between house price growth and income. Comparing those assumptions to the time-series patterns in Figure 8, it is noteworthy that they correspond extremely closely to economic conditions in the 2001 recession. For this reason, we will often describe the comparison between 2008 and the stochastic steady-state as a comparison between the 2008 and 2001 recessions.

We calibrate to the 2008 period rather than the stochastic steady-state for several reasons. First, this is the period for which we have strong empirical evidence on the distribution of house price changes, income and refinancing responses from our event study around QE1. Second, with a long empirical time-series, one could calibrate to match the average distribution of income and house prices in the economy and responses to interest rate shocks in an average year, but our refinancing micro-data are available only for the most recent boom-bust period. It is unclear that any year over this period well represents a "normal" steady-state period. Furthermore, by design, monetary stimulus is correlated with recessions and so any empirical evidence on the effects of interest rate reductions is going to come from periods with negative aggregate conditions. That is, any empirical measure of refinancing elasticities to interest rate reductions will always be primarily identified from recession periods; therefore, we target this elasticity during such a period explicitly in the model. Finally, we focus on the elasticity of refinancing to interest rate reductions across regions rather than the aggregate change in refinancing since aggregate relationships may be contaminated by other confounding unmodeled aggregate shocks. By focusing on matching relative changes across regions, we eliminate any such confounding aggregate shocks and isolate just the effects of local house prices and income on the refinancing response to interest rates that is the focus of our analysis.

6.3 Theoretical Results in the 2008 Baseline Economy

Figure 10 shows the threshold-like property of refinancing policies in a stationary environment with permanently high or low mortgage rates for a household with median income and earnings, as well as in a non-stationary environment right after a permanent mortgage rate decline.[37] The refinancing equity threshold is lower when mortgage rates are permanently low than when they are permanently high, making households refinance more frequently and extract less equity on average. Moreover, right after a mortgage rate decline, households refinance at even lower levels of equity. Intuitively, refinancing is more frequent in an economy where mortgage rates are lower because the cost of borrowing using one's house as collateral—future mortgage payments—is lower, whereas the benefit—equity cash-out net of fixed costs—is independent of mortgage rates.[38]

[37] Since households cannot sell housing, extractable equity is all equity above the fraction $1-\gamma$ required when refinancing.

[38] Appendix Figure A-6 shows how the refinancing threshold changes with assets and earnings. The threshold is generally increasing in each. Intuitively, the marginal utility of an extra unit of resources coming from equity extraction is higher

Figure 10: Refinance Decision Follows a Threshold Policy

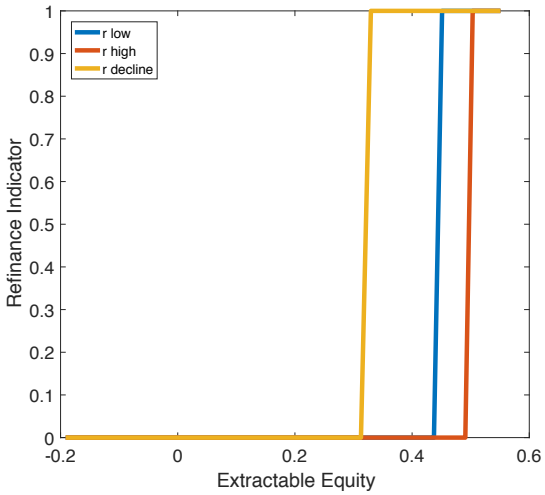


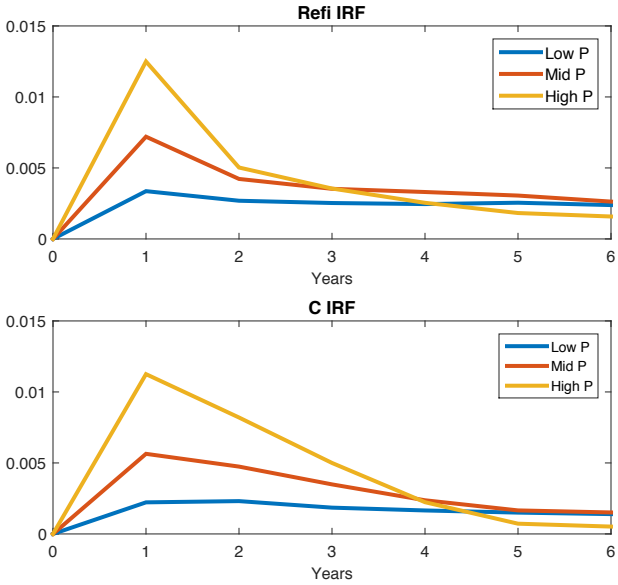
Figure 11 shows the impulse response of each region to a decline in mortgage rates from 0.06 to 0.05 in the baseline economy calibrated to match economic conditions in 2008. The top panel shows the change in the fraction of loans refinancing at a monthly rate. By construction, this closely matches the changes in Figure 4, since the model is calibrated to hit these numbers. Just as in the data, regions with high equity are much more likely to refinance in response to the decline in rates. The bottom panel shows that consumption responses are also larger in high equity regions. While we do not have empirical measures of broad consumption at the regional level, this is consistent with empirical patterns for auto spending in Figure 6.

We next compute aggregate consumption in the economy as well as the variance of log consumption across regions. Figure 12 shows their responses to the decline in mortgage rates and delivers our first important theoretical result: in our baseline economy, a reduction in interest rates increases aggregate consumption but also increases the variance of consumption across regions. This increase in regional variance occurs because consumption increases most in regions with high house prices. Since house prices and income are correlated in this economy, these regions already have the highest income and consumption before the interest rate declines. Thus, while monetary policy increases overall consumption, it does so mainly by stimulating consumption in locations that are already doing relatively well, so that there is an important trade-off between stimulus and inequality.

This trade-off is strongly suggested by our empirical patterns that show a strong correlation between house prices, income, refinancing and auto spending around QE, but the model addresses two potential concerns with that interpretation of the empirical patterns: 1) In the model, we can precisely measure the change in the variance of consumption across regions. Our empirical analysis shows that consumption increases more in high equity regions, but the difference-in-differences nature of the exercise together with a normalization for overall scale makes it ill-suited for measuring how the actual level of inequality changed. More specifically, while consumption increased more in high equity regions in response to QE1, it is possible that high equity regions initially had consumption that was temporarily low, so that QE1 would actually decrease inequality. In our quantitative model, this is

when current assets or earnings are lower, making households want to refinance more often at lower equity levels.

Figure 11: Response to Interest Decline by Region



Note: For the baseline 2008 calibration, the refinancing impulse response function (IRF) shows the change in the monthly fraction of households refinancing in response to a one percentage point reduction in mortgage rates. The consumption (C) IRF shows the change in log consumption in response to the same one percentage point reduction in mortgage rates.

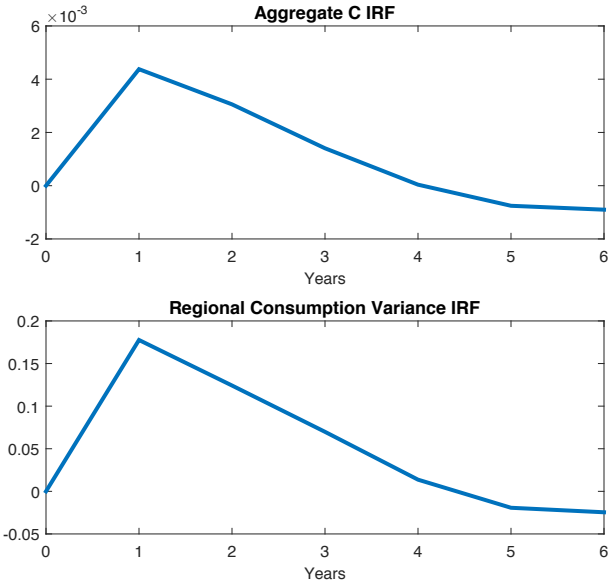
not the case: high equity regions have substantially higher consumption so that a decline in mortgage rates dramatically amplifies inequality. 2) The empirical analysis measures only differential effects of the policy across regions and cannot measure any offsetting aggregate effects from lenders. Our quantitative model explicitly accounts for these effects and shows they do not cancel out. Aggregate consumption rises at least modestly when mortgage rates fall in the baseline economy. There are two reasons for this. First, some of the mortgage debt is owned by foreign lenders. Second, and more important, cash-out refinancing leads to new spending, on net.

The model thus confirms the effects of quantitative easing suggested by our empirical analysis: under the economic conditions present in 2008, monetary policy resulted in a modest increase in aggregate consumption but also an increase in inequality across regions.

6.4 Counterfactual Analysis

That our model reproduces the behavior of the economy just before and after QE1 gives us some confidence in using it for more ambitious counterfactual analysis. We show that the theoretical effects of monetary policy are highly non-linear with respect to the regional distribution of housing equity in the economy. This means that accounting for time variation in this distribution is crucial for correctly predicting the consequences of monetary policy because the effects of interest rate changes on both aggregate economic activity and regional inequality change as the level of housing equity, its variance and its correlation with income move across time. In Section 5, we showed that such time variation is an important feature of the data.

Figure 12: Aggregate Stimulus vs. Regional Inequality



Note: For the baseline 2008 calibration, impulse response functions show the change in log aggregate consumption and in log consumption variance across regions in response to a one percentage point reduction in mortgage rates.

6.4.1 2008 Economic Conditions vs. Stochastic Steady-State

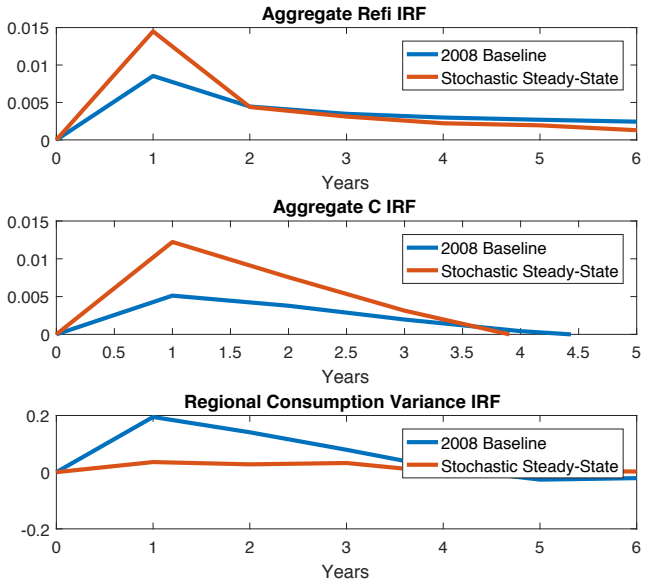
Figure 13 compares the impulse response function of aggregate refinancing activity, aggregate consumption and regional consumption inequality to an interest rate shock in the 2008 baseline to that which would have instead occurred in the stochastic steady-state. In the stochastic steady-state aggregate house prices do not decline, there is no increase in the variance of house prices, and there is no correlation between house prices and income.[39] As described above, this makes conditions in the stochastic steady-state similar to those in the 2001 recession. These counterfactuals, therefore, shed light on whether the effects of QE1 through the refinancing channel would have been different if regional house price patterns in 2008 had instead been similar to those in 2001.

As seen in Figure 13, changing the equity distribution dramatically changes the consequences of monetary policy: in the stochastic steady-state, the same decline in interest rates raises aggregate consumption by almost three times as much as in 2008 (middle panel). This means QE1 would have had much larger stimulative effects if it had been enacted under the equity distribution present in 2001 as opposed to 2008. There are two reasons for the larger consumption response: First, as shown in the top panel, the aggregate refinancing response to the rate decline is nearly twice as large in the stochastic steady-state because fewer borrowers are underwater. Second, households that refinance have more equity to extract in the stochastic steady-state, which contributes to an additional consumption boost. Interestingly, as seen in the bottom panel of the figure, there is a negligible effect on inequality across regions when the policy is implemented at the stochastic steady-state. This implies that there is not always a trade-off between stimulus and regional inequality.

To understand why the 2008 distribution of equity makes consumption respond less to monetary

[39] In the baseline we calibrate the house price shocks across regions to match that of the regional housing equity distribution in the data. In the stochastic steady-state there is not such a natural benchmark for defining this difference; therefore, we define regions as those receiving a +1, 0 and -1 standard deviation house price shock.

Figure 13: 2008 Economic Conditions vs. Stochastic Steady-State



Impulse response functions show the change in log aggregate consumption and log consumption variance across regions in response to a one percentage point reduction in mortgage rates. See text for the description of the baseline economy calibrated to match November 2008 as well as the stochastic steady-state, which is more similar to the 2001 recession.

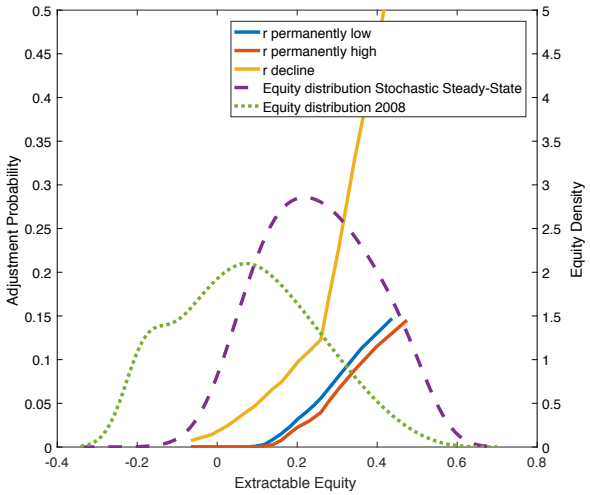
policy than in the stochastic steady-state, it is useful to first explain why interest rate reductions lead to increased consumption. Figure 14 shows the distribution of equity in our baseline 2008 calibration (in dotted green) and in the stochastic steady-state (in dashed purple) as well as the fraction of households adjusting for a given level of equity when the interest rate is permanently 0.06, permanently 0.05, and in the period when it declines from 0.06 to 0.05. This is the analogue to Figure 10 but averaging over the endogenous joint distribution of assets and income for each value of extractable equity. Consistent with the threshold policy shown above, households are more likely to refinance when equity is large. Since $r^m > r$, households that refinance always consume some fraction of their equity when refinancing so that the marginal propensity to consume (MPC) out of equity is always greater than zero and typically large.[40] This means that consumption increases if more households refinance.

