

Regional Heterogeneity and Monetary Policy*

Martin Beraja Andreas Fuster Erik Hurst Joseph Vavra[†]

September 22, 2016

Abstract

We argue that the time-varying *regional* distribution of housing equity plays an important role in shaping *aggregate* consequences of monetary policy. Interest rate cuts can stimulate spending by encouraging households to refinance mortgage debt and cash-out housing equity, but the strength of this channel depends on the joint distribution of housing equity and MPCs. We provide both empirical and theoretical evidence that the distribution of housing equity in 2008 reduced the stimulative effects of monetary policy and led it to increase regional inequality. Using detailed loan-level micro data, we document that QE1 had the smallest effects on mortgage activity and real spending in the parts of the country with the largest employment and house price declines. Using a general equilibrium, heterogeneous household model with realistic refinancing frictions, we match this 2008 evidence, but show that the tradeoff between stimulus and inequality does not hold in general. Monetary policy can have large stimulus effects and potentially reduce inequality under alternative equity distributions such as that observed in the 2001 recession.

*First draft: May 2015. We thank Caitlin Gorbach and Karen Shen for excellent research assistance, and conference participants at the 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

In response to the Great Recession, the Federal Reserve engaged in unprecedented large-scale asset purchases, with the goal of reducing long-term interest rates and stimulating the economy. It is a common view amongst both policy makers and academics that collateralized borrowing in the housing market can play an important role in the monetary transmission mechanism, as lower interest rates incentivize households to refinance their mortgages and extract housing equity to fund current consumption.¹ However, in this paper we argue that this refinancing channel is diminished in recessions that involve large house price declines and where the extent of those declines varies spatially. When there is little housing equity, monetary policy's power to stimulate the aggregate economy by encouraging refinancing and equity extraction is diminished. Furthermore, we show that monetary policy is likely to amplify regional inequality within a monetary union when the hardest-hit regions also experience the largest house price declines and have little home equity to extract.

This exactly describes economic conditions in the US and Europe during the Great Recession, as house prices declined most strongly in locations with the largest declines in economic activity (e.g. Nevada or Spain). This implies that the distribution of housing equity during this period, through its interaction with refinancing frictions in the mortgage market, likely reduced the stimulative effects of monetary policy and lead it to increase regional inequality. However, we argue that both the aggregate effects of monetary policy and its effects on regional inequality vary substantially across time. Under alternative distributions of house prices, such as those in the 2001 recession, the refinancing channel of monetary policy has substantial aggregate effects and can reduce rather than amplify regional inequality. Since the distribution of equity both varies across time and changes the consequences of given monetary policy actions, it is crucial for policy makers to track this variation when making decisions.

In the first part of the paper, we explore the regional response within the United States to the first round of the Federal Reserve's large-scale asset purchase program. This program is commonly known as "quantitative easing" and we will henceforth refer to it as QE1. On November 25th 2008, the Federal Reserve announced that it would initiate a program to purchase \$500 billion of agency mortgage-backed securities (MBS) as well as \$100 billion of direct obligations of housing-related government-sponsored enterprises (GSEs).² Given that short-term interest rates were close to zero, the Federal Reserve action was designed to lower long-term interest rates.³

¹See e.g. <https://www.federalreserve.gov/boarddocs/hh/2004/february/testimony.htm> and <https://www.newyorkfed.org/newsevents/speeches/2012/dud120106.html>.

²Agency MBS include those guaranteed by the GSEs Fannie Mae and Freddie Mac, as well as those guaranteed by the government agency Ginnie Mae. The GSEs securitize a large portion of conforming mortgages within the U.S. Conforming mortgages are mortgages with a size below a fixed amount (the "conforming loan limit") made to borrowers who meet certain quality and loan-to-value thresholds. Borrowers who do not meet the quality and loan-to-value thresholds may in some cases get a loan insured through government programs such as the Federal Housing Administration (which are then securitized in Ginnie Mae MBS).

³In particular, the stated goal of QE1 was to promote mortgage activity. In their announcement of the purchase of mortgage securities, the Fed said "[The policy] is being taken to reduce the cost and increase the availability of credit for the purchase of houses, which in turn should support housing markets and foster improved conditions in financial

30-year fixed-rate mortgage rates, as measured by the Freddie Mac Primary Mortgage Market Survey, hovered between 6.0 and 6.4 percent from July 2008 through late November 2008. The week before the QE1 announcement, the mortgage rate was at 6.04 percent. However, by the first week of December 2008—a week after QE1—the mortgage rate had fallen to 5.53 percent. By December 31st 2008, the mortgage rate had fallen even further to 5.1 percent. The 30-year fixed-rate mortgage rate remained around 5 percent through May of 2009.⁴

Using a variety of loan-level data sources, we document four facts with respect to mortgage originations in response to QE1. First, we show that QE1 did lead to a boom in mortgage originations. The increase in mortgage originations was the result of households refinancing their existing mortgages as opposed to purchasing new homes. Second, we use regional variation to show that the distribution of housing wealth within a metropolitan statistical area (MSA) on the eve of QE1 was important for determining the local refinancing response. In particular, we show that refinancing activity increased most in regions where homeowners had more equity. During this time period, there was a strong correlation between local house price declines and local declines in economic activity.⁵ Given this correlation, we also document that the smallest refinancing response in the aftermath of QE1 took place in the locations that were hit hardest by the recession. Our third fact documents that the amount of equity removed from the house during the refinancing process also varied spatially. For those regions where homeowners had little equity remaining in their home, there was little equity removed during the QE1 induced refinancing boom. This result is not surprising. But, if the link between monetary policy and economic activity through the refinancing channel results from equity extraction, this result confirms that the distribution of housing wealth is important for the efficacy of the collateral channel of monetary policy. Finally, we show a link between local refinancing after QE and a measure of local spending. Using data on new car purchases, we show that areas where borrowers refinanced the most in early 2009 were also the same areas in which car purchases increased the most.

In the next part of the paper, we build a general equilibrium model with collateralized borrowing, long-term mortgages, and idiosyncratic and regional income and house price risk to provide further insight on the interaction between the distribution of housing equity, refinancing and the monetary transmission mechanism. We discipline the baseline parameters in this quantitative model by picking them to match the cross-region effects of QE documented in the first part of the paper. In particular, we pick parameters of the model so that it matches the distribution of housing equity, income, and refinancing propensities both before and after QE. After calibrating the model to match these cross-sectional facts, we aggregate the model in general equilibrium to explore the implications of monetary policy for aggregate spending and inequality.

markets more generally”(<https://www.federalreserve.gov/newsevents/press/monetary/20081125b.htm>).

⁴QE1 was extended on March 18, 2009, by another \$750 billion in agency MBS and \$100 billion in agency debt, to be purchased through the end of 2009. Over 2010 to 2014, additional rounds of asset purchases, including both MBS and Treasury securities, were conducted by the Federal Reserve. We focus on the beginning of QE1 because it was largely unanticipated, thus allowing an “event study,” and arguably had the largest effect in terms of leading to a rapid drop in mortgage rates.

⁵See, for example, [Mian and Sufi \(2014\)](#) and [Stroebel and Vavra \(2016\)](#).

Our first theoretical results simply focus on the implications of the interest rate cuts which were actually observed in 2008. Our empirical event study shows a response of refinancing and cash-out to QE, which should lead to increased spending by mortgage debtors, but [Greenwald \(2016\)](#) shows that accounting for offsetting effects by mortgage lenders can substantially dampen aggregate spending responses in general equilibrium. When interest rates decline, this lowers payments for borrowers who then increase spending but reduces income and spending for lenders. In an environment like the US with a large fraction of fixed rate mortgages, this offset can be almost one-for-one. However, we show that in our environment, this offset is not one-for-one: a decline in interest rates does raise aggregate spending modestly in equilibrium in our baseline economy, which is calibrated to match economic conditions in 2008. This is because our model features an important role for cash-out activity in determining spending. Households accumulate equity across 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.

Aggregate spending rises modestly in the model when interest rates decline under 2008 economic conditions, but this comes at the cost of a large increase in inequality across locations. High equity homeowners have stronger refinancing and spending responses to declines in interest rates than low equity homeowners, and since income and equity were highly correlated during the Great Recession, this leads the variance of consumption across regions to rise.

Thus our model implies that monetary policy in 2008 boosted aggregate spending but did so at the cost of greater inequality. While we think it is interesting to model responses of the economy to monetary policy during the Great Recession, the bulk of our analysis moves beyond this particular episode and shows that time-variation in the collateral distribution plays a crucial role in changing the monetary transmission mechanism across time. Through a series of counterfactual exercises, we show that effects of interest rate changes on both aggregate activity and regional inequality change dramatically with the distribution of housing equity in the economy.

As the average level of housing equity in the economy rises, both aggregate spending and regional inequality respond more strongly to declines in interest rates. However, the spending response grows more rapidly with house prices and equity than does 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 decline in both the 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 in recessions with large variance like 2008. Finally, we show that when the correlation between house prices and income is low, interest rate declines no longer increase inequality. However, the correlation between income and house prices has

almost no effect on the aggregate spending consequences of monetary policy. This means that in recessions like 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 make the theoretical tradeoff between monetary policy and inequality worse during the Great Recession, since this period was characterized by large house price declines, and large spatial variance in prices which was highly correlated with the local magnitude of the recession. However, our results show that the tradeoff between stimulus and inequality does not hold in general and is not stable across time. During other periods such as 2001, there was likely no tradeoff at all. It is also important to note that even under a policy mandate which places no weight on reducing inequality and focuses solely on aggregate conditions, one must still pay attention to the regional distribution of collateral. For example, our results show that the distribution of collateral during the Great Recession substantially reduced the stimulative power of monetary policy through the refinancing channel in addition to worsening the trade-off between stimulus and inequality.