When interest rates decline, households lower their refinancing threshold substantially (i.e., the red line shifts to the left, beyond the blue line in Figure 14). This is because some households that would not choose to extract equity today if their interest rate was already low choose to extract equity sooner so that they can also reset their payment to the new lower rate. Thus, when interest rates decline, there is a burst of (cash-out) refinancing and an increase in consumption. This extra refinancing arises from the mass of households between the red line and the yellow line, since households to the right of the red line will refinance even if interest rates do not decline and households to the left of the yellow line will not refinance even if interest rates do decline.[41] It is not the fraction of households that want to refinance and extract equity that matters for the strength of the interest rate impulse response,

[40] On average this MPC is 0.5 so that half of extracted equity is consumed on impact. Unsurprisingly, this number is higher for households with low liquid assets and smaller for those with high liquid assets.

[41] The change in the fraction of households refinancing with equity x is given by the vertical difference between the red and yellow lines, so the total change in refinancing is given by this difference integrated over the distribution of equity.

Figure 14: Distribution of Equity and Refinancing Probability: 2008 Calibration vs Stochastic Steady-State



Note: This figure shows the simulated distribution of equity and the fraction of households refinancing under two different equity distributions.

it is the change in this fraction that is the relevant object for assessing the consequences of interest rate declines.[42] In 2008, the distribution of equity shifts left and fans out relative to the stochastic steady-state. On net, this leaves many fewer households in the region where refinancing decisions are triggered by the rate decline and so reduces refinancing impulse responses. This then leads to smaller aggregate consumption responses.

The intuition for the differential response of regional consumption inequality in 2008 vs. the stochastic steady-state is more subtle because the respective distributions of equity differ in several dimensions that result in offsetting effects of rate declines. The next section discusses this in detail.

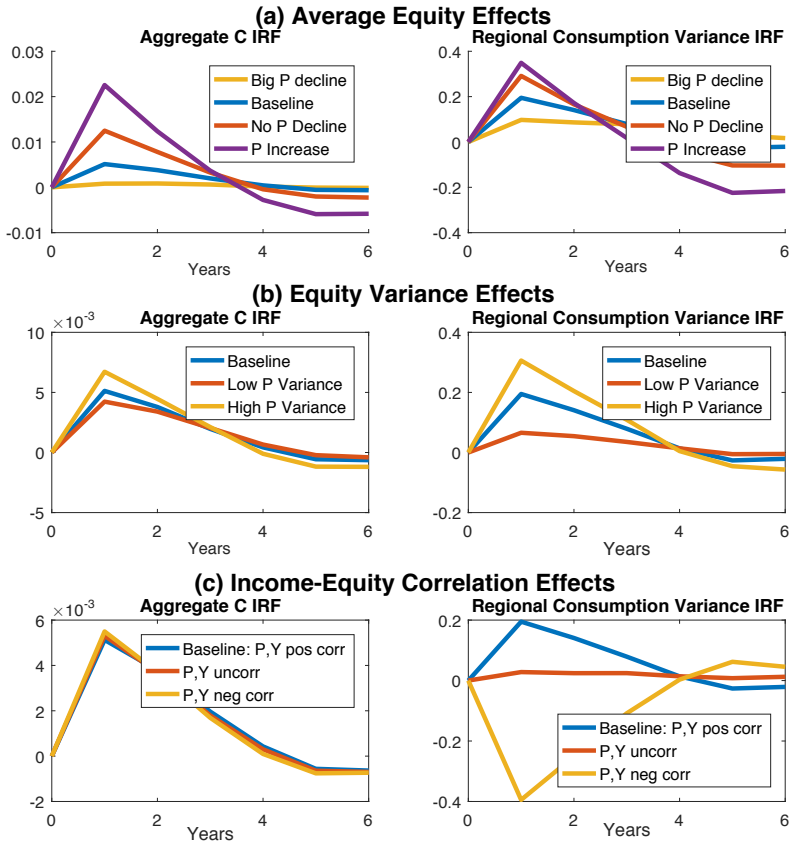
6.4.2 Housing Equity Statistics and the Consequences of Monetary Policy

The distribution of equity in 2008 is different in three ways from the stochastic steady-state. What are the independent roles of changes in the level of equity, the regional variance of equity, and the correlation between regional income and regional equity for aggregate responses? Furthermore, why does a change in interest rates in 2008 increase regional inequality much more than in the stochastic steady-state? We now show the role of each of these components of the distribution of equity in shaping aggregate consumption and regional inequality responses to interest rate declines.

Panel (a) of Figure 15 shows the effect of changing the mean of the equity distribution at the time of a decline in interest rates. As the level of house prices rises, both aggregate consumption and inequality respond more to the same decline in interest rates. Aggregate consumption responses increase with house prices, since more households are pushed into the part of the equity distribution where refinancing decisions respond to rate changes, and households also have more equity to consume con-

[42] That is, only those households whose decision to refinance is triggered by the rate decline contribute to the extensive margin and interest rate impulse response. For example, in a world where there was no cost to refinancing so that all households refinance every period, a decline in interest rates would not generate any increase in refinancing and would have little effect on consumption.

Figure 15: Effects of Changing Equity Distribution



Impulse response functions show the change in log aggregate consumption and in log consumption variance across regions in response to a one percentage point reduction in mortgage rates. The baseline economy includes a 12.5% aggregate house price decline. In the top panel, the "big price decline" calibration features a 25% decline in house prices and the price increase calibration features a 12.5% house price increase. The variance of equity and its correlation with income are fixed at the 2008 calibration across all simulations. In the middle panel, the high variance calibration doubles the difference between high and low house price regions, while the low variance calibration halves it. All economies feature the same baseline decline in house prices and correlation with income. In the bottom panel, the baseline calibration has income and house prices being positively correlated across regions. In the other two calibrations they are uncorrelated or negatively correlated. All simulations feature the same baseline decline in house prices and variance across regions.

ditional on refinancing. Why does the response of regional inequality to interest rates rise with house prices? This occurs because refinancing decisions and the resulting consumption responses are highly convex in equity. In regions with negative equity, few households refinance in response to declines in interest rates so that consumption responses are always near zero. This is true whether households are deeply underwater or only mildly underwater. In contrast, the consumption response to interest rates increases rapidly with equity as equity rises. This means that shifting the distribution of the equity to the right has no effect on consumption responses for those on the left side of the distribution, while it increases them substantially for those on the right side of the distribution. Since initial levels of equity and income are positively correlated in our 2008 baseline, this means that an increase in average equity increases consumption inequality.

Panel (b) of Figure 15 shows the effects of changing the variance of equity across regions. An increase in the variance of equity amplifies the aggregate consumption response to rate declines but also amplifies the response of consumption inequality. The intuition is almost identical to that for

the effects of mean shifts and again follows immediately from Figure 14. Moving households with low equity to even lower equity has no effect on consumption responses to monetary policy, since these households do not refinance anyway. In contrast, additional equity amplifies the consumption response of those households on the right side of the distribution with substantial equity.

Finally, Panel (c) of Figure 15 shows the effects of changing the correlation between income and equity. If income and equity are uncorrelated, then interest rate declines have almost no effect on regional inequality and when income and equity are negatively correlated, declines in interest rates substantially reduce regional inequality. In contrast, the correlation between income and equity has almost no effect on aggregate consumption impulse responses. The intuition for inequality effects is straightforward: Consumption *levels* are higher in high income than in low income regions. Consumption *responses* to interest rate changes are higher in high equity than in low equity locations. When income and equity are correlated, this means that interest rate declines exacerbate the initial consumption inequality. If income and house prices are instead uncorrelated, as in the stochastic steady-state, then changes in consumption when interest rates fall are largely uncorrelated with initial levels of consumption, and if they are negatively correlated then inequality is reduced.

The intuition for the lack of aggregate effects is slightly more subtle and reflects two offsetting forces. Overall, consumption growth is largest for regions with high house price growth and low income, since they have more equity and are also more liquidity constrained. However, low income regions also have lower initial consumption levels than high income regions. This means that the change in consumption *levels* for high equity high income regions is similar to that of high equity low income regions, so that changing the proportion of such regions by altering the correlation between income and house prices has a negligible effect on aggregate consumption responses.

7 Robustness and Model Extensions

In this section, we consider many extensions of our baseline model and show that time variation in the consequences of monetary policy continues to hold. Ultimately, this is because our conclusions are crucially driven by two simple and robust model features, which also hold in the data: 1) Underwater households cannot refinance without putting new cash into the house, which makes policy functions highly non-linear in equity. 2) The regional distribution of equity changes across time. Getting time variation in the fraction of households that respond to interest rate changes relies only on the interaction between these two features. Since this non-linearity is highly robust across models (and indeed in any reasonable model, it should always be the case that underwater households cannot refinance and tap into home equity), it is not surprising that our conclusions are equally robust, but we now elaborate on several extensions that are particularly interesting.

7.1 Accounting for Additional Heterogeneity

7.1.1 ARM Shares

Our baseline analysis assumes that all mortgages have fixed rates. In reality, a substantial fraction of mortgages have adjustable rates that reset as current interest rates decline even if households do

not refinance. The presence of adjustable-rate mortgages (ARMs) has some potential to change the refinancing channel of monetary policy. On the one hand, with ARMs, payments will decrease when rates decline even if households do not refinance. To the extent that borrowers have higher MPCs than lenders, this should amplify the spending response to monetary policy (e.g., Auclert, 2015). On the other hand, the presence of ARMs reduces the interaction between cash-out decisions and interest rate declines since households do not need to accelerate equity extraction in order to take advantage of lower rates today. This makes cash-out-based spending less responsive to monetary policy.

Accounting for ARMs is also potentially important for regional inequality, since panel (a) of Appendix Figure A-7 shows that ARM shares of outstanding loans in November 2008 were higher in MSAs with lower equity.[43] This is because ARM shares increased more strongly during the boom years in areas with larger price increases, which subsequently experienced larger busts. Since ARMs were more prevalent in low equity regions, it is possible more mortgages experienced rate declines in these regions despite our previous evidence that fixed-rate mortgages were less likely to adjust.

The importance of ARM offsets in the data depends crucially on what fraction of ARM borrowers actually saw rate resets between November 2008 and mid-2009. There are a number of reasons why not all ARM borrowers benefit from rate declines: (i) Most ARMs are "hybrids" with initial fixed-rate periods of 3, 5, 7 or 10 years, during which borrowers must refinance to lower rates. (ii) ARMs typically have "rate floors," which are often set at the initial interest rate of the loan. (iii) The length of the ARM fixed-rate period often coincides with the length of an interest-only (IO) period during which the borrower only pays interest but does not amortize principal. When the IO period ends, required payments jump up, which can more than offset simultaneous rate decreases. Panel (b) of Appendix Figure A-7 shows the share of ARMs that experience significant rate reductions of 1 percentage point or more from November 2008 – June 2009 against $E^{med}_{j,Nov2008}$.[44] Clearly, differences across MSAs in rate resets are muted relative to differences in ARM shares in (a). These differences in ARM resets with equity are also less than half the differences in FRM refinancing propensities.[45] Thus, overall, declines in effective rates were still larger in MSAs with higher equity.

To explore the role of ARMs for our theoretical conclusions, we solve a version of the model in which some households borrow using FRMs while others borrow using ARMs. We assume that ARMs in the model adjust every period one-for-one with the current mortgage rate, and we calibrate the share of ARMs across regions to match the variation in panel (a) of Appendix Figure A-7. In light of the above discussion, this substantially overstates the actual regional variation in ARM resets in response to QE1, and so is a very conservative upper bound on the extent to which ARMs change our conclusions. Panel (a) of Appendix Figure A-8 shows that even under this conservative calibration, modeling ARMs has a negligible effect on our conclusions. This is because even in the low equity regions, FRM shares remain large, and overall spending responses to interest rate declines are dominated by cash-out effects. Low house price regions have more ARM but less FRM rate resets. On net,

[43] Calza, Monacelli, and Stracca (2013) also show that ARMs are more prevalent in Spain and Italy, which experienced large recessions, than in Germany and France.

[44] We also require that the recorded required monthly payment not increase over the same period (which may indicate that the loan's IO period expired).

[45] The regression line in panel (b) has a slope of -0.059 (s.e. = 0.009). Meanwhile, regressing the 6-month refinancing propensity (from January-June 2009) on $E^{med}_{j,Nov2008}$ yields a coefficient of +0.147 (s.e. = 0.028).

the FRM effect dominates so that there are more rate reductions in high equity regions. Moreover, in low equity regions there is no cash-out activity since there is no equity to remove, while in high equity regions there is a significant cash-out and spending response to rate declines.

7.1.2 Preceding Booms

We also explore whether the boom-bust nature of the Great Recession substantially affects our conclusions. Since regions with the largest house price declines previously had the largest house price booms, perhaps the Great Recession itself dampened rather than amplified inequality. This could in turn change the implications of monetary policy for inequality. In panel (b) of Appendix Figure A-8 we repeat our baseline exercise but in a model where the house price bust is preceded by a boom.[46] This figure shows that our conclusions are unchanged. Monetary policy is slightly more effective and the trade-off with inequality is slightly reduced relative to the economy that experiences a bust with no boom, but the difference relative to our base case is small.[47]

7.2 Endogenous Cash-out

In our baseline model, we abstract from the distinction between cash-out and non-cash-out refinancing by assuming households always extract all available equity when refinancing. We make this assumption largely for tractability, but we now show it makes little difference for our conclusions. In particular, panel (a) of Appendix Figure A-9 shows that results are very similar in a model where households can choose between a cash-out refi, modeled as before, and a pure rate refi, in which they lower their rate but do not cash-out any equity. Allowing households to choose between cash-out and non-cash-out refi makes little difference because households in high equity locations typically extract equity when refinancing anyway, and households in low equity locations on average have little equity to extract, so that the distinction between a cash-out and a rate refi is less relevant.[48]

7.3 Interest Rate Process

In our baseline model, we assume mortgage rates are constant across time and explore responses to one-time unanticipated declines in these rates. In response to QE1, mortgage rates declined and remained low for an extended period of time. In addition, this one-time shock illustrates the mechanism in the simplest environment and increases the computational tractability of the model, which allows some of the robustness exercises in this section. However, we have explored robustness under alternative assumptions for the stochastic process on r^m. In particular, panel b) of Appendix Figure A-9 shows results in a version of the model where r^m follows an AR process, with persistence of 0.89 and

[46] Specifically, we assume that in the period before the Great Recession, there is an aggregate house price increase of 10% and a regional shock of $\pm 7.5\%$ that is perfectly negatively correlated with the shock during the bust. This roughly captures house price movements in the last year of the housing boom.

[47] This difference arises because the fraction of underwater households during the bust is slightly reduced. Note also that we do not recalibrate parameters in this exercise to match the effects of QE1, but they are similar to the baseline.

[48] In this model, the overall cash-out share of refinancing when mortgage rates decline is 65 percent as compared to roughly 50 percent in the data just after QE1. However, our model does not allow households to access equity by selling housing, so it is not surprising that cash-out refinancing is a larger share in the model. Including capital gains when loans are prepaid due to moving in measured equity extraction would raise the empirical equity extraction share substantially.

standard deviation of 0.0055, picked to match the behavior of 30-year mortgage rates from 1990-2015. The results with this interest rate process are similar to those of the baseline model. Since this has little effect on our conclusions, we use the simpler one-time shock as our baseline experiment.

Our baseline model also assumes that monetary policy lowers the long-term mortgage rate r^m but has no effect on the short-term rate r. This describes the behavior of rates during QE1 since short-term rates were at the zero lower bound. However, more conventional monetary policy would typically result in both rates falling simultaneously. Panel (c) of Appendix Figure A-9 shows that results are similar when we instead lower both r and r^m to maintain a constant spread.

7.4 General Equilibrium Effects

Our baseline model assumes that 50 percent of mortgage payments are ultimately made to non-US consumers. Since large fractions of mortgage debt are held by commercial banks, which are owned in part by foreign shareholders, it makes it difficult to precisely measure the ultimate recipient of mortgage payments. In panel (d) of Appendix Figure A-9, we show the aggregate response of the economy under two extreme assumptions about who holds mortgage debt. This also allows us to assess the importance of some general equilibrium forces for our results. As the share of foreign lenders declines, the importance of equilibrium effects rises and aggregate consumption responds less to interest rate reductions. However, in the short-run, these effects remain positive even in the unrealistic case where all mortgage payments go to domestic households, so that none of our conclusions about the trade-off between short-run stimulus and inequality are altered. Again, the reason for this is that cash-out refinancing is an important part of the stimulus effects. The behavior of lenders is irrelevant for cross-region inequality, since we assume that lenders are equally distributed geographically.

We choose to model income and house prices as exogenous, as it simplifies the analysis substantially and allows us to provide more transparent intuition for our main mechanism. However, endogenizing income and house prices in a more fully fledged DSGE framework should only amplify all of our conclusions. We find that refinancing activity and consumption responses to interest rates are stronger in regions that are already doing relatively well. If some portion of spending is on non-tradable goods and if greater mortgage borrowing drives up house prices, then income and house prices will rise more in initially well-off locations, amplifying initial inequality. Similarly, endogenizing aggregate income and house prices will only amplify aggregate time variation as long as greater aggregate spending generates greater aggregate income, as in New Keynesian models. In fact, in the representative agent model of Greenwald (2016) it can be shown quantitatively that endogenizing house prices indeed mildly amplifies the feedback from equity shocks to monetary stimulus.[49]

Our model also holds housing fixed and does not allow households to buy or sell housing. However, introducing a construction sector and endogenizing housing should again amplify our results, for the same reasons that endogenizing income and housing would amplify them. When households have more equity, there is more scope to purchase larger houses and increase housing demand, which will both drive up house prices and amplify initial equity differences and drive up construction and income, amplifying initial income differences.

[49] We thank Dan Greenwald for computing these results.

Finally, our model takes both r and r^m as given and does not impose that the liquid asset market a clears. However, we find that when r^m falls, savings in liquid assets modestly increase. If we imposed asset market clearing, this would lead to a small decline in r, which would increase consumption through standard intertemporal substitution channels. We also find that the strength of these a responses increases in the strength of the refinancing response to r^m. This means that imposing asset market clearing would complicate the model but again mildly amplify our effects.

8 Interaction with Fiscal and Macroprudential Policies

In this section, we study mortgage modification policies, which capture some elements of policy implemented in the Great Recession, as well as macroprudential policies that alter LTV caps in response to economic conditions. The goal is to gauge whether such policies can mitigate the adverse effects of monetary policy on inequality, while strengthening its effectiveness for aggregate stimulus.

We explore two forms of mortgage modification policies: "debt forgiveness" and "relaxed refinancing requirements." While we intentionally implement these policies in a very stylized fashion in order to starkly illustrate their interactions with monetary policy, one can think of the first policy as capturing some elements of the mortgage write-downs available to some borrowers through HAMP, while the second is more similar to the HARP program.[50] We model debt forgiveness by assuming that a portion of mortgage debt for any household that is underwater during the Great Recession is forgiven. In particular, all households with an LTV greater than γ have their loans forgiven so that LTV=γ.[51] Under the relaxed refinancing requirements policy, we allow underwater households to refinance their mortgage rate without meeting the LTV requirement. To reflect the fact that the practical implementation of these policies explicitly eliminated various appraisal and other fees associated with refinancing, we also assume that under both policies underwater households are able to refinance their mortgages without paying the fixed refinancing cost.

Panel (a) of Figure 16 shows the response to these modification programs, holding interest rates constant. That is, it shows the effects of the programs alone with no simultaneous monetary policy change. The debt reduction program increases total consumption in the economy as it redistributes resources from unconstrained lenders to more constrained borrowers.[52] It also reduces regional inequality, since debt forgiveness is available only to underwater households. In contrast, relaxing refinancing requirements has no effect when interest rates are held constant because underwater households have no reason to refinance, even if it is costless, when rates are constant. Thus, this policy has no effect in

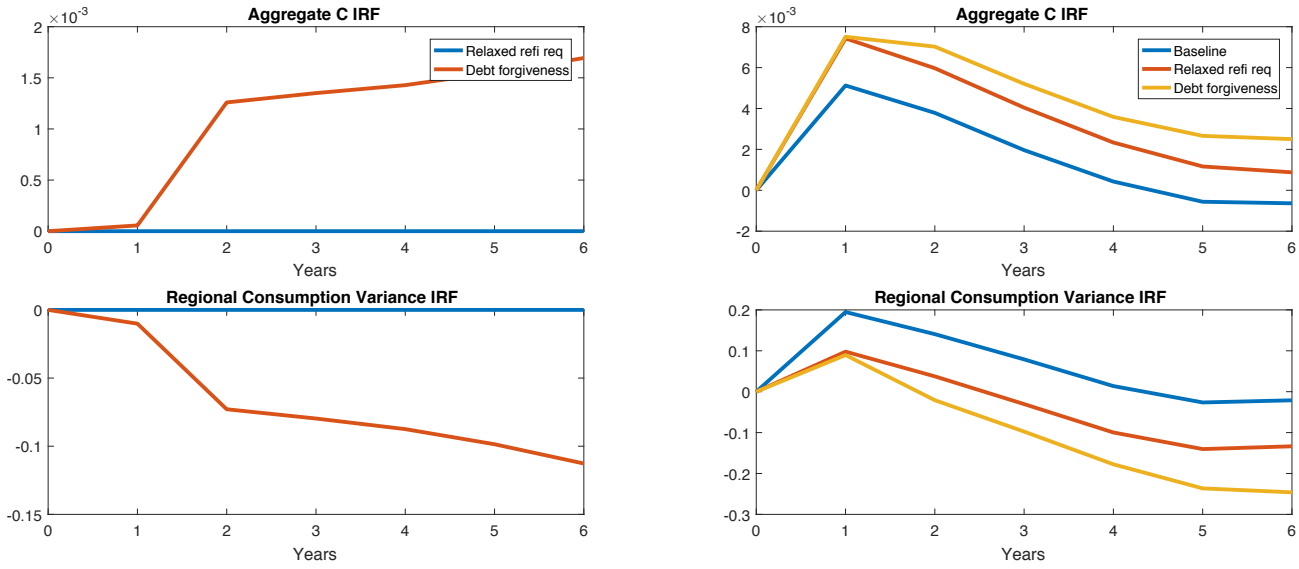
[50] Existing research has shown that institutional features such as servicer participation and market power matter for the consequences of these policies as actually implemented (Agarwal et al., 2015, 2016). Also, especially with modification programs, which in practice usually focus on delinquent borrowers, moral hazard is an important concern that we do not consider. We are not evaluating the specifics of program implementations in the Great Recession or the detailed institutional design of any such programs. We are instead interested in the broad ways in which such programs, independently implemented by the fiscal authority, might affect the consequences of monetary policy.

[51] We account for the negative effect of this policy on lenders, although in reality, lenders would likely be compensated by the government, which in turn would raise taxes. But these taxes would likely be borne disproportionately by the richer lenders. Even with lump-sum taxes from all households, borrowers with forgiven debt would receive a net transfer.

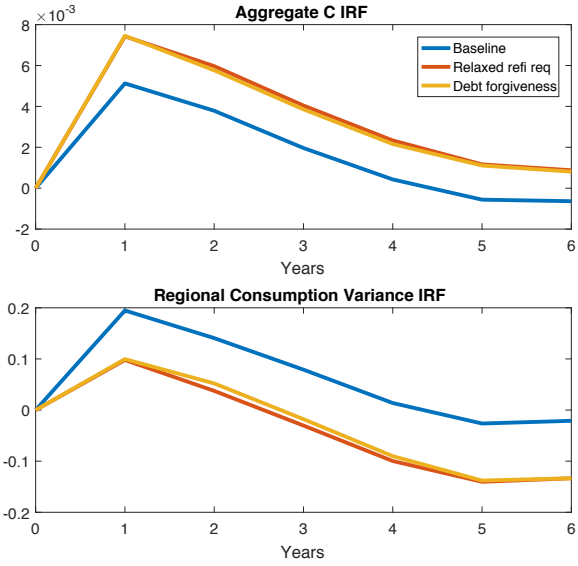
[52] Note that these effects are negligible on impact but grow with time. This is consistent with the empirical and theoretical conclusions in Ganong and Noel (2017) that debt forgiveness has little effect on consumption if, as in our experiment, households still have no equity after forgiveness.

Figure 16: Mortgage Modification Effects

(a): Response to Mortgage Modification w/ Constant r^m **(b):** Response to Mortgage Modification + r^m Decline



(c): Response to r^m Decline, Taking Mortgage Modification Programs as Given



Panel (a) of this figure shows the effects of the debt reduction and relaxed refinancing requirement policies described in the text when interest rates are held fixed. Panel (b) shows the effects of simultaneously reducing rates and implementing the modification policies. Panel (c) shows the change in output and inequality from reducing interest rates and implementing mortgage modification relative to an economy that implements mortgage modification but has no decline in rates.

the model unless accompanied by a reduction in mortgage rates.

Panel (b) of Figure 16 shows the response to *simultaneously* lowering rates and implementing mortgage modification. That is, it shows the combined effects of these policies. For comparison, we also show the baseline economy with a rate decline but no mortgage modification. Relative to the baseline, the combined policies lead to larger increases in spending and smaller increases in inequality. When monetary policy is accompanied by very aggressive debt forgiveness, inequality effects on impact are nearly eliminated, and inequality actually falls in later periods.

Since impulse responses in panel (b) of Figure 16 are computed relative to an economy without either mortgage modification or monetary policy, they tell us the combined effects of these policies. In contrast, panel (c) isolates impacts of monetary policy from direct effects of mortgage modification. In particular, we compute the level and dispersion of consumption with both rate declines and mortgage modification relative to an economy with mortgage modification but no rate decline. Clearly, the presence of either modification program increases the effectiveness of monetary policy and reduces its inequality effects. Interestingly, from the perspective of monetary policy, these two mortgage modification programs work nearly identically. Both policies increase the sensitivity of underwater households to changes in rates and so amplify the response of spending to interest rate declines. The fact that debt forgiveness has larger effects than relaxed refinancing requirements in panel (b) reflects that this policy has direct effects on the economy independently of rate changes, while the relaxed refinancing requirements policies work only through their interaction with interest rates. Since debt forgiveness has both direct effects as well as interaction effects with monetary policy, the combined effects are ultimately larger. However, implications for the efficacy of monetary policy are nearly identical.