Why do changes in the distribution of collateral change the response of the economy to interest rate changes? These effects arise from the non-linear interaction between household equity and refinancing decisions. When interest rates fall, households have an incentive to pay the cost of refinancing immediately so that they can reduce their interest payments as soon as possible. However, households must also satisfy a collateral requirement in order to refinance. Households who are currently underwater must put up additional cash today in order to refinance while households with substantial equity can reduce their interest payments while also extracting equity and increasing consumption today. This ultimately means that many households with positive equity refinance when interest rates fall while almost no households with negative equity do. Households who do not refinance have negligible consumption responses regardless of their housing equity while consumption responses are strongly increasing in equity for households who do refinance. This leads to a consumption response to interest rates which is highly convex in equity as households who are mildly underwater exhibit the same zero response as those substantially underwater while households with substantial positive equity exhibit much stronger consumption changes than those with mildly positive equity. This convexity explains why changing the distribution of collateral affects the response of the economy to interest rate declines. For example, an aggregate decline in house prices which reduces all individuals' equity proportionately shifts reduces the consumption response of households with initially high equity but leaves the response of those with low equity at zero.

These model implications continue to hold under a variety of alternative assumptions. In these extensions we account for the presence of adjustable-rate mortgages, calibrated to match the observed regional share in the data, 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, and explore alternative assumptions on the lender side of the economy about the share of mortgage debt ultimately held by domestic consumers. None of

these changes has any substantive effect on our conclusions.

Our work relates to many separate literatures. First, a vast New Keynesian literature emphasizes intertemporal substitution as the main reason why unanticipated interest rate changes affect household consumption behavior.⁶ While we also emphasize the response of household spending to interest rate changes, we depart from this literature in our modeling of household capital markets. In particular, standard New Keynesian models assume frictionless household capital markets which feature only one-period borrowing and lending. This standard modeling abstraction stands in stark contrast to the reality of the bulk of actual household borrowing.

The vast majority of household borrowing occurs through the mortgage market. Loans in this market are subject to collateral requirements, are typically long-term with fixed nominal payments and can only be refinanced subject to some costly adjustment process. Each of these features differentiates this borrowing from that in the standard model. Together they give rise to what we call a "collateralized lending channel" channel of monetary policy. Importantly, the strength of this channel depends crucially on the distribution of collateral 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, since the same decline in interest rates can have very different consequences in different recessions. 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 distribution of agents in the economy. For example, [Berger et al. \(2015\)](#) argue that changes in the distribution of household mortgage debt and leverage during the housing boom likely contributed to the large decline in spending when house prices subsequently crashed.⁷ Interestingly, we show here that these same leverage patterns also hampered the ability of monetary policy to stimulate the economy. Thus, the large level of leverage not only made the recession worse but also hindered the ability of monetary policy to correct it.

A large existing literature studies the role of collateral constraints in a "credit channel" of monetary policy. 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 output response to a decrease in nominal interest rates by increasing collateral values and relaxing these constraints for net borrowers.⁸ This channel is complementary but distinct from ours as it arises from monetary policy changing the value of collateral which in turn influences economic activity. In contrast, we take the distribution of collateral at a point in time as given and show that this distribution affects the transmission from interest rates to spending. We think both channels are important and that exploring their interactions is an interesting avenue for future work.

⁶See [Woodford \(2005\)](#) and [Galí \(2009\)](#) for canonical expositions.

⁷See also [Vavra \(2014\)](#) and [Berger and Vavra \(2016\)](#) for prices, [Caballero and Engel \(1999\)](#) and [Winberry \(2016\)](#) for investment, [Berger and Vavra \(2015\)](#), and [Bloom et al. \(2014\)](#) for the implications of microeconomic volatility.

⁸See also the literature on the related "balance-sheet" channel of monetary policy. For example, [Gertler and Karadi \(2011\)](#) study unconventional monetary policy in a world where financial intermediaries face balance-sheet constraints.

We are not the first paper 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 \(2016\)](#) study the transmission of monetary policy in adjustable-rate mortgage (ARM) and fixed-rate mortgage (FRM) environments. However, these papers assume a representative borrower and generally abstract from the costs of refinancing, in contrast to our environment which accounts for heterogeneity, incomplete markets and fixed costs of refinancing. This means that their models have no role for the distribution of collateral across borrowers which is the focus of our paper. [Wong \(2015\)](#) uses a model closer to our own which includes borrower heterogeneity and costly refinancing but in partial equilibrium, and she focuses on the role of aging in affecting monetary policy. Since the age distribution changes slowly across time, these age effects are more relevant for cross-country comparisons and long-run secular trends than for shorter-run changes in the efficacy 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 fixed-rate mortgages, which are important in many countries including the U.S., and from cash-out refinancing. We show that refinancing frictions involved with cash-out refinancing lead to an important role for the time time-varying distribution of collateral.

On the empirical front, [Fuster and Willen \(2010\)](#) measure the effects of QE1 on the primary U.S. mortgage market and emphasize differential effects on borrowers with different levels of creditworthiness, while here we emphasize regional disparities.⁹ [Chen, Michaux, and Rousanov \(2013\)](#) investigate the link between macroeconomic uncertainty and cash-out refinancing and [Calza, Monacelli, and Stracca \(2013\)](#) document the importance of variation in mortgage structure for monetary policy across countries. [Bhutta and Keys \(2014\)](#) show that low interest rates increase the likelihood and magnitude of home equity extraction. [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 \(2012\)](#) and [Tracy and Wright \(2012\)](#) studying 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. \(2015a\)](#) use detailed 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. Focusing on an earlier period, [Caplin, Freeman, and Tracy \(1997\)](#) emphasize how in the early 1990s, drops in housing values in some regions impeded the ability of homeowners to refinance, thereby deepening regional recessions.

⁹There is also a growing literature specifically studying the effects of QE. Most of this work focuses 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 on the macroeconomy through the lens of a DSGE model.

Finally, [Fratantoni and Schuh \(2003\)](#) propose a heterogeneous-agent VAR model that incorporates regional heterogeneity in housing markets to study time variation in the pass-through of monetary policy.

Our evidence on the interaction between the distribution of collateral, monetary policy and regional inequality contributes to the literature studying regional stabilization in currency and fiscal unions. The theoretical side includes the pioneering work on optimal currency areas (e.g. [Mundell \(1961\)](#) and [McKinnon \(1963\)](#)) and the more recent efforts by [Gali and Monacelli \(2008\)](#) and [Farhi and Werning \(2012\)](#). On the empirical side, [Nakamura and Steinsson \(2014\)](#) and [Suarez-Serrato and Wingender \(2010\)](#) concentrate on local fiscal multipliers. [Beraja \(2015\)](#) studies federal transfers rules in fiscal unions. [Hurst et al. \(2015\)](#) focus on regional redistribution through the mortgage market. [Lustig and Van Nieuwerburgh \(2010\)](#) study how regional risk-sharing varies with local housing collateral values.

While our analysis focuses on the implications of U.S. mortgage debt, since this is the source of our detailed micro data, it is important to note that the role of the heterogeneous collateral distribution in the monetary transmission mechanism applies more broadly. We concentrate on regional heterogeneity since house price shocks, which move housing equity, have a very large regional component. Variation in the distribution of other types of collateral can generate many of the same implications for monetary policy, but this variation may be driven by very different sources of heterogeneity.¹⁰ For example, industry-level shocks may change the distribution of collateral across firms and alter the response of investment to monetary policy. Our conclusions also extend beyond the U.S.: 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.

The remainder of the paper proceeds as follows: Section 2 describes the data sources for our empirical work. Section 3 explores the spatial distribution of refinancing activity around QE1, while Section 4 explores the spatial distribution of cash-out and spending activity in this same period. Section 5 explores the theoretical implications of these patterns in a quantitative model of refinancing, and Section 6 concludes.

2 Data

To help document the extent to which refinancing patterns differed spatially after QE1, we use a variety of different data sources. Below, we briefly discuss each of these data sources; the Data Appendix that accompanies the paper provides additional detail.

Throughout the paper, we use two measures of refinancing activity computed from two different data sources. Our first measure of refinancing activity uses data made available as part

¹⁰See also [Agarwal et al. \(2015b\)](#) for evidence of incomplete interest rate passthrough for unsecured credit.

of the Home Mortgage Disclosure Act (HMDA), which requires mortgage lenders to report information on mortgage applications and originations. The HMDA data is generally perceived to be the most comprehensive and representative source of information on mortgage applications and originations, with market coverage estimated to be around 90 percent.¹¹ 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.¹² While the public-use HMDA data only contains calendar year indicators, the private-use version of the dataset (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 either 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 this paper, we use the high frequency data to explore the extent to which refinancing activity differed across locations in the months surrounding QE1.

While the HMDA data is ideal for measuring the flow volume of mortgage origination activity across locations, it has two prominent limitations. First, for refinance loans, the HMDA data does 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 household are removing cash from their mortgage during the refinancing process. Second, the HMDA data does 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 overcome the limitations of the HMDA data, we supplement our analysis with additional data sources. First, to obtain an estimate of the number of outstanding mortgages in each MSA, we use 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 only samples 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 either per household or per number of outstanding mortgages for each location within the U.S.

To obtain measures of cash out refinancing and to create a second measure of local refinancing propensities, we supplement our analysis with data from Equifax's Credit Risk Insight Servicing McDash (CRISM) dataset. This dataset merges McDash mortgage servicing records (from Black Knight Financial Services) with credit bureau data (from Equifax). The structure of the dataset makes it possible to link multiple loans by the same borrower together, something that is not possible with servicing data alone, and thus allows us to accurately measure refinancing activity.

¹¹See, for instance, [Avery et al. \(2010\)](#) or [Dell'Araccia, Igan, and Laeven \(2012\)](#).

¹²There are actually three designated loan types within the HMDA: origination, refinancing, and home improvement. We combine the home improvement loans with the refinancing ones in our work below.

Servicing data alone indicates whether an existing loan pays off, but there is no way of knowing whether the borrower refinanced the mortgage or moved to another home. The structure of the CRISM data avoids this problem because multiple loans can be linked to the same borrower. Furthermore, the CRISM data allow us to study cash-out refinancing much better than with servicing data alone. Since we know 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 from the home during the refinancing process. Finally, unlike the HMDA data, the CRISM data provides us with a natural denominator to scale the refinancing activity given that we can measure the stock amount of loans outstanding in a given area in the previous month.¹³ That said, CRISM has somewhat lower coverage than HMDA (it is estimated to cover roughly 65% of the market during the period we study), and does not contain loan application dates. Given that both datasets have different limitations, we use both the HMDA and CRISM data to explore refinancing activity around QE1 and during the broader 2000s.