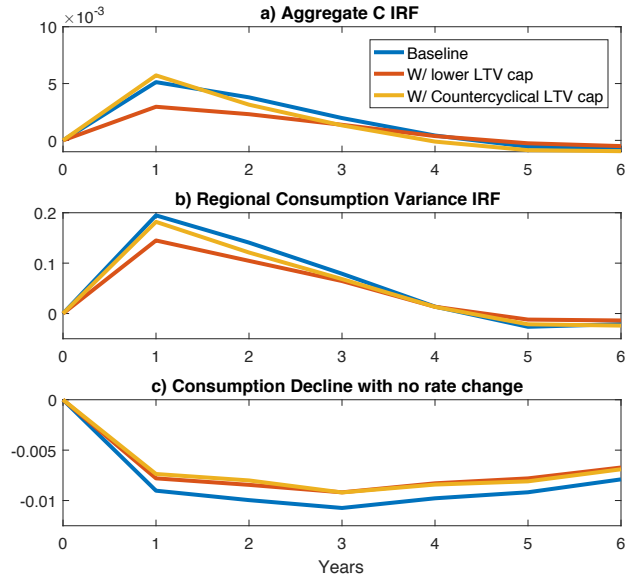
It is important to reiterate the caveat in footnote 50 that this exercise is not intended to evaluate the efficacy of mortgage modification per se, but instead to illustrate the channels through which such policies influence the monetary transmission mechanism. Our results show that certain fiscal policy programs targeting the housing market can increase the aggregate effects of monetary policy while dampening their effects on regional inequality.

The experiments thus far explore the interaction of fiscal and monetary policy in the mortgage market. In our final policy exercise, we explore how a simple form of macroprudential policy can interact with the refinancing channel of monetary policy. In particular, we consider two forms of risk regulation. In the first experiment, we simply lower the LTV cap, γ, from 0.8 to 0.7. In the second experiment, we implement a countercyclical LTV cap that is set to 0.7 during normal times but then rises to 0.9 during the Great Recession. That is, the central bank limits risk during good times but then relaxes constraints and increases liquidity in response to bad shocks.[53]

The red line compared to the blue line in Figure 17 shows the effect of a permanently tighter LTV cap. The first two panels show this reduces both the aggregate stimulus power and the increase in regional inequality generated by monetary policy. This is unsurprising, as tighter borrowing constraints just reduce the extent to which the refinancing channel of monetary policy matters. The third panel of the figure shows that this risk-reduction policy does reduce the depth of the recession, as reduced

[53] In the results shown here, we assume that this increase in the LTV cap during the Great Recession is completely unanticipated and lasts for a single year. We have solved a version of the model where households are aware of the countercyclical LTV policy ex-ante and recessions that trigger this LTV change occur with some small but non-zero probability and it delivered nearly identical results; therefore, we present results for the simpler environment.

Figure 17: Effects of Countercyclical LTV Cap



Panel (a) shows consumption effects of reducing rates in the Great Recession in the baseline economy with a 0.8 LTV compared to an economy with a permanently lower LTV cap of 0.7 and to an economy that has an LTV cap of 0.7 prior to the Great Recession which is then raised to 0.9 for one year. Panel (b) shows effects of monetary policy on inequality in the same three scenarios. Panel (c) shows the depth of the recession *without* any interest rate decline under the three LTV policies.

leverage means that consumption declines less with house prices. The more interesting results are shown in yellow, when the LTV cap is lowered ex-ante but is then increased ex-post after large declines in aggregate house prices. This policy mildly increases the stimulative power of monetary policy and reduces its effects on inequality relative to the baseline (yellow vs. blue), and also reduces the depth of the recession. In this sense, the countercyclical LTV cap dominates either of the other policies: it reduces ex-ante risk, which reduces the depth of the recession. But it also relaxes LTV requirements when and where they bind most, which amplifies the effectiveness of monetary policy and reduces the trade-off between stimulus and inequality.

Thus, while empirically relevant variation in the equity distribution can substantially hamper monetary policy making, there is at least some scope for mitigating these effects. Policies that help underwater households refinance can interact importantly with interest rate changes to amplify the effectiveness of monetary policy. Similarly, well-designed macroprudential policy has scope to reduce risk in the economy while maintaining the strength of monetary policy when it is needed most.

9 Conclusions

The Great Recession led to a prolonged period of monetary stimulus throughout much of the developed world. The effects of these policies are typically studied through the lens of representative agent New Keynesian models, which emphasize the importance of intertemporal substitution. In this paper, we explore a complementary channel of monetary transmission through collateralized lending

and show that understanding this channel requires moving beyond a representative borrower. Nonlinear interactions between collateral constraints, refinancing and spending mean that the *distribution* of housing equity plays a crucial role in the economy's response to interest rate declines.

Using a general equilibrium, heterogeneous agent model of household mortgage borrowing, we argue that the regional distribution of housing equity during the Great Recession led to a substantial dampening of this refinancing channel of monetary policy. Furthermore, large variation in house price growth that was strongly correlated with local economic activity during this recession means that monetary stimulus likely exacerbated existing regional consumption inequality. These theoretical conclusions rest importantly on the distribution of equity, which is assumed away in typical representative agent analyses: under alternative distributions of housing equity, such as that observed in 2001 (which was similar to our stochastic steady-state), monetary policy is much more powerful and can potentially mitigate regional inequality.

We provide evidence of these collateral effects using novel household-level data that include comprehensive information on mortgage debt and refinancing. We show that after QE1, there was an aggregate increase in refinancing but that there was little response in the hardest hit regions, where many households were underwater. The empirical distribution of house price growth was quite different during the 2001 recession: aggregate house price growth was positive throughout the recession, and regional house price growth was uncorrelated with local unemployment. Consistent with our theoretical predictions, there was much more refinancing activity during the 2001 easing cycle than during the Great Recession, and refinancing was actually more common in regions with high unemployment. Thus, the data confirm that time-varying heterogeneity in the distribution of collateral is important for understanding the consequences of monetary policy across time.

Our data come from the US mortgage market, which means our analysis focuses on heterogeneity across regions, since regional house price movements are the dominant source of shocks to households' housing equity. Variation in the distribution of other types of collateral will generate many of the same implications for monetary policy, but the relevant sources of shocks and heterogeneity may differ. For example, sectoral shocks may play an important role in influencing the distribution of collateral across firms and will likely influence the response of investment to monetary policy through similar mechanisms. Variation across time in economic activity and its correlation with housing equity and other forms of collateral is also not unique to the US. Europe has recently experienced persistent cross-country differences in economic growth that are highly correlated with house price movements. While the prominence of fixed-rate mortgages and other institutional features of mortgage contracts differs between the US and Europe and across countries within Europe, Section 7.1.1 shows that our conclusions are not particularly sensitive to variation in fixed-rate shares. We leave a more thorough analysis of the effects of the collateral distribution in these alternative contexts to future work, but our analysis suggests that central banks are likely to face substantial headwinds and important tradeoffs between stabilization and inequality when collateral values are low but highly dispersed and correlated with economic activity. Since the distribution of collateral varies across time, tracking its evolution is crucial for accurately assessing the effects of policy at any point in time (see Fuster, Guttman-Kenney, and Haughwout, 2016, for an effort along those lines).

References

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. 2016. "Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program." *Journal of Political Economy* forthcoming.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2015. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence From the Home Affordable Refinance Program." Working Paper 21512, NBER.

Auclert, Adrien. 2015. "Monetary Policy and the Redistribution Channel." Working Paper, MIT.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2016. "The Aggregate Implications of Regional Business Cycles." Working Paper, University of Chicago.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2015. "House Prices and Consumer Spending." Working Paper 21667, NBER.

Berger, David and Joseph Vavra. 2015. "Consumption Dynamics During Recessions." *Econometrica* 83 (1).

———. 2016. "Dynamics of the U.S. Price Distribution." Working Paper 21732, NBER.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1999. "The financial accelerator in a quantitative business cycle framework." In *Handbook of Macroeconomics*, vol. 1. Elsevier, 1341–1393.

Bhutta, Neil and Benjamin J. Keys. 2016. "Interest Rates and Equity Extraction during the Housing Boom." *American Economic Review* 106 (7):1742–1774.

Brady, Peter J., Glenn B. Canner, and Dean M. Maki. 2000. "The Effects of Recent Mortgage Refinancing." *Federal Reserve Bulletin* 86 (July):441–450.

Caballero, Ricardo and Eduardo Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach." *Econometrica* 67.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. "Housing finance and monetary policy." *Journal of the European Economic Association* 11 (s1):101–122.

Canner, Glenn, Karen Dynan, and Wayne Passmore. 2002. "Mortgage Refinancing in 2001 and Early 2002." *Federal Reserve Bulletin* 88 (December):469–481.

Caplin, Andrew, Charles Freeman, and Joseph Tracy. 1997. "Collateral Damage: Refinancing Constraints and Regional Recessions." *Journal of Money, Credit and Banking* 29 (4):496–516.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2013. "Manufacturing Decline, Housing Booms, and Non-Employment." Working Paper 18949, NBER.

Chen, Han, Vasco Cúrdia, and Andrea Ferrero. 2012. "The Macroeconomic Effects of Large-Scale Asset Purchase Programmes." *The Economic Journal* 122 (564):F289–F315.

Chen, Hui, Michael Michaux, and Nikolai Roussanov. 2013. "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty." Working Paper 19421, NBER.

Dell'Ariccia, Giovanni, Deniz Igan, and Luc Laeven. 2012. "Credit booms and lending standards: Evidence from the subprime mortgage market." *Journal of Money, Credit and Banking* 44 (2-3).

Di Maggio, Marco, Amir Kermani, and Christopher Palmer. 2016. "How Quantitative Easing Works: Evidence on the Refinancing Channel." Working Paper 22638, NBER.

Di Maggio, Marco, Amir Kermani, and Rodney Ramcharan. 2014. "Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging." Working Paper, Columbia Business School.

Fratantoni, Michael and Scott Schuh. 2003. "Monetary policy, housing, and heterogeneous regional markets." *Journal of Money, Credit and Banking* 35 (4):557–589.

Fuster, Andreas, Benedict Guttman-Kenney, and Andrew Haughwout. 2016. "Tracking and Stress-Testing U.S. Household Leverage." Staff Report 787, Federal Reserve Bank of New York.

Fuster, Andreas and Paul S. Willen. 2010. "$1.25 Trillion is Still Real Money: Some Facts About the Effects of the Federal Reserve's Mortgage Market Investments." Public Policy Discussion Paper 10-4, Federal Reserve Bank of Boston.

———. 2016. "Payment Size, Negative Equity, and Mortgage Default." *American Economic Journal: Economic Policy* forthcoming.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. 2011. "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases." *International Journal of Central Banking* 7 (1).

Galí, Jordi. 2009. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press.

Ganong, Peter and Pascal Noel. 2017. "The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession." Working Paper, Harvard University.

Garriga, Carlos, Finn E. Kydland, and Roman Sustek. 2013. "Mortgages and Monetary Policy." Working Paper 19744, NBER.

Gertler, Mark and Peter Karadi. 2011. "A model of unconventional monetary policy." *Journal of Monetary Economics* 58 (1):17–34.

Greenwald, Daniel L. 2016. "The Mortgage Credit Channel of Macroeconomic Transmission." Working Paper, MIT Sloan.

Hancock, Diana and Wayne Passmore. 2011. "Did the Federal Reserve's MBS purchase program lower mortgage rates?" *Journal of Monetary Economics* 58:498–514.

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control* 17 (5):953–969.

Hurst, Erik, Benjamin J. Keys, Amit Seru, and Joseph S. Vavra. 2016. "Regional Redistribution Through the US Mortgage Market." *American Economic Review* 106 (10).

Hurst, Erik and Frank Stafford. 2004. "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption." *Journal of Money, Credit, and Banking* 36 (6):985–1014.

Iacoviello, Matteo. 2005. "House prices, borrowing constraints, and monetary policy in the business cycle." *American economic review* :739–764.

Kaplan, Greg and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2 (4):53–87.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vincent Yao. 2014. "Mortgage Rates, Household Balance Sheets, and the Real Economy." Working Paper 20561, NBER.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity* :215–287.

———. 2013. "The Ins and Outs of LSAPs." Presented at the Federal Reserve Bank of Kansas City's Jackson Hole Symposium.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4):1687–1726.

Mian, Atif and Amir Sufi. 2014. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82 (6):2197–2223.

Rubio, Margarita. 2011. "Fixed- and Variable-Rate Mortgages, Business Cycles, and Monetary Policy." *Journal of Money, Credit and Banking* 43 (4):657–688.

Stroebel, Johannes and John B. Taylor. 2012. "Estimated Impact of the Federal Reserve's Mortgage-Backed Securities Purchase Program." *International Journal of Central Banking* 8 (2):1–42.

Tracy, Joseph and Joshua Wright. 2016. "Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses." *Journal of Urban Economics* 93:60–70.

Vavra, Joseph. 2014. "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation." *Quarterly Journal of Economics* 129 (1).

Winberry, Thomas. 2016. "Lumpy Investment, Business Cycles and Stimulus Policy." Working Paper, University of Chicago.

Wong, Arlene. 2016. "Population Aging and the Transmission of Monetary Policy to Consumption." Working Paper, Northwestern University.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Online Appendix

A.1 MSA Groups Used in Figures 4, 5, and 6

Note: For large MSAs that are subdivided into Metropolitan Divisions, we use the latter (throughout the paper).