We also use CRISM data to estimate the fraction of mortgages in different MSAs that have combined loan-to-value (CLTV) ratios above certain thresholds (80 or 100). A CLTV ratio combines the outstanding balances on all mortgages outstanding relative to the home value attached to the loans. To estimate a mortgage's CLTV, we combine balances of first mortgages and potential second liens (closed-end second liens or home equity lines of credit; see e.g. [Lee, Mayer, and Tracy, 2012](#)), and divide by the estimated property value (given by the appraisal value at the time the mortgage was granted, updated using a house price index from CoreLogic).¹⁴ Additional details are provided in the Data Appendix. We use this data to create an indicator of the ability and incentive of local homeowners to refinance their mortgage. Obtaining a new mortgage with a CLTV above 0.8 is generally more expensive than if the CLTV is below this threshold, because the borrower needs to take out mortgage insurance (through a private mortgage insurance company for loans securitized through Fannie Mae or Freddie Mac, or by paying insurance premia to FHA). Furthermore, due to tightened underwriting standards in 2008/9, there were also lenders who were simply unwilling to make high-CLTV loans even for borrowers who were otherwise creditworthy. Mortgage financing with no down payment (i.e., a CLTV of 100), which was relatively easy to obtain during the housing boom, has been practically non-existent since 2007, except through special programs such as the Home Affordable Refinancing Program (HARP). HARP was introduced in March 2009 to help borrowers with mortgages guaranteed by Fannie Mae or Freddie Mac refinance even if they are underwater (or nearly so). However, due to various implementation issues (see, for instance, [Goodman, 2012](#)) the program initially did little to increase refinancing volumes among such borrowers.¹⁵

¹³A detailed description of how we use CRISM to measure origination and refinance propensities as well as how we measure cash-out refinancings is provided in the Data Appendix.

¹⁴We use the zip-code-level index if available, and the MSA-level index if not.

¹⁵Changes to the program in late 2011, often referred to as HARP 2.0, substantially increased the volume of refinancings through the program. One of the main changes was to remove the cap on admissible LTV ratios, which was initially set at 105, then extended to 125 in summer 2009. See for instance <http://www.fhfa.gov/AboutUs/Reports/>

Finally, we supplement our analysis with five other types of data. First, we use the 2007 and 2008 American Community Survey (ACS) to define local demographic controls for each U.S. sub-location. We merge the 2007 and 2008 data to ensure the sample sizes are large enough to minimize measurement error. In terms of demographics, we measure the age composition of each area, the education composition of each area, the fraction of each area that are homeowners, the racial composition of each area, local employment (and unemployment) rates, and the fraction of each area that is a naturalized citizen. In the Data Appendix, we define specifically how these control variables are constructed. Second, we also show aggregate refinancing trends within the U.S. using published statistics from the Mortgage Bankers Associations (MBA) Refinance Index. The index is aggregate and is created to measure refinancing volume within the U.S. based on refinancing applications. Third, as alluded to above, we use house price data from CoreLogic at the MSA level to measure local house price appreciations. Fourth, we measure employment and unemployment rates for each MSA using data from the BLS's Local Area Unemployment Statistics. Finally, we use data from R.L. Polk to measure new car purchases at the MSA level. The data are collected from new auto registrations at the zip code level and can be combined to measure total new car purchase activity for individuals residing in a given location.

3 The Spatial Variation in Mortgage Activity in Response to QE

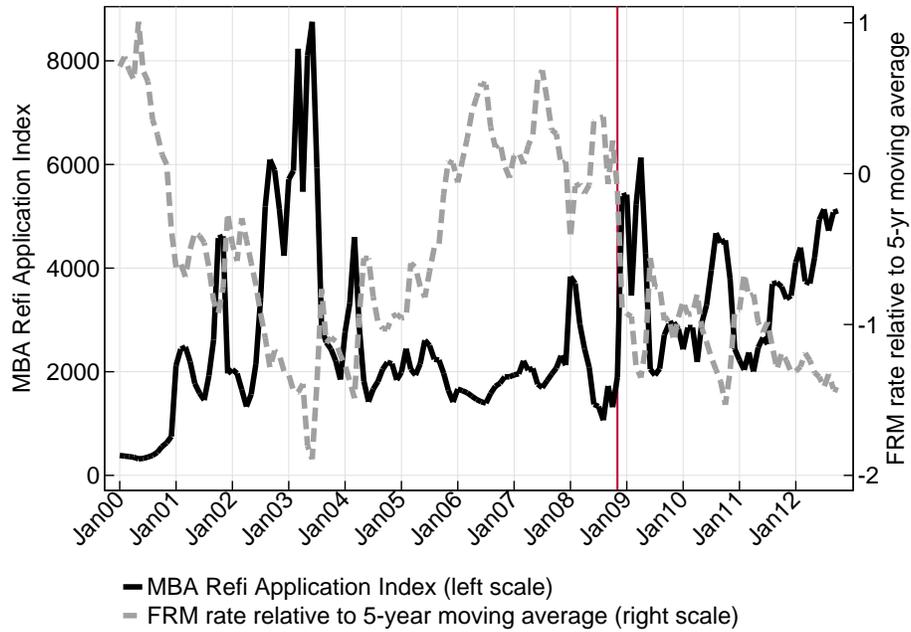
3.1 Aggregate Trends in Mortgage Activity Around QE1

Panel (a) of 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 rate (also from MBA) in month t and the average of the 30-year mortgage rate over the prior five years (dashed line). This metric indicates periods when the 30-year mortgage rate changes discretely in a given month relative to average rates over the prior five years. A few things are noticeable from Figure 1. First, there is a very strong relationship between refinancing activity and 30-year 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 in December 2008 through April of 2009—as measured by the MBA Refinance Index—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 where Federal Reserve policy moves mortgage rates. We primarily focus on the QE1 announcement simply because it was largely unexpected and led to such a sharp drop in mortgage rates.

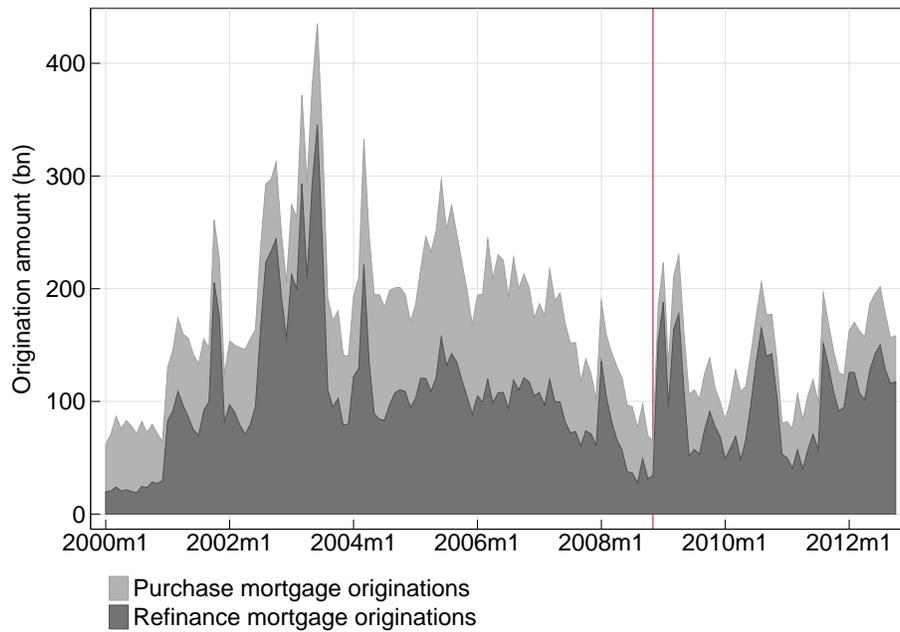
In panel (b) of Figure 1, we plot the time series of monthly mortgage origination activity within the HMDA data over the same 2000 to 2012 period, by month in which the borrower

Figure 1: Mortgage Activity in the U.S. over 2000-2012.

(a) MBA Refinancing Index and Mortgage Rate



(b) HMDA Origination Volume



Panel (a) 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. Panel (b) shows mortgage originations on 1-4 unit homes in HMDA, by month in which the borrower applied for the loan. In both panels, the vertical line indicates the month of the QE1 announcement (November 2008).

applied for the mortgage (not the month in which the loan was ultimately originated, which is usually 1-3 months later). The dark shaded area measures the national monthly dollar volume of refinancing originations. The light shaded area measures the national monthly dollar volume of purchase mortgage originations. The sum of these two areas is total mortgage origination activity. Three things are of interest from this figure. First, like in the MBA Refinance Index, refinancing originations and also total new mortgage originations are highly correlated with 30-year mortgage rates within the HMDA data. The simple correlation of the HMDA refinancing originations (total mortgage originations) with the 30-year mortgage rate relative to its 5-year average is -0.59 (-0.41).¹⁶ Second, like the MBA Refinance Index, mortgage originations increased sharply in December of 2008 after the announcement of QE1. Finally, 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 due to seasonality. For this reason, we focus much of our analysis in this paper on refinancings.

3.2 Spatial Variation in Loan-To-Value Ratios Prior to QE1

To explore spatial variation in loan origination around QE1, we need to define our notion of space. Throughout our paper, we use metropolitan statistical areas (MSAs) to define locations.¹⁷ We begin exploring the local heterogeneity in equity values by documenting the extent to which cumulative loan-to-value ratios (CLTVs) evolved differentially across different areas during the 2007 and 2008 period.

Figure 2 shows the distribution of households with different CLTVs for five different 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). The distributions we show are balance-weighted within MSA.

Panel (a) of Figure 2 shows the CLTV distribution for our five example MSAs in January of 2007. We restrict our analysis to only those individuals in the MSA with a mortgage during that period. We choose January of 2007 because it is a period prior to the national house price decline. For all five of the MSAs, the CLTV distributions are quite similar. To summarize the distribution, we define two variables. $CLTV_{80}$ is the fraction of mortgage owners within the MSA that have a CLTV greater than 0.8. Likewise, $CLTV_{100}$ is the fraction of mortgage holders within the MSA that has a CLTV greater than 1. In January 2007, $CLTV_{100}$ was below 5 percent and $CLTV_{80}$ was around 20 percent for four of the five highlighted MSAs. Borrowers in Las Vegas were highly levered, and also experienced some local property price declines prior to 2007

¹⁶Our HMDA series is a slightly different construct than the MBA Refinance Index. The MBA Refinance Index measures refinance application activity, while in HMDA we only retain applications that ultimately lead to originations. Also, the MBA series does not include broker/correspondent/wholesale origination channels.

¹⁷For large MSAs that are subdivided into Metropolitan Divisions, we use the latter. Many of our control variables are defined at the public use microdata area (PUMA) level. We aggregate the PUMA data to MSAs using the crosswalk available at http://mcdc.missouri.edu/data/georef/zcta_master.Metadata.html.

Figure 2: Combined Loan-to-Value (CLTV) Distributions of Borrowers in 5 MSAs

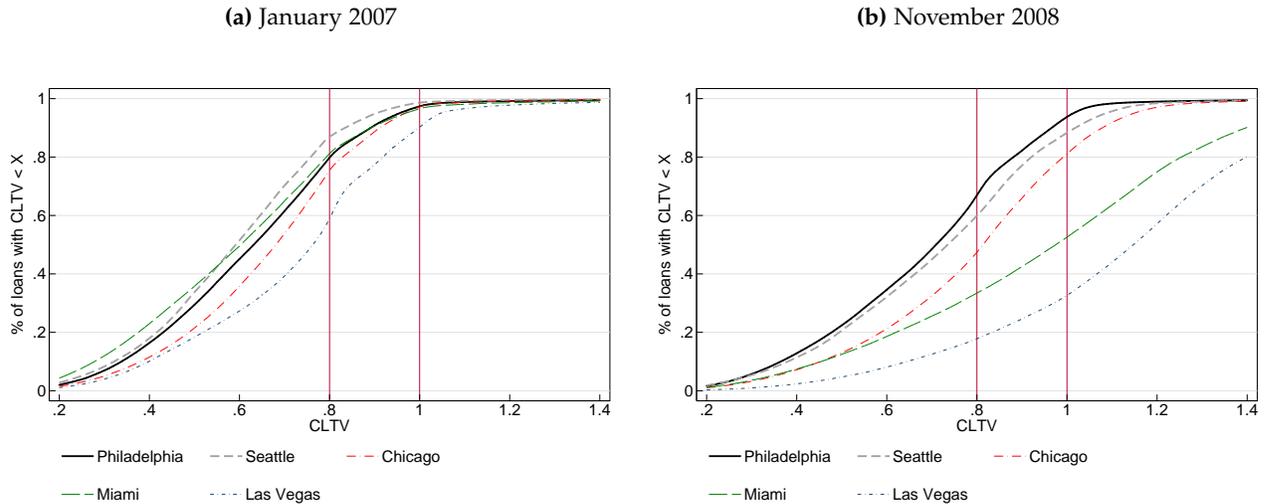


Figure shows the cumulative distribution function of the CLTV distribution of borrowers in five illustrative MSAs in January 2007 (panel a) and November 2008 (panel b). CLTV is measured for each household in an MSA using CRISM data on total current mortgage debt divided by an estimate of their current house value (see text for details). The vertical lines represent CLTV values of 0.8 and 1.0.

so that by January 2007, roughly 10 percent of mortgage owners had CLTVs above 100 percent and roughly 40 percent had CLTVs above 80 percent.

By November of 2008 when QE1 was implemented, there was large variation in the CLTV distribution across MSAs. This can be seen in panel (b) of Figure 2. About 50 percent of mortgage holders in Miami and 70 percent in Las Vegas had a CLTV greater than 1.0. In contrast, Philadelphia and Seattle had only around 10 percent greater than 1.0 while Chicago was in between, with roughly roughly 20 percent of mortgage holders at this high CLTV level. Also seen from panel (b), there was large variation in the fraction of households with a CLTV above 0.8 between Miami/Las Vegas (70 to 85 percent) and Philadelphia/Seattle (35 to 40 percent). In other words, a majority of homeowners in cities like Miami and Las Vegas were underwater at the time of QE1, and even more had high CLTVs that would have made it difficult and expensive to refinance. In cities like Philadelphia or Seattle, a much larger proportion of homeowners had sufficient collateral to easily refinance.

Table 1 shows descriptive statistics for all 381 MSAs in our analysis sample, including the distribution of $CLTV_{80}$ and $CLTV_{100}$ 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 the fraction of households who own their own home and the fraction of homeowners with a mortgage in 2007 and 2008.¹⁸ In January of 2007, essentially all mortgage owners had a CLTV below 1.0 and the vast majority

¹⁸To compute the fraction of households who own a home and the fraction of homeowners who have a mortgage, we use data from the 2007 and 2008 American Community Survey (ACS). The ACS cannot be used to compute monthly statistics. As a result, we show statistics averaged over the entire year.

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

		Mean	St. Dev.	p10	p25	p50	p75	p90
CLTV_80	Jan. 2007	0.316	0.106	0.194	0.238	0.305	0.380	0.460
CLTV_80	Nov. 2008	0.527	0.134	0.375	0.419	0.508	0.617	0.710
CLTV_100	Jan. 2007	0.054	0.038	0.020	0.029	0.044	0.064	0.100
CLTV_100	Nov. 2008	0.212	0.163	0.057	0.095	0.162	0.275	0.480
Homeownership rate		0.779	0.053	0.712	0.751	0.785	0.813	0.839
Fraction of owners with mortgage		0.698	0.077	0.583	0.656	0.707	0.755	0.791
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

had a CLTV below 0.8. For example, in the mean MSA, only 5 percent of all mortgage holders had a CLTV greater than 1.0. The standard deviation across the MSAs with respect to the share of mortgage owners underwater was also small (3.8 percent). This is consistent with the five MSAs we highlighted in Figure 2. However, by November 2008, there was large variation across MSAs in CLTV distributions. Averaging across the MSAs, the mean MSA had 52.7 percent (21.2 percent) of mortgage owners with a CLTV above 0.8 (1.0), with standard deviations of 13.4 percent and 16.3 percent, respectively. This large variation in housing equity across locations will be key to the regional differences in refinancing that we highlight below.

The variation in CLTV across the regions prior to QE1 is driven by the fact that house prices evolved differentially across MSAs during 2007 and 2008. Table 1 shows that there was wide variation across MSAs in house price declines during the Great Recession, and Figure 3 shows that these house price declines were the main determinant of high-LTV shares. On average, a 10 percent decrease in house prices from January 2007 to November 2008 is associated with a 8.8 percentage point higher fraction of mortgage owners with a CLTV above 0.8 in November 2008.

Likewise, there is a large literature showing that house price declines were associated with weakening labor markets during this period (Charles, Hurst, and Notowidigdo, 2013; Mian and Sufi, 2014). Figure 4 shows a simple scatter plot of the percentage point change in the unemployment rate within the MSA between January 2007 and November 2008 against CLTV_80 in November 2008. The places with the largest increases in local unemployment rates were also associated with having relatively more mortgage owners with a CLTV above 0.8.

3.3 Spatial Variation in Mortgage Activity Around QE1

In the month after QE, refinancing activity was much higher in regions where individuals had sufficient equity in their home. As seen from Figure 4, these are also the same places where the unemployment rate was relatively low. To summarize the amount of equity individuals have in their home within a region, we use our measures of CLTV_80 and CLTV_100 (defined above).

Figure 3: House Price Growth January 2007-November 2008 vs. CLTV₈₀ in November 2008

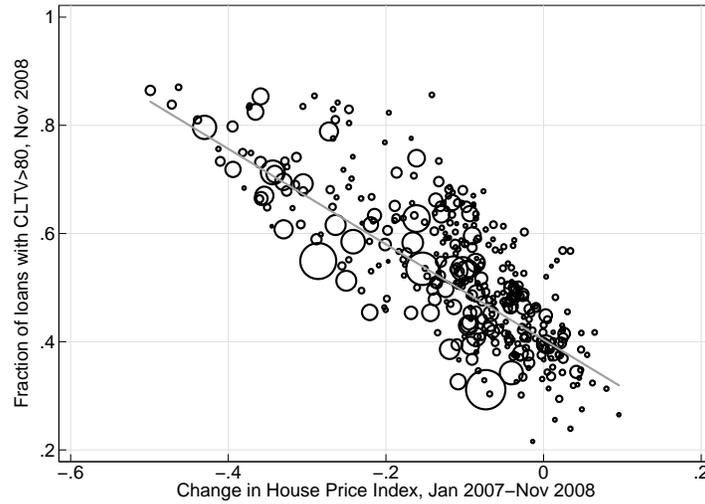


Figure shows MSA house price growth between January of 2007 and November 2008 vs. the share of loans in the MSA in November 2008 with estimated CLTV greater than 0.8 ($CLTV_{80}$). Each observation is an MSA, with 381 in total. The size of the circle represents the 2008 population of the MSA. The figure also shows the simple (population weighted) regression through the scatter plot: a 1 percent decline in house prices is associated with a 0.88 percentage point increase in the fraction of borrowers within the MSA with a CLTV over 0.8 (standard error 0.07) with an R-squared of 0.61.

We focus primarily on $CLTV_{80}$ as our measure of the state of the local housing market prior to QE1, but it makes little difference whether we use $CLTV_{80}$ or $CLTV_{100}$ since they are highly correlated with each other (as shown in Appendix Figure A-1.)