MSAs in the quartile with lowest E^{med} in November 2008:
Akron, OH; Anderson, IN; Bakersfield, CA; Bangor, ME; Battle Creek, MI; Bay City, MI; Bradenton-Sarasota-Venice, FL; Canton-Massillon, OH; Cape Coral-Fort Myers, FL; Carson City, NV; Cleveland-Elyria-Mentor, OH; Dalton, GA; Danville, IL; Dayton, OH; Deltona-Daytona Beach-Ormond Beach, FL; Detroit-Livonia-Dearborn, MI; El Centro, CA; Elizabethtown, KY; Elkhart-Goshen, IN; Fairbanks, AK; Flint, MI; Fort Lauderdale-Pompano Beach-Deerfield Beach, FL; Fort Walton Beach-Crestview-Destin, FL; Fort Wayne, IN; Fresno, CA; Grand Rapids-Wyoming, MI; Greeley, CO; Hagerstown-Martinsburg, MD-WV; Hanford-Corcoran, CA; Holland-Grand Haven, MI; Indianapolis-Carmel, IN; Jackson, MI; Jacksonville, FL; Kalamazoo-Portage, MI; Kankakee-Bradley, IL; Lake Havasu City-Kingman, AZ; Lakeland-Winter Haven, FL; Lansing-East Lansing, MI; Las Vegas-Paradise, NV; Madera-Chowchilla, CA; Mansfield, OH; Memphis, TN-MS-AR; Merced, CA; Miami-Miami Beach-Kendall, FL; Modesto, CA; Monroe, MI; Muskegon-Norton Shores, MI; Napa, CA; Naples-Marco Island, FL; Niles-Benton Harbor, MI; Oakland-Fremont-Hayward, CA; Ocala, FL; Orlando-Kissimmee, FL; Oxnard-Thousand Oaks-Ventura, CA; Palm Bay-Melbourne-Titusville, FL; Palm Coast, FL; Panama City-Lynn Haven-Panama City Beach, FL; Pensacola-Ferry Pass-Brent, FL; Phoenix-Mesa-Scottsdale, AZ; Port St. Lucie, FL; Providence-New Bedford-Fall River, RI-MA; Punta Gorda, FL; Redding, CA; Reno-Sparks, NV; Riverside-San Bernardino-Ontario, CA; Sacramento–Arden-Arcade–Roseville, CA; Saginaw-Saginaw Township North, MI; Salinas, CA; San Diego-Carlsbad-San Marcos, CA; Santa Rosa-Petaluma, CA; Sebastian-Vero Beach, FL; Springfield, OH; St. George, UT; Stockton, CA; Sumter, SC; Tampa-St. Petersburg-Clearwater, FL; Terre Haute, IN; Toledo, OH; Vallejo-Fairfield, CA; Visalia-Porterville, CA; Warren-Troy-Farmington Hills, MI; Weirton-Steubenville, WV-OH; West Palm Beach-Boca Raton-Boynton Beach, FL; Wheeling, WV-OH; Winchester, VA-WV; Worcester, MA; Youngstown-Warren-Boardman, OH-PA; Yuba City, CA; Yuma, AZ.

MSAs in the quartile with highest E^{med} in November 2008:
Albany-Schenectady-Troy, NY; Alexandria, LA; Anderson, SC; Asheville, NC; Athens-Clarke County, GA; Austin-Round Rock, TX; Baltimore-Towson, MD; Barnstable Town, MA; Baton Rouge, LA; Beaumont-Port Arthur, TX; Bellingham, WA; Bethesda-Frederick-Rockville, MD; Billings, MT; Binghamton, NY; Bismarck, ND; Blacksburg-Christiansburg-Radford, VA; Boulder, CO; Bridgeport-Stamford-Norwalk, CT; Buffalo-Niagara Falls, NY; Burlington-South Burlington, VT; Cambridge-Newton-Framingham, MA; Cedar Rapids, IA; Charleston, WV; Charlottesville, VA; Cleveland, TN; College Station-Bryan, TX; Corvallis, OR; Cumberland, MD-WV; Dubuque, IA; Duluth, MN-WI; Durham-Chapel Hill, NC; Edison-New Brunswick, NJ; Elmira, NY; Eugene-Springfield, OR; Fargo, ND-MN; Florence, SC; Fort Smith, AR-OK; Glens Falls, NY; Grand Forks, ND-MN; Grand Junction, CO; Great Falls, MT; Greenville-Mauldin-Easley, SC; Harrisburg-Carlisle, PA; Harrisonburg, VA; Hartford-West Hartford-East Hartford, CT; Honolulu, HI; Hot Springs, AR; Houma-Bayou Cane-Thibodaux, LA; Huntsville, AL; Iowa City, IA; Johnson City, TN; Kingsport-Bristol-Bristol, TN-VA; Lafayette, LA; Lake Charles, LA; Lancaster, PA; Lawrence, KS; Lebanon, PA; Longview, WA; Lynchburg, VA; Midland, TX; Missoula, MT; Mobile, AL; Mount Vernon-Anacortes, WA; Nassau-Suffolk, NY; New Orleans-Metairie-Kenner, LA; New York-White Plains-Wayne, NY-NJ; Ocean City, NJ; Philadelphia, PA; Pittsburgh, PA; Pittsfield, MA; Portland-Vancouver-Beaverton, OR-WA; Raleigh-Cary, NC; Reading, PA; Roanoke, VA; Salem, OR; San Angelo, TX; San Francisco-San Mateo-Redwood City, CA; San Jose-Sunnyvale-Santa Clara, CA; Seattle-Bellevue-Everett, WA; Sioux City, IA-NE-SD; Sioux Falls, SD; Spokane, WA; State College, PA; Trenton-Ewing, NJ; Tulsa, OK; Victoria, TX; Wenatchee-East Wenatchee, WA; Williamsport, PA; Wilmington, NC; Yakima, WA; York-Hanover, PA.

A.2 Additional Data Description

A.2.1 CRISM

We start with a 50% sample of all McDash (also known as LPS) mortgages linked to Equifax credit records that were outstanding for at least one month between January 2007 and December 2010. The CRISM data set provides the linked Equifax credit records for each of these loans for the lifetime of the loan, including an additional 6 months before origination and after termination. Equifax is reported as a panel at the consumer level, providing total outstanding debt amounts in different categories (first-lien mortgages, second-lien mortgages, home equity lines of credit [HELOCs], auto loans, etc.). Additionally, in any month, Equifax provides the origination date, amount, and remaining principal balance of the two largest (in balance terms) first mortgages, closed-end seconds, and HELOCs outstanding for a given consumer.

We convert these records into a loan-level panel with each loan's type, origination month, origination amount, termination month, and remaining principal balance for all months that the loan is outstanding. We restrict our sample to those consumers who start our sample with two or fewer loans in each category and never have more than three of any of these types of loans outstanding.[1] This amounts to about 96% of the population of Equifax borrower IDs, and these IDs cover about 90% of the loans in McDash. In creating this loan-level data set, we assume that the month in which the loan stops appearing in Equifax is the month that it was terminated.

The variables that McDash provides are already in the form of a loan-level panel and include: origination date, origination amount, remaining principal balance, termination date, termination type, lien type, interest type, property zip code, and purpose type. We match these to our Equifax panel. We consider an Equifax loan/McDash loan pairing a match if the origination date of the Equifax loan is within 1 month and the origination amount is within $10,000 of the McDash loan. If more than one loan is matched, we use the origination amount, date, termination date, zip code, and termination balance as tiebreakers. We are able to match more than 93% of McDash loans using these restrictions, with more than 80% matching the origination information perfectly (up to $1 in balances due to rounding).

We use the set of Equifax/McDash matched loans as our universe in our analysis. Owing to the restrictions above, this amounts to about 82% of the McDash universe. We also verify that we are correctly measuring the termination date and termination balance using the Equifax records by checking these variables against their McDash counterparts for the matched loans.

A.2.2 Measuring Refinancing Propensities

Our goal is to measure the proportion of outstanding loans in an MSA that were refinanced in a given month. For the denominator, we start with all outstanding first liens (where lien type is measured using the McDash variable) in our Equifax/McDash matched universe, but exclude in each month loans that terminate in the next month because they were transferred to another servicer or terminate for unknown reasons (since we will be looking at the proportion of loans that are voluntarily paid off and refinanced).

We count a loan as being refinanced if: (1) its McDash termination type is a "voluntary payoff," and (2) for that consumer, there is another loan that is opened around the time of the first loan's termination on the same property (i.e. the new loan is a refinance, rather than a new purchase loan). More specifically, the most clear indicator that the new loan was a refinance is if the loan has a

[1] This restriction allows us to infer the origination month, origination balance, and balance of the third largest loan of any loan type even though this information does not appear explicitly in Equifax, where if the third largest loan is also the newest loan, we assume its origination month to be the first month it appears in Equifax. We also drop loans that do not have complete consecutive Equifax records.

matching McDash loan (about 70%), and that McDash loan is marked as a refinance loan (in McDash's purpose type variable). On the other hand, the loan is clearly a new purchase loan if the purpose type is marked as such. However, about 25% of McDash loans have purpose type "Unknown" or "Other," and about 30% of the new loans are not matched in McDash (they only appear in Equifax, since McDash does not cover the entire market) and thus have no purpose type attached.

We thus use the following rules to identify refinances. We start by looking for any loan in the Equifax data set that has an open date within 4 months of the McDash loan's termination date. We find at least one such loan for about 81% of the voluntary terminations in 2008 and 2009. We classify these new loans as a refinance if either:

- The loan also appears in McDash and is tagged as a refinance in the purpose-type variable (61% of the McDash-matched loans).
- The loan also appears in McDash and is tagged as an "Unknown" or "Other" purpose type, and has the same property zip code as the original loan.
- The loan appears only in Equifax but the borrower's Equifax address does not change in the 6 months following the termination of the original loan.

This allows us to compute our measure of interest, the balance-weighted refinance propensity, as (balance outstanding in t-1 of loans that were refinanced in month t) / (balance outstanding in McDash in month t-1 that does not terminate for unknown reason in month t)

As a check, we also calculate the refinance propensities separately for the three different cases above (McDash, known purpose; McDash, unknown purpose; Equifax), and find that these refinance propensities are very similar.

A.2.3 Measuring Cash-outs

To measure cash-out refinancings, we need to both identify refinances and how the balance of the new loan compares to the outstanding balace of the loan(s) paid off in the process.

We begin with Equifax/McDash first liens (again using the McDash lien type variable), and keep only those loans that have a McDash purpose type of refinance or unknown/other. Our algorithm to identify whether our new loan is a refinance is similar to the algorithm above. This time, we look for a loan (or loans) in Equifax that terminate(s) around the time when the new loan is originated and check that this loan looks like it was refinanced. We use McDash refinances rather than outstanding loans as our point of reference for these statistics so that we can better represent all refinances, rather than introducing potential bias through only seeing refinances of McDash loans.

Specifically, we call any loan in the Equifax data set that terminates between -1 and 4 months from our new loan's close date a "linked" loan, including first mortgages as well as closed-end seconds and HELOCs, and we call the new loan a refinance if:

- The loan is a known refinance in McDash. (For 86% of these, we find a linked loan in either McDash or Equifax. For the remaining 14%, we would consider these refinances where there was no previous loan on the property.)
- The loan has an "Unknown" or "Other" purpose type in McDash and a linked loan in McDash that has a matching property zip code.
- The loan has an "Unknown" or "Other" purpose type in McDash and a linked loan that appears only in Equifax, but the consumer's Equifax address does not change in the 6 months after the new loan was opened.

If there is more than one linked loan that is a first mortgage in Equifax, we link only the loan that is closest in balance to the origination amount of the new mortgage. We also allow to be linked only those Equifax loans that exist in the Equifax data for at least three months to prevent the refinanced loan balance from being counted in the old balance of the loan.

For each of these cases, we can then calculate the cash-out amount as the difference between the origination amount on the refinance loan and the balance of the linked loan(s) at termination. In order to capture the correct origination amount on the refinance loan, we want to ensure that we are also including any "piggyback" second liens that are opened with the refinance loan that we find in McDash. Thus, we look for any loan in the Equifax record linked to our refinance loan that has an Equifax open date within three months of our refinance loan and an origination balance of less than 25% of our loan's origination balance if labeled a first mortgage and less than 125% of the refinance loan's origination balance if labeled a HELOC or CES, and add the balance of these piggyback seconds to the refi origination amount when calculating cash-out amounts.[2] To eliminate outliers, we also drop cash-out and "cash-in" amounts that are greater than $1,000,000. These amount to dropping less than 0.05% of the refinance loans.

At the MSA level, this allows us to calculate the amount cashed out relative to the total outstanding balance in month t-1. To estimate total dollar amounts cashed out, we scale up the amount cashed out by the ratio of total housing debt outstanding in an MSA according to the FRBNY Consumer Credit Panel (CCP) relative to the total outstanding balance in our CRISM sample. (The CCP amounts are available as end-of-quarter snapshots, so we interpolate between them to get a monthly series.)

In Figure A-5, we compare the total estimated quarterly cash-out amounts to those estimated on prime conventional loans by Freddie Mac.[3] The figure shows that the two series co-move closely and also show similar levels. The higher level in CRISM is expected, since the Freddie Mac series does not include subprime/Alt-A as well as FHA and VA loans.

A.2.4 Measuring CLTVs/Equity

We start with all matched first-lien McDash loans. For a given month, we take the corresponding Equifax record and assign all outstanding second liens to the outstanding first liens in Equifax using the rule that each second lien is assigned to the largest first lien (in balance terms) that was opened on or before the second lien's opening date. We then add the assigned second lien balance(s) to the McDash balance of our original loan as our measure of secured debt on a property, which is the numerator of CLTV. For the denominator, the estimated updated property value, we start from the appraisal amount of the property at the time of the McDash loan origination and update it based on the local home price index from CoreLogic (using the zip-code-level index if available, and the MSA-level index if not).

Equity is simply defined as $1 - CLTV$. When taking medians within an MSA (our measure $E^{med}_{j,t}$), we weight individual observations by their current outstanding first-lien loan balance (from McDash).

A.2.5 HMDA

For robustness, in Figure A-4 we use a different measure of refinancing activity based on data made available as part of the Home Mortgage Disclosure Act (HMDA), which requires mortgage lenders to report information on mortgage applications and originations. The HMDA data are generally perceived to be the most comprehensive and representative source of information on mortgage appli-

[2] We impose these upper bounds because we want to avoid picking up other first lien mortgages (to purchase another property) the borrower might originate at the same time.