Figure 5 shows different refinancing activity measures for MSAs in the bottom $CLTV_{80}$ quartile and MSAs in the top $CLTV_{80}$ quartile in November 2008. We make quartiles that are population-weighted based on 2008 population numbers from the Census. This ensures that there are the same number of people within each quartile. The top quartile of $CLTV_{80}$ in November 2008 include MSAs like Las Vegas and Miami that had very the vast majority of mortgage owners with a CLTV of greater than 0.8. The bottom $CLTV_{80}$ quartile includes MSAs where most mortgage owners had CLTVs below 0.8. These MSAs include, for example, Philadelphia and Seattle.¹⁹

Panel (a) shows the raw monthly refinancing volume (in billions of dollars) in the HMDA data between January 2008 and December 2009. For the HMDA data, we are focusing on refinancing application dates for originated mortgages so as to better exploit the high frequency response to QE1.²⁰ Refinancing volumes evolved the same between high and low $CLTV_{80}$ MSAs up to

¹⁹Appendix A.2.1 lists the MSAs within each of the $CLTV_{80}$ quartiles.

²⁰Although we show the results monthly, we could have explored weekly refinancing totals. Fuster and Willen (2010) show that refinancing applications in HMDA jumped starting the day of the QE1 announcement. Also, Figure A-2 in the Appendix shows that patterns are nearly identical if we focus on total originations rather just refinancing originations (which is not surprising, given that most mortgage origination activity during this period is refinancing

Figure 4: Percentage Point Change in the MSA Unemployment Rate January 2007–November 2008 vs CLTV_80 in November 2008

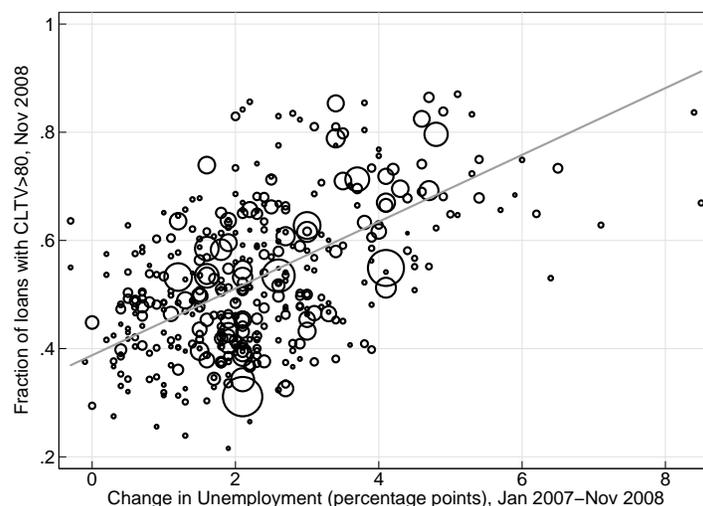


Figure shows the change in an MSA’s unemployment rate between January of 2007 and November 2008 vs. the share of loans in the MSA in November 2008 with CLTV greater than 0.8 (CLTV_80). The simple regression line shows that a 1 percentage point increase in the unemployment rate is associated with a 0.060 percentage point decline in the fraction of borrowers within the MSA with a CLTV over 0.8 (standard error 0.008) with an R-squared of 0.29.

November 2008. Once QE1 was implemented, refinancing activity jumped—but it jumped much more in the high *CLTV_80* MSAs relative to the low *CLTV_80* MSAs.

Panel (b) also uses HMDA to explore the regional variation in refinancing applications to QE1 but looks at refinancing propensities instead of refinancing volumes. We compute refinancing propensities by counting all the refinancing loan applications in the HMDA data during the given month and dividing them by the number of mortgage holders in the MSA (as computed by the 2008 American Community Survey). The patterns in refinancing propensities are similar to the ones in refinancing volumes.

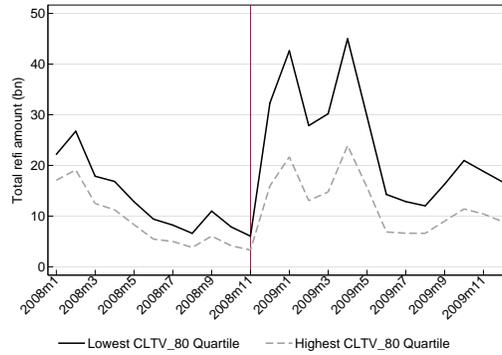
Panel (c) documents the regional difference in refinancing propensities using the CRISM data.²¹ The difference between the timing response in the HMDA data and the CRISM data is driven by the fact that the CRISM data measures actual refinancing originations as opposed to applications. The HMDA data shows that applications jumped immediately in response to QE1. However, the majority of actual originations did not take place until January and February. This is exactly what one should expect given that there is a delay of about 4 to 12 weeks between when

activity, as discussed above).

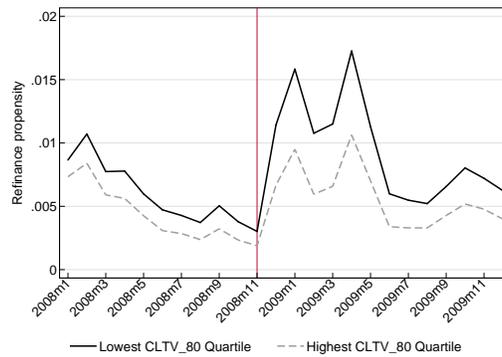
²¹The measures of MSA refinancing propensities in late 2008/early 2009 are very highly correlated between the HMDA data and the CRISM data, once we account for the lag in CRISM relative to HMDA. The population-weighted cross-sectional correlation between the HMDA refinance propensity in December 2008 and the CRISM refinance propensity in January (February) 2009 is 0.86 (0.88). Pooling the second half of 2008 and the first half of 2009, the correlation between HMDA and CRISM propensities is 0.82 for both one-month or two-months forward CRISM propensities.

Figure 5: Mortgage Refinance Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008

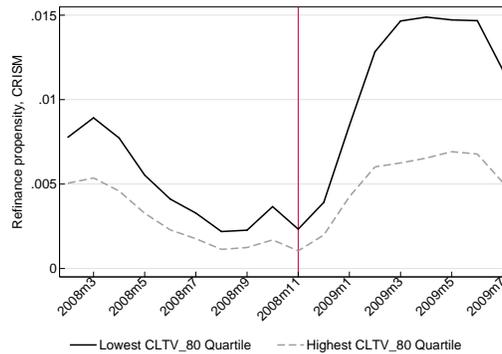
(a) HMDA: Refinance Origination Volumes



(b) HMDA: Refinance Propensities



(c) CRISM: Refinance Propensities



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 refinance propensities in CRISM, defined as the dollar amount of new refinance mortgage originations divided by the outstanding amount of mortgages in the previous month. 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).

a mortgage application is initially made and when the actual mortgage origination takes place.²² The key point from Figure 5 is that regardless of the metric for measuring local refinancing activity, refinancing activity was significantly higher in MSAs where $CLTV_{80}$ was low.

Table 2 shows results from the following regression:

$$\Delta refi_{t,k} = \beta_0 + \beta_1 CLTV_{t,k} + \beta_2 \Delta Unemp_{t,k} + \Gamma X_{t,k} + \varepsilon_{t,k}$$

where $\Delta refi_{t,k}$ is the increase in CRISM refinancing propensity between November 2008 and February of 2009 in MSA k , $CLTV_{t,k}$ is either $CLTV_{80}$ or $CLTV_{100}$ within the MSA during November of 2008, $\Delta Unemp_{t,k}$ is the change in MSA k 's unemployment rate between January 2007 and November 2008, and $X_{t,k}$ is a vector of MSA k control variables defined from the pooled 2007/2008 American Community Survey. The control variables include four variables representing the educational composition of the MSA (e.g., fraction of individuals with schooling equal to 12, 13-15, 16, and 16+), three variables representing the age distribution of the MSA (e.g., fraction of individuals aged 31-45, 46-60, and 61+), the fraction black, the fraction who are citizens, the fraction who own their home and the fraction of homeowners with a mortgage.

Column 1 of Table 2 shows the simple relation between the growth in the MSA refinancing rate after QE1 relative to the month before and the fraction of mortgage owners with CLTV greater than 0.8. With no additional controls ($\Delta Unemp_{t,k}$ or $X_{t,k}$), our estimate of β_1 is -1.43 (standard error = 0.51). To help interpret the magnitudes, we consider a one standard deviation change in $CLTV_{80}$ which was 13.4 percentage points. The simple linear regression with no additional controls implies that a one standard deviation change in $CLTV_{80}$ is associated with a 0.19 percentage point decline in the local refinancing rate. Given the population-weighted mean change in the refinancing rate in the CRISM data was 0.97 percent, a one standard deviation change in $CLTV_{80}$ thus represents about a 20 percent difference in refinancing rates. Put another way, if one went from the 10th percentile of the $CLTV_{80}$ distribution (0.375) to the 90th percentile (0.71), refinancing activity would decline by 0.48 percentage points. Column 2 shows the regression results when $CLTV_{100}$ is the only control variable. The simple relationships between CLTV and refinancing rates is even stronger when we use $CLTV_{100}$ instead of $CLTV_{80}$. For example, a one standard deviation increase in $CLTV_{100}$ results in 0.27 percentage point decline in refinancing activity within the MSA (corresponding to a 28 percent decrease relative to the mean).

In columns 3 and 4, we add $\Delta Unemp_{t,k}$ to the regressions in columns 1 and 2. The coefficients on CLTV weakens in both specifications and the coefficient on the change in the unemployment rate comes in with a negative sign and is statistically significant. In places where unemployment grew the most, refinancing was lower; this occurs above and beyond the effect of differences in equity across regions. We note that these two variables alone explain between 25 and 30 percent of the cross-MSA variation in refinancing propensities.

In columns 5 and 6, we add our X vector of controls to columns 3 and 4. Controlling for

²²During this time both the loan underwriting and closing procedures take place.

Table 2: The Responsiveness of Local Refinancing Activity to Local Collateral and Unemployment around QE1

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
CLTV_80	-1.433*** (0.510)		-0.447 (0.527)		-1.430*** (0.240)		-4.114*** (0.576)
CLTV_100		-1.679*** (0.250)		-1.014*** (0.333)		-1.473*** (0.216)	
D_UR			-0.207*** (0.036)	-0.136*** (0.036)	-0.071*** (0.027)	-0.044 (0.027)	-0.163** (0.068)
Constant	1.738*** (0.307)	1.375*** (0.099)	1.727*** (0.257)	1.554*** (0.092)	0.542 (2.027)	-1.159 (1.972)	2.450 (4.796)
Other Controls	No	No	No	No	Yes	Yes	Yes
N	381	381	381	381	381	381	381
Adj. R ²	0.122	0.266	0.269	0.312	0.578	0.598	0.605
Mean(dep. var.)	0.973	0.973	0.973	0.973	0.973	0.973	2.498

Column 1 shows the results of a regression of the change in the local monthly refinancing propensity (balance weighted) between November 2008 and February 2009 on the fraction of local borrowers with a CLTV greater than 0.8 in November of 2008. Both variables are measured using the CRISM data. Column 2 is the same regression as in column 1 except the independent variable is the fraction of borrowers with a CLTV greater than 1.0 in November 2008. Columns 3 and 4 replicate columns 1 and 2 with the addition of the percentage point change in the unemployment rate between January 2007 and November 2008. Columns 5 and 6 replicate columns 3 and 4 with the addition of a detailed set of local controls (measured in 2008). These controls include variables representing the education, age, citizenship, and racial composition of the MSA as well as the fraction of households that are homeowners within the MSA, and, among those, the fraction that has a mortgage. Column 7 is the same as column 5 except the dependent variable is the change in the local 3 month refinancing rate between January-March 2009 relative to September-November 2008. MSAs are weighted by their 2008 population. Robust standard errors in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

differences in the local composition of households and differences in local home ownership and mortgage ownership rates increase the magnitude of the coefficients on *CLTV_80* and *CLTV_100* relative to columns 3 and 4. Going forward, our preferred specification is column 5, using *CLTV_80* as our measure of local collateral values. The regression coefficient, and thus the economic magnitude of the effects, is almost exactly the same as in column 1 discussed above.

Finally, in the last column of Table 2 we change the dependent variable so it is the cumulative refinancing rate between January and March of 2009, relative to the September-November 2008 cumulative refinancing rate. (We start in January since Figure 5 suggests that some of the QE1-induced refinancing starts being reflected in CRISM in January already.) Otherwise, the specification is the same as in column 5. The mean of the new dependent variable over this period is 2.5 percent. The coefficient on *CLTV_80* implies that a one standard deviation increase in *CLTV_80* reduces refinancing by about 0.55 percentage points, or about 20 percent of the mean refinancing rate. This size of the response to a one standard deviation change in *CLTV_80* over

three months is nearly identical to the one month response highlighted in column 5.

Collectively, the results in this section show that there were large spatial differences in refinancing activity in response to QE1. Places that received the largest house price declines and had the largest declines in employment were the least responsive to QE1 in terms of their subsequent mortgage refinancing behavior.

4 Spatial Variation in Equity Extraction and Spending Around QE1

One natural question is whether the differences in refinancing activity across space resulted in differences in spending differences across space. There are no easily available broad measures of local spending. To overcome the data limitations, we proceed in two ways. First, we explore the extent to which households removed equity from their home at the time of refinancing. Research has shown that households, on average, spend a large amount the equity they remove during the refinancing process on current consumption and home improvements.²³ 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.²⁴

Figure 6: Cash Out Refinancing Activity in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008, CRISM Data.

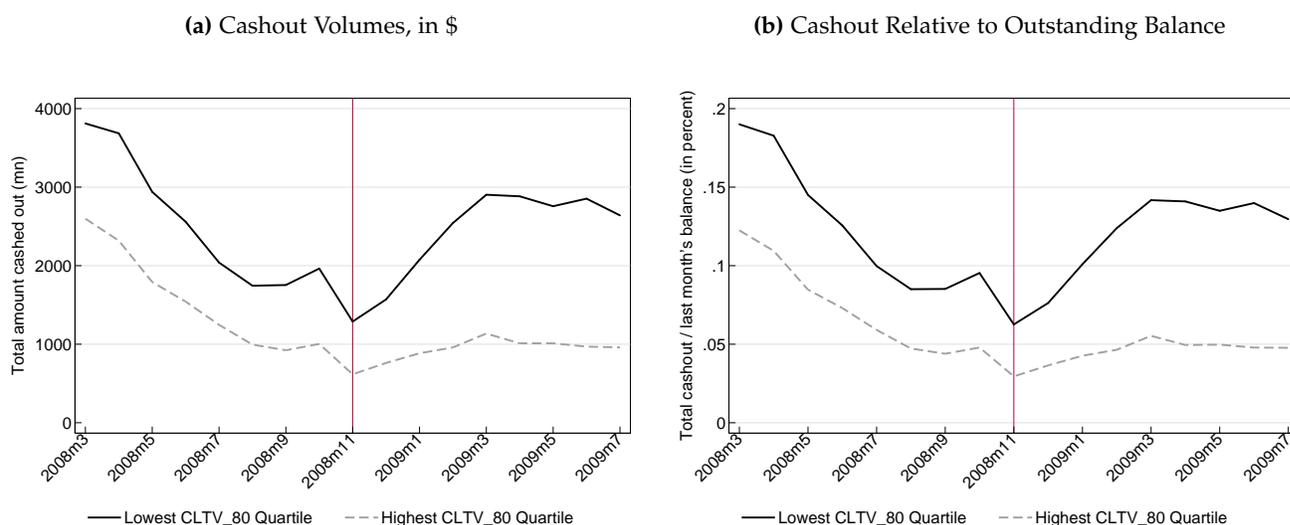


Figure shows the total cash removed from the mortgage during the refinancing process for the top and bottom quartiles of MSAs with respect to CLTV_80 in November 2008. Panel (a) shows dollar volumes, which are calculated by scaling up amounts from the CRISM data to the market level, as explained in the Appendix. Panel (b) shows cash out amounts in percent relative to outstanding mortgage balances (all lien types) in the previous month. MSA quartile groups are defined the same way as in Figure 5. In both panels, the vertical line indicates the month of the QE1 announcement (November 2008).

²³See, for example, Brady, Canner, and Maki (2000), Canner, Dynan, and Passmore (2002), Hurst and Stafford (2004) and Bhutta and Keys (2014).

²⁴See, for example, Mian, Rao, and Sufi (2013).

Figure 6 shows the amount of equity removed during the refinancing process for both the top and bottom quartiles of the $CLTV_{80}$ distribution around the QE period. In panel (a) we show estimated dollar amounts, while in panel (b) we show the equity removed through refinancings in a given month relative to the total outstanding mortgage balance in the prior month.²⁵ We use the CRISM data to measure the total amount of equity removed during the refinancing process. This sums over the people who removed no equity during the refinancing, who put equity into their home during refinancing, and those who extracted equity during the process. In both high and low $CLTV_{80}$ quintiles, borrowers removed equity during the refinancing process. However, in the low $CLTV_{80}$ quartiles there was a much larger increase in equity removed than in the high $CLTV_{80}$ MSAs. Within the low $CLTV_{80}$ MSAs, more than 7.5 billion dollars of equity were removed during the refinancing that took place in three months after QE1 (January-March). Conversely, for the high $CLTV_{80}$ MSAs, only 3.0 billion dollars of equity were removed. In the former case, this was an increase in monthly cashouts of about 0.08 percent of outstanding mortgage balances, while in the latter case (those MSAs in the top quartile of the $CLTV_{80}$ distribution), the increase in monthly cashouts was only about 0.025 percent of outstanding mortgage balances. If we integrate the amount of equity removed during the refinancing process over the 6 month period between January 2009 and June 2009, the bottom quartile of the $CLTV_{80}$ distribution removed an estimated 16.0 billion dollars while the top quartile only removed 6.0 billion dollars.

Table 3 shows results from regressions similar to the ones in Table 2, but using the change in equity removed (relative to outstanding balance) between November 2008 and February 2009 as the dependent variable. Column 1 demonstrates a strong relationship between this variable and $CLTV_{80}$ in November 2008; a one standard deviation change in $CLTV_{80}$ is associated with roughly a 0.016 percentage point change in cashout fractions, or roughly 32 percent of the mean of that variable. Column 2 adds $\Delta Unemp_{t,k}$ (defined above) as an additional control. In column 3, we add both $\Delta Unemp_{t,k}$ and our X vector of local controls. If anything, adding controls only strengthens the coefficient on $CLTV_{80}$. The coefficient on $\Delta Unemp_{t,k}$ comes in negative in both column 2 and column 3. However, it is only significant in column 2; once additional demographic controls are added in column 3, it is no longer significant. How much of the differences in cashouts are due to the fact that high $CLTV_{80}$ places refinance less and how much is due to the fact that high $CLTV_{80}$ remove less equity conditional on refinancing? To answer this question, in column 4 of Table 3 we additionally control for the change in refinancing propensity over the same period. The coefficient on the change in refinancing propensity and on $CLTV_{80}$ are both very significant. This implies that high $CLTV_{80}$ MSAs both refinanced less and removed less equity during the refinancing process conditional on refinancing. This is intuitive, since these places have less equity to remove conditional on refinancing.

Instead of looking at equity removed and implicitly discussing the potential for differential

²⁵Since the CRISM data does not cover the whole mortgage market, we scale up the dollar amounts we get in CRISM as explained in Appendix A.1.2.2.