[3] To make the two comparable, we multiply our CRISM total by 1/0.9175, where 0.9175 is the share of mortgage balances in CRISM that is in MSAs (as opposed to micropolitan statistical areas or rural areas) as of November 2008.

cations and originations, with market coverage estimated to be around 90%.[4] For each application, HMDA reports the geographic location of the property, the desired loan amount, the loan purpose (purchase or refinance), and whether the loan application led to an origination, was rejected by the lender, or was withdrawn by the borrower.[5] While the public-use HMDA data contain only calendar year indicators, the private-use version of the data set (available to users within the Federal Reserve System) also contains the exact application date and the exact action date. The action date is the date on which the loan is originated, the application is rejected, or the application is withdrawn. These exact dates make the data suitable for high frequency event studies (see, for example, Fuster and Willen, 2010). In all our analysis using HMDA data, we retain only applications that led to originations (action code $= 1$), and always use the application date (rather than the action date). We drop multifamily properties and mortgages with an origination amount >\$3 million (about 0.015% of loans).

While the HMDA data are ideal for measuring the flow volume of mortgage origination activity across locations, it has two prominent limitations. First, for refinance loans, the HMDA data do not include any information on the loan that is paid off. As a result, we cannot use the HMDA data to estimate the extent to which households are removing cash from their mortgage during the refinancing process—a limitation we overcome with the CRISM data we focus on in the main paper. Second, the HMDA data do not include any information on the loans after they are originated. Thus, HMDA is not informative about how many outstanding mortgages there are in an MSA. The stock of outstanding mortgages is necessary to measure a refinancing propensity.

To obtain an estimate of the number of outstanding mortgages in each MSA, we supplement HMDA with data from the 2008 American Community Survey (ACS), which reports the number of outstanding mortgages (but not their amount) and the number of households for fine geographic areas. Since the ACS samples only a fraction of the population, we scale up the number of households based on Census information on the overall number of households in the US in 2008. We use the same scaling factor for the number of mortgages in each location. By combining the ACS data with the HMDA data, we can compute the number of loan originations per number of outstanding mortgages for each location within the U.S.

The measures of MSA-level refinancing propensities in late 2008/early 2009 are very highly correlated between the HMDA data and the CRISM data, once we account for the lag in CRISM relative to HMDA. The population-weighted cross-sectional correlation between the HMDA refinance propensity in December 2008 and the CRISM refinance propensity in January (February) 2009 is 0.87 (0.88). Pooling the second half of 2008 and the first half of 2009, the correlation between HMDA and CRISM propensities is 0.84 for the one-month-forward and 0.78 for the two-month-forward CRISM propensities.

[4] See, for instance Dell'Ariccia, Igan, and Laeven (2012).

[5] There are actually three designated loan types within the HMDA data: origination, refinancing, and home improvement. We combine the home improvement loans with the refinancing ones in our work below.

A.3 Model Proofs and Description of Solution

We provide here the proof of Claim 1 as well as a description of our computational procedure.

Claim 1 *For fixed* $\{r, r^m, r^0\}$*, the value function is homogeneous of degree* $1-\sigma$*.*
This implies that $V(a, y, q_0, r_0^m, q, r^m, r, F) = J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F)q^{1-\sigma}$ *where* $\tilde{a} \equiv \frac{a}{q}, \tilde{y} \equiv \frac{y}{q}, \tilde{x} \equiv \frac{q_0}{q}$ *and* $J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F) \equiv V(\frac{a}{q}, \frac{y}{q}, \frac{q_0}{q}, r_0^m, 1, r^m, r, F)$*.*

Proof. To show the claim, we proceed by guess and verify. The value functions for refinancing and not refinancing are

$$\begin{aligned}
V^{norefi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F)q^{1-\sigma} &= \max_{\{\tilde{a}'\}} \frac{(\tilde{a}(1+r) + \tilde{y} - \gamma r_0^m \tilde{x} - \tilde{a}')^{1-\sigma}}{1-\sigma} q^{1-\sigma} \\
&\quad + \beta \mathbb{E}[V(\tilde{a}' \frac{q}{q'}, \tilde{y}', \tilde{x}_0'^m, 1, r^{m'}, r', F')(q'^{1-\sigma}] \\
V^{refi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F)q^{1-\sigma} &= \max_{\{\tilde{a}'\}} \frac{(\tilde{a}(1+r) + \tilde{y} - \gamma r_0^m \tilde{x} + \gamma(1-\tilde{x}) - F - \tilde{a}')^{1-\sigma}}{1-\sigma} q^{1-\sigma} \\
&\quad + \beta \mathbb{E}[V(\tilde{a}' \frac{q}{q'}, \tilde{y}', \frac{q}{q'}, r^m, 1, r^{m'}, r', F')(q'^{1-\sigma}] \\
&s.t. \quad \tilde{a}' \geq 0 \\
&log(\tilde{x}_t) = -\mu_q + log(\tilde{x}_{t-1}) - \nu_t \\
&log(\tilde{y}_t) = \mu_y - \mu_q + log(\tilde{y}_{t-1}) + \varepsilon_t - \nu_t
\end{aligned}$$

To complete the proof, notice that the solution is independent of q, so we can eliminate it as a state variable to obtain the transformed value function,

$$\begin{aligned}
J(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F) &= \max\{J^{norefi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F), J^{refi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F)\} \\
J^{norefi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F) &= \max_{\{\tilde{a}'\}} \frac{(\tilde{a}(1+r) + \tilde{y} - \gamma r_0^m \tilde{x} - \tilde{a}')^{1-\sigma}}{1-\sigma} \\
&\quad + \beta \mathbb{E}\left[\left(\frac{\tilde{x}}{\tilde{x}'}\right)^{1-\sigma} J(\tilde{a}' \frac{\tilde{x}'}{\tilde{x}}, \tilde{y}', \tilde{x}_0'^m, r^{m'}, r', F')\right]
\end{aligned}$$

$$\begin{aligned}
J^{refi}(\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, r^m, r, F) &= \max_{\{\tilde{a}'\}} \frac{(\tilde{a}(1+r) + \tilde{y} - \gamma r_0^m \tilde{x} + \gamma(1-\tilde{x}) - F - \tilde{a}')^{1-\sigma}}{1-\sigma} \\
&\quad + \beta \mathbb{E}\left[\left(\frac{\tilde{x}}{\tilde{x}'}\right)^{1-\sigma} J(\tilde{a}' \frac{\tilde{x}'}{\tilde{x}}, \tilde{y}', \frac{\tilde{x}'}{\tilde{x}}, r^m, r^{m'}, r', F')\right] \\
&s.t. \quad \tilde{a}' \geq 0 \\
&log(\tilde{x}_t) = -\mu_q + log(\tilde{x}_{t-1}) - \nu_t \\
&log(\tilde{y}_t) = \mu_y - \mu_q + log(\tilde{y}_{t-1}) + \varepsilon_t - \nu_t
\end{aligned}$$

■

In order to solve the transformed value function, we discretize $\tilde{x}$ using 40 grid points evenly spaced with width *mu* between 0.175 and -0.825, 64 grid points for $\tilde{a}$ between 0 and 1, with more grid points near the lower asset values to account for the concavity of the value function, and 46 grid points for $\tilde{y}$ evenly spaced in logs between -0.5 and 0.5. The stochastic shock can take on 3 values: -1SD, 0, and

+1SD with probabilities computed using the Tauchen algorithm. The model is then solved using value function iteration. Finally, the model is simulated using 50,000 households and 9 regions.

A.4 Appendix Tables

Table A-1: Descriptive Statistics. Data on 381 MSAs; statistics are unweighted.

	Mean	SD	p10	p25	p50	p75	p90
Median equity share (E^{med}), Jan 2007	0.293	0.066	0.209	0.243	0.289	0.345	0.383
Median equity share (E^{med}), Nov 2008	0.160	0.110	0.015	0.120	0.190	0.230	0.261
Δ(House Price Index), Jan 2007-Nov 2008	-0.111	0.113	-0.286	-0.161	-0.086	-0.033	0.010
Δ(Unemployment), Jan 2007-Nov 2008	2.297	1.370	0.700	1.400	2.100	3.000	4.100
Average mortgage characteristics as of Nov 2008							
FICO score	701.8	17.4	679.4	689.4	701.6	714.9	723.5
Current interest rate (%)	6.23	0.16	6.02	6.13	6.23	6.34	6.43
Loan age (months)	39.3	3.8	34.6	36.6	39.1	41.6	44.8
Share jumbos (based on current balance)	0.046	0.048	0.007	0.014	0.029	0.063	0.115
Share adjustable-rate mortgages	0.136	0.098	0.057	0.070	0.098	0.171	0.272
Share GSE securitized	0.615	0.097	0.498	0.572	0.631	0.676	0.721
Share FHA/VA	0.150	0.093	0.051	0.087	0.144	0.194	0.254
Share privately securitized	0.158	0.086	0.083	0.103	0.127	0.180	0.262

Table shows descriptive statistics for all 381 MSAs in our analysis sample, including the distribution of median equity shares (E^{med}) in January 2007 and November 2008, and the distribution of house price and unemployment rate changes (in percent and percentage points, respectively) between January 2007 and November 2008. Finally, the table shows distributions of a number of characteristics of the outstanding mortgages in each MSA as of November 2008 (taking balance-weighted averages within each MSA).

Table A-2: Relationships between Median Equity and Unemployment Increases as well as Various Average Characteristics of Outstanding Mortgages

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ(Unemployment), Jan 2007-Nov 2008	-0.0597***									-0.0107**
	(0.00877)									(0.00473)
FICO score		0.00404***								0.00442***
		(0.000661)								(0.000773)
Current interest rate (%)			-0.242***							0.0142
			(0.0572)							(0.0854)
Loan age (months)				0.00424						-0.00487***
				(0.00426)						(0.00146)
Share jumbos (based on current balance)					-0.240					0.528***
					(0.254)					(0.154)
Share adjustable-rate mortgages						-0.342***				-1.712***
						(0.111)				(0.185)
Share GSE securitized							0.207*			-2.210***
							(0.121)			(0.267)
Share FHA/VA								0.153		-2.034***
								(0.139)		(0.277)
Share privately securitized									-0.271**	-1.501***
									(0.131)	(0.347)
Constant	0.302***	-2.690***	1.650***	-0.0123	0.172***	0.221***	0.0338	0.137***	0.212***	-0.702
	(0.0208)	(0.467)	(0.353)	(0.171)	(0.0131)	(0.0149)	(0.0794)	(0.0274)	(0.0195)	(1.154)
N	381	381	381	381	381	381	381	381	381	381
Adj. R^2	0.36	0.30	0.09	0.01	0.01	0.11	0.03	0.01	0.06	0.79

Table shows the results of regressions of median equity as of November 2008 ($E^{med}_{j,Nov2008}$) on various other variables observed as of November 2008 (and shown in Table A-1). Sample consists of 381 MSAs. MSAs are weighted by their 2008 population. Robust standard errors in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A-3: The Responsiveness of Local Refinancing Activity to Equity and Other Local Characteristics around QE1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
$E^{med}_{Nov2008}\times$ postQE	1.711***	1.161***	0.920***	1.373***	1.742***	1.673***	1.624***	1.574***	1.714***	1.583***	0.365**
	(0.232)	(0.258)	(0.237)	(0.217)	(0.240)	(0.219)	(0.260)	(0.211)	(0.248)	(0.229)	(0.166)
$\Delta UR_{Jan2007-Nov2008}\times$ postQE		-0.0911***									-0.0528***
		(0.0188)									(0.0111)
FICO $\times$ postQE			0.0107***								0.00654***
			(0.00101)								(0.00160)
Current int. rate $\times$ postQE				-0.889***							-0.795***
				(0.115)							(0.172)
Loan age $\times$ postQE					-0.00914**						-0.0182***
					(0.00404)						(0.00407)
Jumbo share $\times$ postQE						-0.913**					-1.038***
						(0.435)					(0.400)
ARM share $\times$ postQE							-0.259				-0.114
							(0.209)				(0.444)
GSE share $\times$ postQE								0.922***			-0.362
								(0.230)			(0.928)
FHA/VA share $\times$ postQE									-0.0437		-0.553
									(0.312)		(0.949)
Private sec. share $\times$ postQE										-0.569***	-1.403
										(0.213)	(1.013)
Fixed effects	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month
Adj. R^2	0.87	0.87	0.88	0.88	0.87	0.87	0.87	0.88	0.87	0.87	0.90
Adj. R^2 (within)	0.18	0.22	0.28	0.25	0.19	0.19	0.19	0.22	0.18	0.20	0.37
Observations	4572	4572	4572	4572	4572	4572	4572	4572	4572	4572	4572

Table shows results from regressions of monthly refinancing propensities in 381 MSAs from August 2008 to July 2009, measured in CRISM, on median equity in each MSA interacted with a dummy "postQE" that equals one for the period February–July 2009, and other variables also interacted with postQE. All regressions also include MSA and month fixed effects. For reference, the average monthly refinancing propensity over August 2008–January 2009 was 0.4 percent, while over February–July 2009 it was 1.2 percent. Descriptive statistics of explanatory variables are provided in Table A-1. MSAs are weighted by their 2008 population. Robust standard errors (clustered by MSA) in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A-4: The Responsiveness of Cash-Out Refinancing Activity to Equity and Other Local Characteristics around QE1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$E^{ned}_{Nov2008}\times$ postQE	0.167***	0.138***	0.103***	0.137***	0.170***	0.166***	0.169***	0.161***	0.170***	0.167***	0.0743***	0.0626***
	(0.0224)	(0.0266)	(0.0233)	(0.0210)	(0.0227)	(0.0229)	(0.0266)	(0.0221)	(0.0254)	(0.0253)	(0.0207)	(0.0137)
$\Delta UR_{Jan2007-Nov2008}\times$ postQE		-0.00479***									-0.00327**	
		(0.00176)									(0.00129)	
FICO × postQE			0.000862***								0.000343	
			(0.0000968)								(0.000222)	
Current int. rate × postQE				-0.0789***							-0.0820***	
				(0.00888)							(0.0196)	
Loan age × postQE					-0.000999**						-0.00123***	
					(0.000398)						(0.000410)	
Jumbo share × postQE						-0.0126					-0.00432	
						(0.0407)					(0.0462)	
ARM share × postQE							0.00709				-0.183***	
							(0.0181)				(0.0503)	
GSE share × postQE								0.0422**			0.225**	
								(0.0210)			(0.106)	
FHA/VA share × postQE									-0.0482*		0.166	
									(0.0286)		(0.108)	
Private sec. share × postQE										0.00162	0.380***	
										(0.0207)	(0.110)	
Refinancing propensity												0.0610***
												(0.00331)
Fixed effects	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month
Adj. R^2	0.80	0.80	0.81	0.81	0.80	0.80	0.80	0.80	0.80	0.80	0.81	0.84
Adj. R^2 (within)	0.12	0.12	0.15	0.15	0.12	0.12	0.12	0.12	0.12	0.12	0.19	0.32
Observations	4572	4572	4572	4572	4572	4572	4572	4572	4572	4572	4572	4572