Table 3: The Responsiveness of Cash Out Refinancing Activity to Local Collateral and Unemployment around QE1

	(1)	(2)	(3)	(4)
CLTV_80	-0.121*** (0.032)	-0.091** (0.037)	-0.159*** (0.021)	-0.106*** (0.019)
D_UR		-0.006** (0.002)	-0.001 (0.002)	0.002 (0.002)
Δ refinance prop.				0.037*** (0.004)
Constant	0.114*** (0.019)	0.114*** (0.018)	0.137 (0.175)	0.117 (0.163)
Other Controls	No	No	Yes	Yes
N	381	381	381	381
Adj. R^2	0.160	0.183	0.318	0.422
Mean (dep. var.)	0.050	0.050	0.050	0.050

Column 1 shows the results of a regression of the change in cash out amount relative to outstanding balance between November 2008 and February 2009 on the fraction of local borrowers with a CLTV greater than 0.8 in November of 2008. We measure the cash out amounts using the CRISM data and divide it by the previous month's mortgage balance outstanding (see Appendix for details). Column 2 replicates the results from column 1 with the addition of the percentage point change in the unemployment rate between January 2007 and November 2008 as an additional control. Columns 3 further adds controls for local demographics and mortgage ownership rates. Column 4 also adds the change in local refinancing propensity between November 2008 and February 2009 as an additional control. The change in refinancing propensity is defined using the CRISM data (and is the same as the dependent variable in columns 1-6 of Table 2). Robust standard errors in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

local spending responses to the QE1 induced refinancing, we now explore more direct measures of spending. Figure 7 shows the normalized auto sales differences across the low and high CLTV_80 MSAs. There were level differences in auto sales across the different CLTV_80 quartiles. To facilitate exposition, we normalize auto sales in both quartiles to 1 in November 2008. As a result, the figure shows percent deviation in monthly auto sales from November 2008. A few things are noticeable from Figure 7. First, prior to QE1, the trend in auto sales was nearly identical in percentage terms between the high and low CLTV_80 quartiles. In both groups of MSAs, new auto sales were 80 percent in early 2008 relative to November of 2008. Second, the trajectory of new auto sales remained constant through February of 2009. This is not surprising given that the refinancing applications that took place in December of 2008 did not result in new mortgage originations until January or February of 2009. Third, and most importantly, after February 2009, auto sales started diverging sharply between the low and high CLTV_80 groups. Averaging over March, April and May of 2009, the increase in auto sales relative to November 2008 averaged 32 percent in the low CLTV_80 quartile MSAs and only 17 percent in the high CLTV_80 quartile MSAs (p-value of difference < 0.01). The timing lines up perfectly with the

Figure 7: New Auto Sales in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008

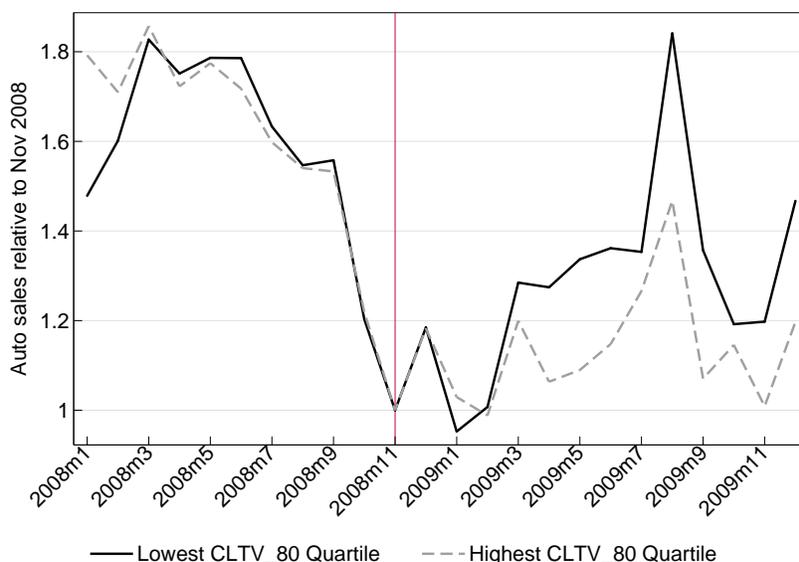


Figure shows the total volume of new car purchases (relative to November 2008) for the top and bottom quartiles of MSAs with respect to CLTV_80 in November 2008. Car purchase volumes come from R.L. Polk. The vertical line indicates the month of the QE1 announcement (November 2008).

expected spending response of the QE1 policy. Moreover, it shows sizable differences in spending at the local level in response to QE1 that are consistent with the differences in the refinancing behavior.²⁶

Table 4 shows MSA-level regressions using the change in the log of new car sales from pre-QE1 (September-November 2008) to post-QE1 (March-May 2009). We take 3-month averages to reduce noise. The first column shows the strong raw relationship with *CLTV_80*. A one standard deviation change in the latter is associated with roughly a one-third of a standard deviation change in $\Delta \log(\text{autosales})$ (since the standard deviation of that variable is 0.155). The second column shows that adding unemployment and MSA characteristics as controls lowers the coefficient somewhat, but the coefficient on *CLTV_80* remains statistically significant at the 5% level. The coefficient on the change in unemployment also comes in negative but it is not significant at standard levels. In columns 3 to 5 we instead directly test the link between refinancing, equity extraction and auto sales, since we are arguing that mortgage refinancing was associated with stronger auto sales over this period. Columns 3 and 4 show that both refinancing and cash-out propensities are individually significantly correlated with higher growth in auto sales. In the last column, we find that if we control for both of them simultaneously, only the former remains

²⁶In addition to the refinancing differences, other differences in mortgage-related outcomes may have contributed to the variation in consumption behavior across MSAs. For instance, defaults in high-CLTV MSAs were substantially higher, which in turn led to a decrease in credit scores and increased difficulty in profiting from lower interest rates (e.g. on car loans) for affected borrowers.

significant; however, this should not be too surprising given that the two variables are strongly positively correlated (with a balance-weighted correlation of above 0.6).

Table 4: The Responsiveness of New Vehicle Spending around QE1

	(1)	(2)	(3)	(4)	(5)
CLTV_80	-0.407*** (0.095)	-0.237** (0.106)			
D_UR		-0.023 (0.016)	-0.022 (0.014)	-0.028* (0.017)	-0.022 (0.014)
Δ refinancing prop.			0.073** (0.030)		0.065** (0.030)
Δ (cashout/balance)				0.488** (0.228)	0.167 (0.186)
Constant	0.207*** (0.048)	2.182** (0.868)	2.048** (0.892)	2.302*** (0.838)	2.116** (0.875)
Other controls	No	Yes	Yes	Yes	Yes
N	381	381	381	381	381
Adj. R^2	0.119	0.274	0.286	0.268	0.286
Mean (dep. var.)	-0.010	-0.010	-0.010	-0.010	-0.010

Columns 1 and 2 show the results of a regression of the change in log auto sales (from R.L. Polk) between the three month period between February and April 2009 relative to the three month period between September and November 2008 at the MSA level on the fraction of local borrowers with CLTV greater than 0.8. Column 1 includes no other controls with column 2 also controls for the change in unemployment between January 2007 and November 2008 and for local demographic and homeownership controls. Column 3-5 remove as a control the fraction of local borrowers with CLTV greater than 0.8. Instead, column 3 includes the change in cash-out share (dependent variable in Table 2) as an additional control. Column 4 includes the change in the refinancing propensity (dependent variable in Table 3) as an additional control. Column 5 includes both the change in the cash-out share and the change in refinance propensity as additional controls. Robust standard errors in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 house price declines during the 2007 and 2008 period.

4.1 Time-Series Variation in Empirical Patterns

The previous results show that locations with large house price declines and low resulting equity had substantially less refinancing activity and associated spending in response to QE1. Figure 4 shows that these locations also had the largest increases in unemployment, so that monetary policy stimulated spending least in the hardest hit locations. We now show that there is substantial variation across time in the distribution of house price changes, and in the next part of the paper

we argue that this leads to important changes in the effects of monetary policy across time.

Since our CRISM data begins in 2005, we cannot construct the distribution of housing equity which would be necessary to simply replicate our previous empirical analysis for earlier recessions. However, Figure 3 documents a striking correlation between house price changes and housing equity in this recession. Motivated by this relationship, we now show that the regional distribution of house price changes exhibits several patterns which change 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 1976-2013.²⁷ While the large declines in house prices during the Great Recession stand out, there is substantial variation across time. For example, during the 2001 recession, real house prices were growing rather than falling, and house price declines during the previous three recessions were mild. Next, for each year we calculate state-level house price growth and then calculate the standard deviation of this growth across states rather than the mean. Panel b) plots the evolution of the 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 the relationship between state-level unemployment and state-level house price growth; a moment which will be particularly important for determining the implications of monetary policy for inequality. In particular, we run a regression of one-year house price growth on the state-level one-year change in unemployment (in percentage points) interacted with annual dummies for each year from 1976-2013 and including year and MSA fixed effects so that results are not driven by slow moving trends at the state-level or by aggregate trends in unemployment or house price growth:

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

Since identification is purely cross-sectional, β_t provides a measure of the comovement between house prices and unemployment in the cross-section in each year. The time-series plot in Panel c) shows that in the Great Recession, unemployment and house prices were unusually correlated, with locations with large increases in unemployment also exhibiting 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 that is 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 following the 2001 recession than in the

²⁷State-level house prices are deflated using the CPI-U and 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