Table shows results from regressions of monthly equity removed (relative to outstanding balances) in 381 MSAs from August 2008 to July 2009, measured in CRISM, on median equity in each MSA interacted with a dummy "postQE" that equals one for the period February–July 2009, and other variables also interacted with postQE. All regressions also include MSA and month fixed effects. For reference, the average monthly equity removed over August 2008–January 2009 was 0.09 percent of outstanding balances, while over February–July 2009 it was 0.14 percent of outstanding balances. Descriptive statistics of explanatory variables are provided in Table A-1. MSAs are weighted by their 2008 population. Robust standard errors (clustered by MSA) in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A-5: The Responsiveness of New Vehicle Spending around QE1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$E^{med}_{Nov2008}\times$ postQE	0.389***	0.246**	0.227***						
	(0.0851)	(0.113)	(0.0574)						
$E^{med}_{Nov2008}\times$ time trend (monthly)		0.0103	0.0110***						
		(0.00918)	(0.00316)						
(Refinancing propensity)$_{t-2}$				0.0994***	0.116***			0.0977***	0.112***
				(0.0244)	(0.0206)			(0.0224)	(0.0194)
(Cashout/balance)$_{t-2}$						0.485***	0.369***	0.0250	0.0447
						(0.165)	(0.0923)	(0.107)	(0.0626)
Fixed effects	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month	MSA&month
Adj. R^2	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99
Adj. R^2 (within)	0.04	0.03	0.05	0.04	0.04	0.01	0.01	0.04	0.04
Observations	4572	9144	18288	4572	9144	4572	9144	4572	9144
Date range	200808-200907	200801-200912	200701-201012	200808-200907	200801-200912	200808-200907	200801-200912	200808-200907	200801-200912

Column 1 shows results from regression of monthly log(auto sales) in 381 MSAs on median equity (E^{med}) in each MSA (as of November 2008) interacted with a dummy "postQE" that equals one from February 2009 onward. In columns 2 and 3, the sample period is extended and a monthly trend interacted with $E^{med}_{Nov2008}$ is added. In columns 4 and 5, log(auto sales) is directly regressed on two-months-lagged local refinancing propensities; in columns 6 and 7 the same is repeated for local cash-out amounts (relative to outstanding balance). Columns 8 and 9 add both refinancing measures simultaneously. All regressions also include MSA and month fixed effects. MSAs are weighted by their 2008 population. Robust standard errors (clustered by MSA) in parentheses. Auto sales are from R.L. Polk; refinancing propensities and cash-out fractions are measured in CRISM. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Appendix Figures

Figure A-1: Median Equity vs. Fraction of Borrowers over CLTV 80 or CLTV 100 Thresholds, as of November 2008.

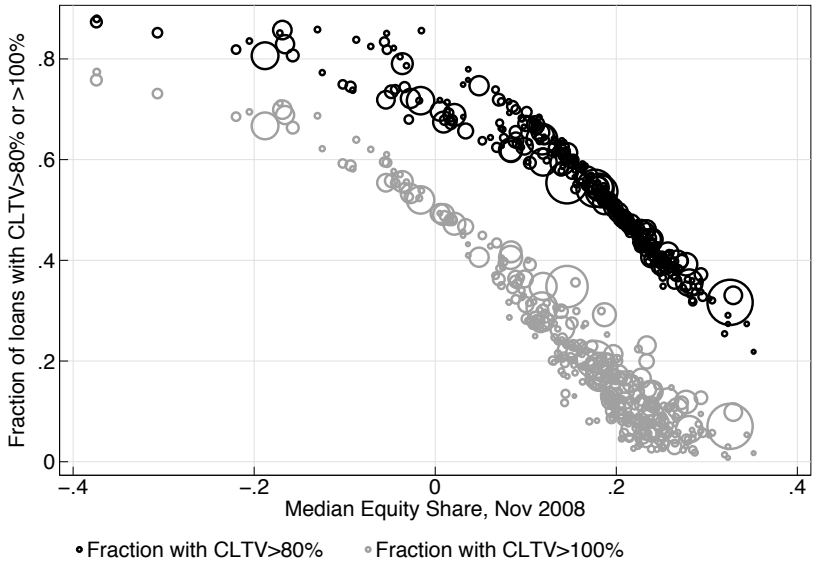


Figure shows the relationship between median equity in an MSA as of November 2008 ($E^{med}_{j,Nov2008}$) and the share of borrowers in the MSA who have an estimated CLTV in the same month higher than 80 percent (black circles) or higher than 100 percent (gray circles). Equity/CLTVs are measured based on CRISM. Each observation is an MSA, with 381 in total (per series). The size of the circle represents the 2008 population of the MSA.

Figure A-2: Mortgage Activity in the US over 2000-2012, Based on HMDA Data

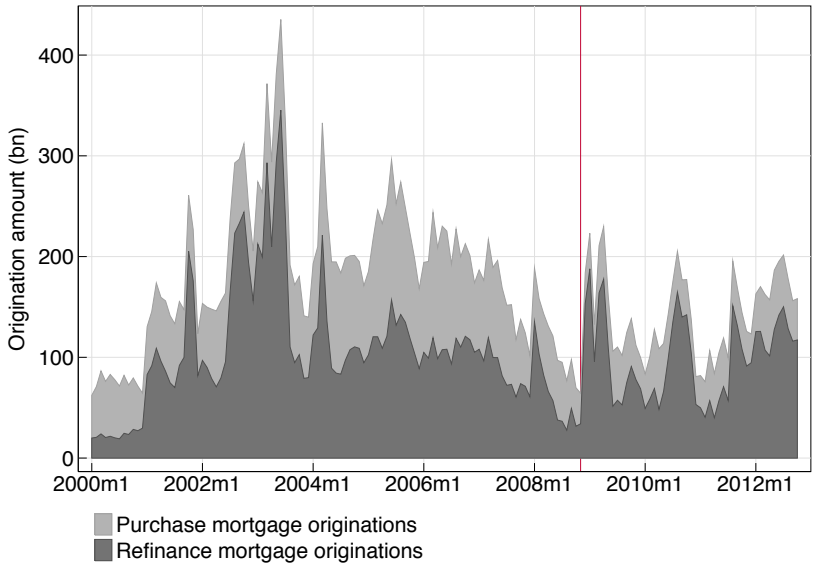
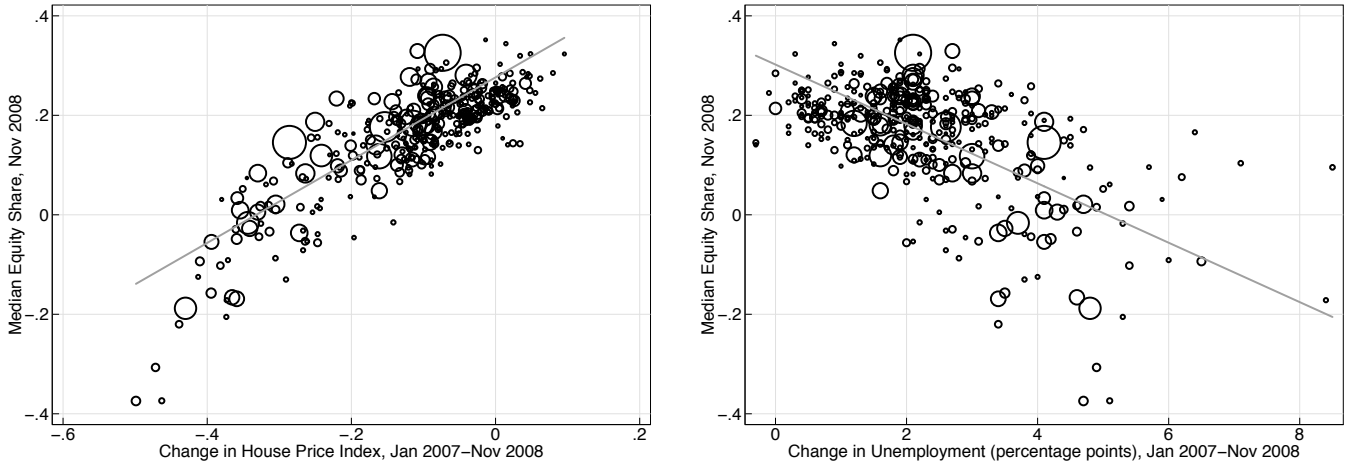


Figure shows mortgage originations on 1-4 unit homes in HMDA, by month in which the borrower applied for the loan. The vertical line indicates the month of the QE1 announcement (November 2008).

Figure A-3: Relationship between Equity and Other Measures of Economic Activity.

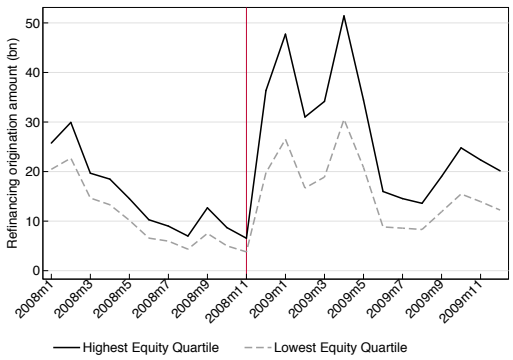
(a): House Price Growth vs. Median Equity **(b):** Unemployment Rate Change vs. Median Equity



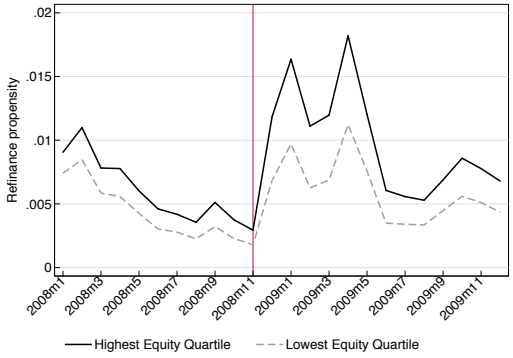
Panel (a) shows MSA house price growth between January 2007 and November 2008 vs. the median borrower's equity (as a share of estimated home value) in the MSA in November 2008. Each observation is an MSA, with 381 in total. The size of the circle represents the 2008 population of the MSA. The figure also shows the simple (population weighted) regression through the scatter plot: a 1 percent decline in house prices is associated with a 0.83 percentage point decrease in median equity (standard error 0.07) with an R-squared of 0.69. Panel (b) shows the change in an MSA's unemployment rate between January 2007 and November 2008 vs. the median borrower's equity (as a share of estimated home value) in the MSA in November 2008. The simple regression line shows that a 1 percentage point increase in the unemployment rate is associated with a 6.0 percentage point decline in median equity (standard error 0.09) with an R-squared of 0.36.

Figure A-4: Mortgage Refinance Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by Median Equity in November 2008, Based on HMDA Data

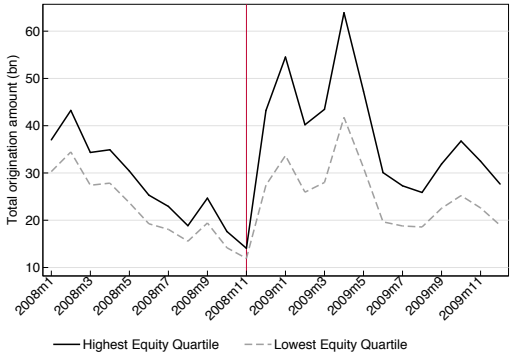
(a): Refinance Origination Volumes



(b): Refinance Propensities



(c): Total Origination Volumes (Refinance and Purchase)



Panel (a) shows total mortgage refinance volume in HMDA by month in which borrower applied for the mortgage, where months are re-defined such that they start on the 25th day of the prior month. Panel (b) shows corresponding refinance propensities, defined as the number of refinance originations in HMDA divided by the total number of mortgages outstanding as measured in the 2008 American Community Survey. Panel (c) shows total mortgage origination volume (purchase and refinance) in HMDA by month in which borrower applied for the mortgage. In all three panels, calculations are done at the level of MSA quartile groups and vertical lines indicate the month of the QE1 announcement (November 2008).

Figure A-5: Estimated Cash-out Amounts from Freddie Mac vs. in our CRISM data

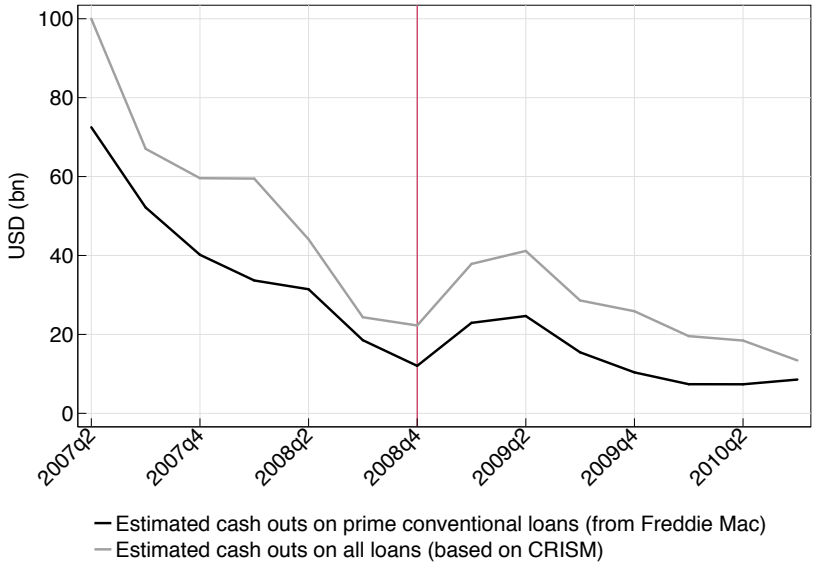


Figure shows estimated quarterly cash-out volumes on prime conventional (non-government) mortgages estimated by Freddie Mac (obtained from http://www.freddiemac.com/finance/docs/q4_refinance_2014.xls), as well as those we obtain based on the CRISM data (which also include FHA/VA loans) after scaling up as explained in Section A.2.3.

Figure A-6: Relationship between Refinancing Threshold, Equity and Income

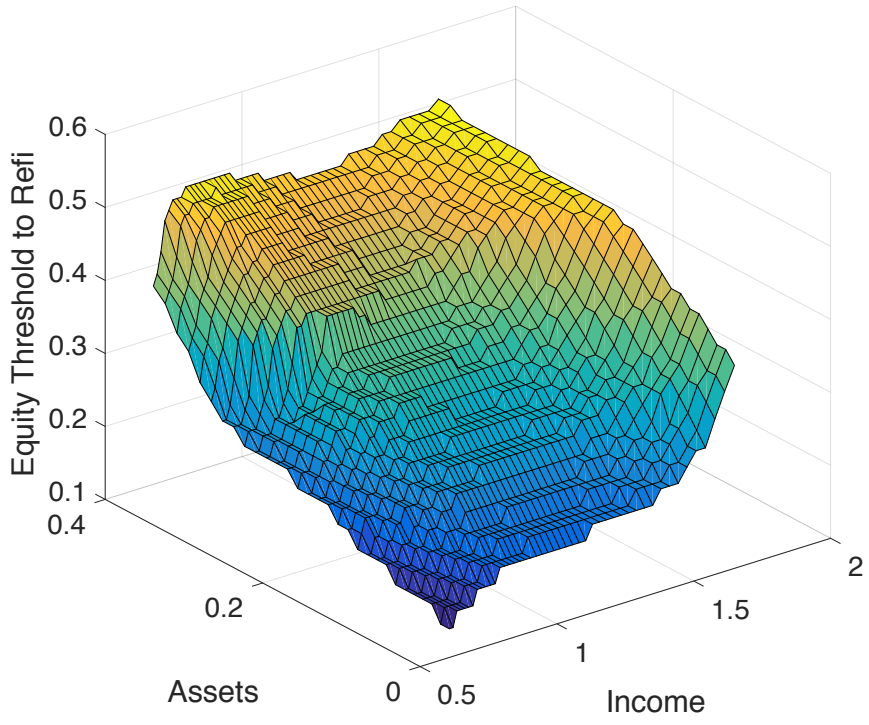


Figure shows refinancing threshold for each asset and income value for low fixed cost.

Figure A-7: Adjustable-Rate Mortgage Shares pre-QE1 and Payment Reductions over Nov 2008 – June 2009

(a): ARM Shares vs. Median Equity **(b):** Sizeable ARM Payment Reductions vs. Median Equity

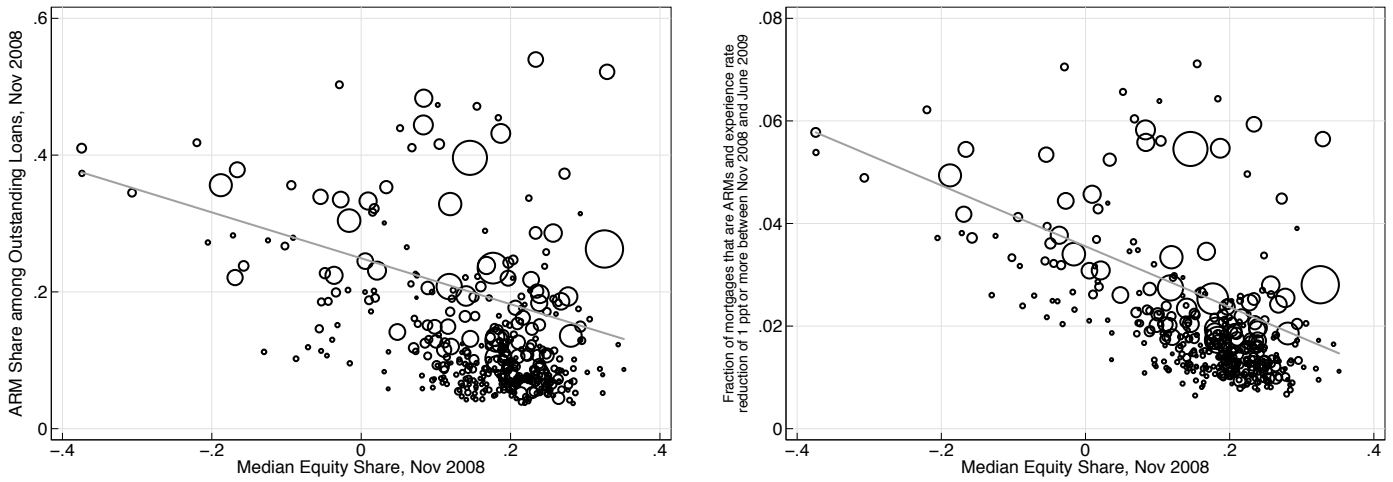
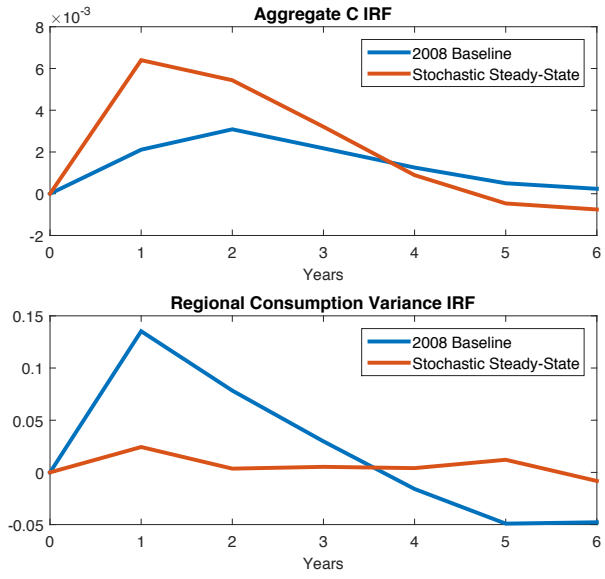


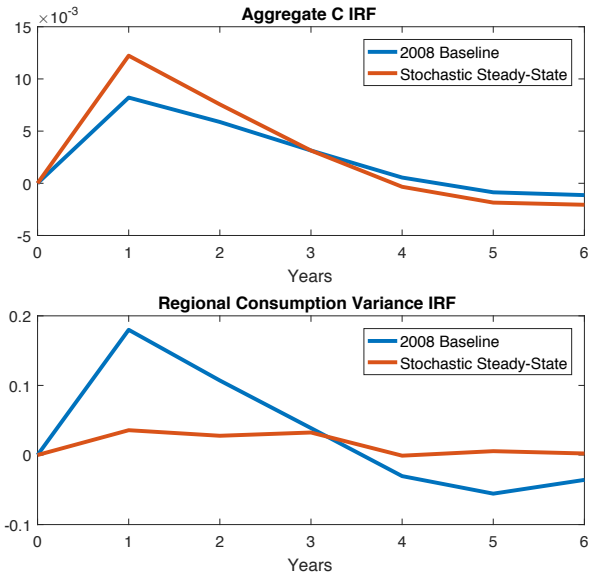
Figure shows scatter plots of the balance-weighted fraction of adjustable-rate mortgages (ARMs) in a given MSA (panel a) or the balance-weighted fraction that are ARMs and experience a rate reduction of 1 percentage point or more over November 2008 – June 2009 (panel b), as measured in the CRISM data, versus the median equity share (E^{med}) within the MSA as of November 2008. The size of the circle represents the 2008 population of the MSA. The gray lines represent simple regression lines fitting the scatter plot (population weighted).

Figure A-8: Accounting for Ex-Ante Heterogeneity

(a): Match ARM-Share



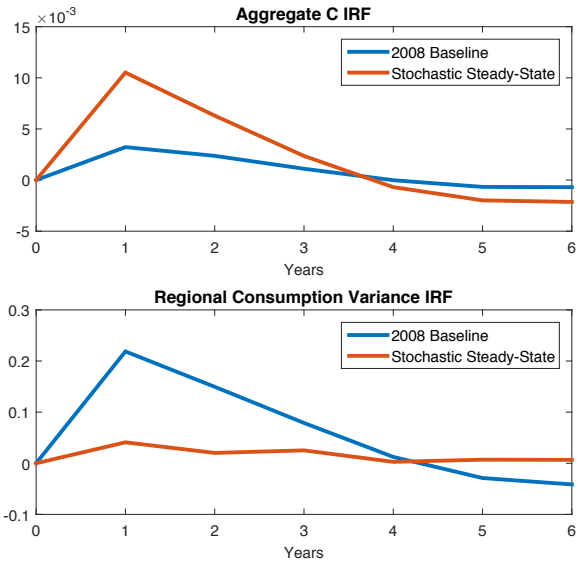
(b): Boom-Bust



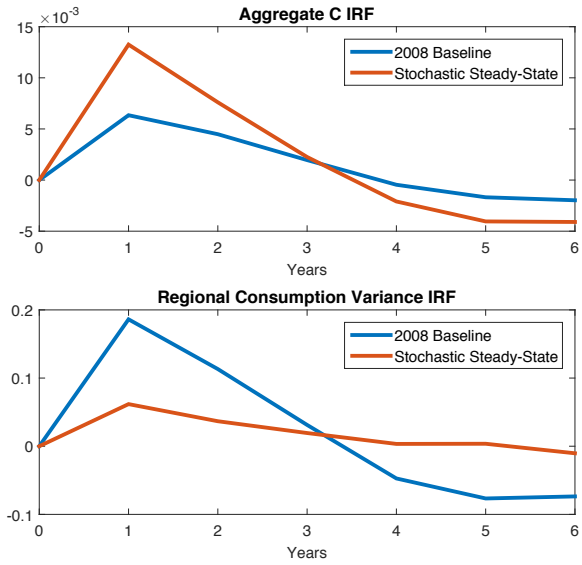
Panel (a) repeats Figure 13 in a model with both FRM and ARMs calibrated to match the regional differences with house prices in the data. Panel (b) repeats Figure 13 in a model where regions with the largest house price declines experience a prior house price boom.

Figure A-9: Robustness Results

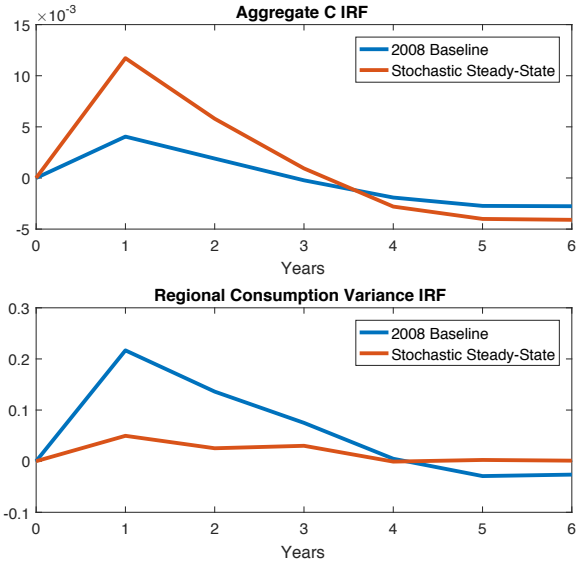
(a): Relaxing Full Cash-out Assumption



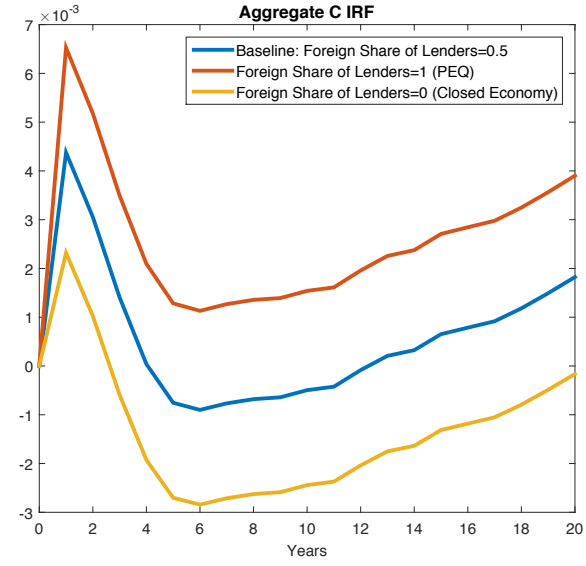
(b): Stochastic Mortgage Rate



(c): Reduce r and r^m



(d): Importance of GE Effects from Lenders' Consumption



This figure repeats Figure 13 in two alternative models. In panel (a) households can choose between full cash-out and no cash-out refinancing. In panel b) interest rate movements are stochastic instead of a one-time shock. In panel (c) r and r^m both decline, with a constant spread. Panel (d) shows aggregate consumption IRFs for alternative assumptions on what share of mortgage debt is held by domestic households, whose income will fall when interest rates decline.