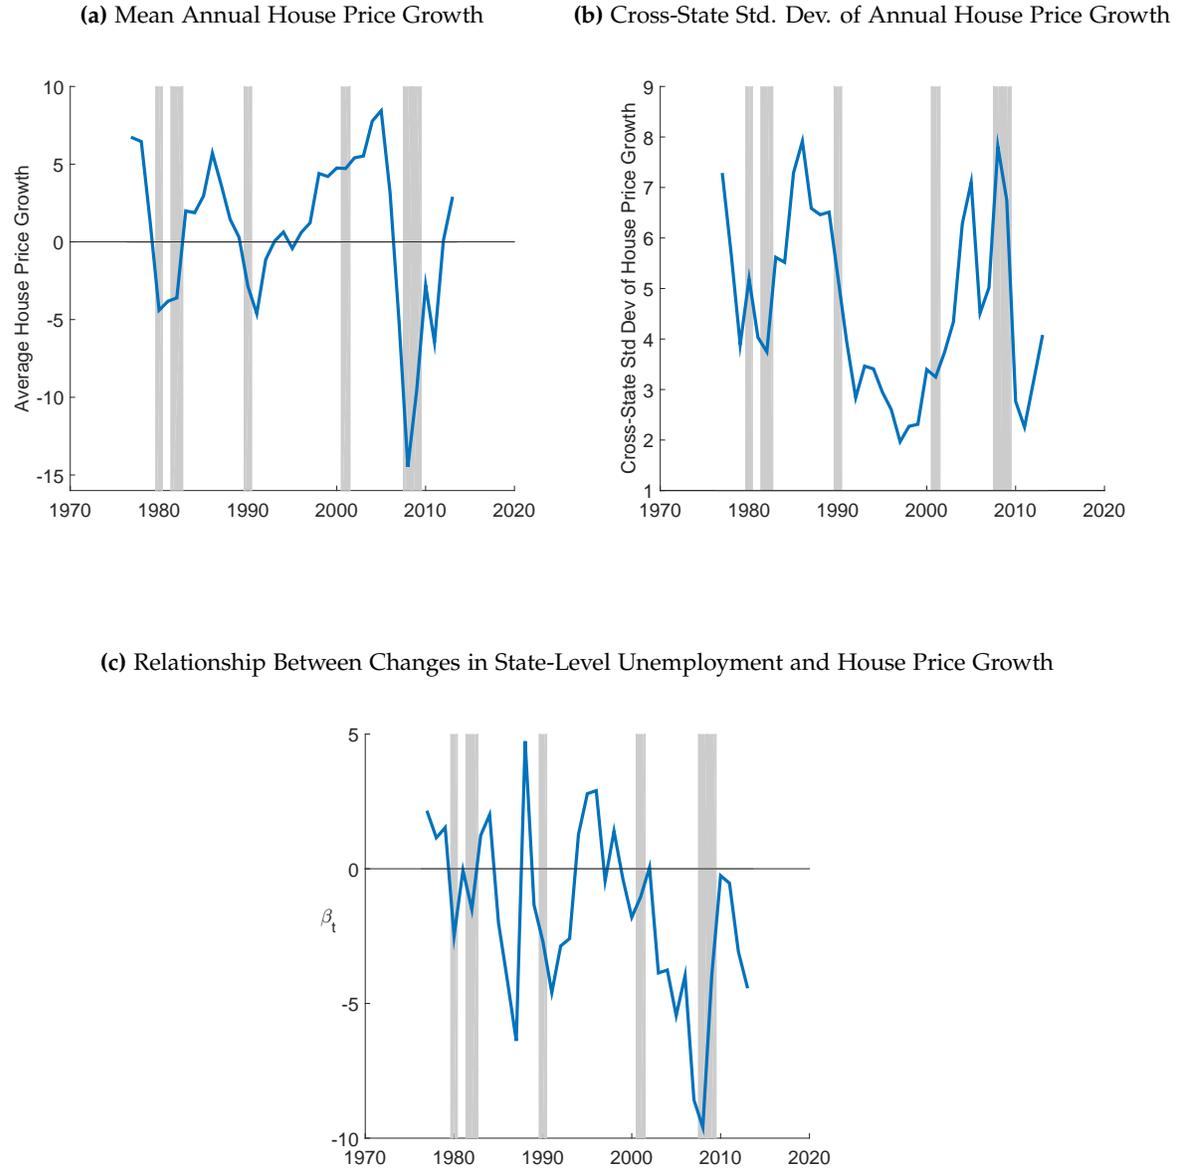
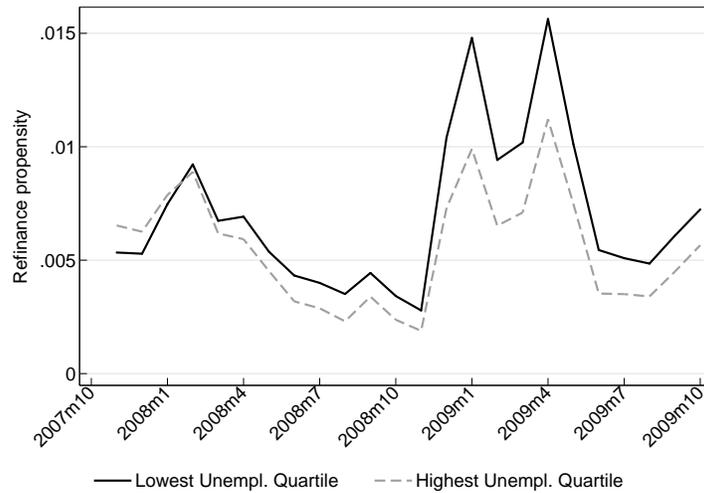


Figure shows the mean, and standard deviation of the state-level house price growth distribution across time as well as the cross-state relationship between house price growth and changes in the unemployment rate. All calculations are population weighted by state.

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

(a) Refinance Propensities 2007-2009



(b) Refinance Propensities 2001-2003

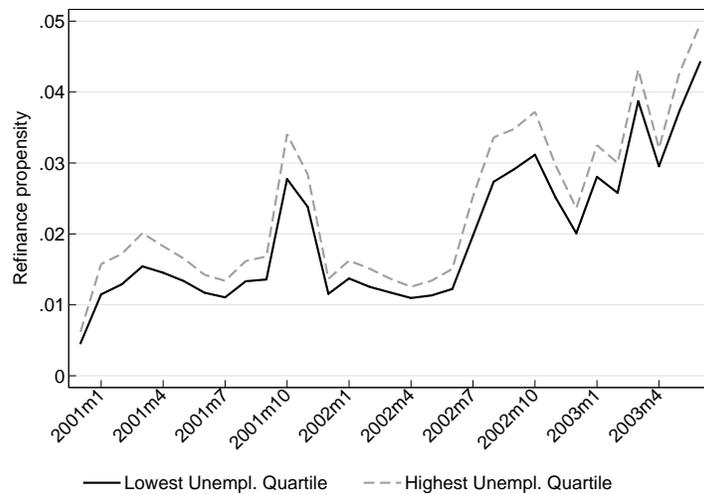


Figure shows monthly refinancing propensities for the top and bottom quartiles of MSAs with respect to their unemployment change over the period shown in the charts. Specifically, for each MSA we compute the change in unemployment between November 2007 and October 2009 (panel a) and December 2000 to 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 the 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).

2008 Recession, and that regional effects may have been very different as well. Panel (a) shows the times series of monthly refinancing propensities in HMDA for both the low and high unemployment quartile MSAs during the 2008 recession. MSAs are sorted into quartiles based on the total unemployment increase 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 period²⁸ 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 CLTV_80 quartiles in Figure 5.

Panel (b) shows that the time-series patterns for monthly refinancing by unemployment quartiles are quite different during the 2001 recession.²⁹ The sorting is again done by total unemployment increases 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.³⁰ Even more importantly, 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). There was arguably no single main monetary policy “event” during this episode—rather, the federal funds rate target decreased gradually from 6.5 percent in December 2000 to 1 percent by the end of June 2003.

5 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 general equilibrium model of household refinancing. 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 fluctuations, and an individual can borrow against the value of his house using a fixed rate mortgage which can be refinanced at any time by paying a fixed closing cost. We account for equilibrium interactions between borrowers and lenders by assuming that mortgage payments are in turn received by lenders in the economy who will potentially alter their consumption when borrowers refinance.

We focus on inequality across regions, as opposed to across households within regions, because we think the mortgage-refinancing channel of monetary policy has less relevance for within-region household heterogeneity: other factors such as income and wealth are much more

²⁸The periods chosen for calculating unemployment differences are determined not by the official NBER recession dates, but by (rough) turning points in the overall seasonally-adjusted U.S. 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.

²⁹To obtain propensities, we normalize refinancing counts by the number of outstanding mortgages in 2000 in the ACS. The CRISM dataset only goes back to 2005, so we measure refinancing propensities using HMDA data.

³⁰[Chen, Michaux, and Roussanov \(2013\)](#) document 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).

important determinants of household heterogeneity 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. In addition, our previous empirical evidence on the heterogeneous response of refinancing and consumer spending to QE1 provides important discipline for our quantitative model, and is inherently regional.

Throughout the analysis, we take the regional distribution of house prices and income as exogenous and focus on the behavior of consumption, savings and refinancing in response to interest rate changes. We look at these responses both in the aggregate as well as across regions and ask how the ex-ante regional distribution of heterogeneity affects both aggregate stimulus and ex-post regional consumption inequality in response to monetary policy.³¹

5.1 Model Description

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

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

where ε_t^{ij} is a zero-mean, i.i.d. shock over time. However, we allow for the possibility that the shock is 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 subject to a no-borrowing constraint $a_t^{ij} \geq 0$. They are endowed with one housing unit with price q_t^j , which can be used as collateral for issuing mortgage debt. House prices follow a random-walk with a drift,

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

where v_t^j is a zero-mean, shock that is i.i.d. over time but potentially correlated across regions.

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 will show in [Claim 1](#). The second

³¹Treating the regional distribution of income as exogenous but the regional distribution of consumption as endogenous amounts to an assumption that all consumption is fully tradeable. This dramatically simplifies the model solution, but 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 relatively high initial income, inequality will be further amplified if income then rises by more due to the greater consumption response. Similarly, endogenizing house prices would also amplify our conclusions as reductions in interest rates amplify regional house price inequality.

is because it simplifies aggregation: we collect households into regions and then aggregate them in a straightforward manner because aggregate, regional and idiosyncratic household shocks enter symmetrically in the problem, e.g., 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 point of view of the household.

Since our empirical evidence focuses on the behavior of refinancing behavior of non-moving households, we assume for simplicity that agents cannot buy or sell houses and that mortgage debt is infinite maturity. Mortgage debt requires a constant mortgage payment $r_{\tau_0}^m m_{\tau_0}^{ij}$ every period, which is determined at the moment of debt issuance τ_0 . However, we allow households to refinance their mortgage at any time $\tau > \tau_0$ if they pay a fixed monetary cost $F_{\tau}^{ij} q_{\tau}^j$.³² We assume that F_{τ}^{ij} is an iid stochastic process which is uncorrelated with house prices or income in order to introduce additional heterogeneity in refinancing decisions across households.

When refinancing, households lock in the current interest rate r_{τ}^m and future mortgage payments $r_{\tau}^m m_{\tau}^{ij}$. Furthermore, we assume that when refinancing, households always borrow up to the maximum amount allowed by the loan-to-value requirement γ . This substantially simplifies the household decision problem and increases tractability in our benchmark model, and we show in Section 5.4.2 that relaxing this assumption has little effect on our conclusions. This implies that the new mortgage balance is $m_{\tau}^{ij} = \gamma q_{\tau}^j$ and that cashed-out equity is $m_{\tau}^{ij} - m_{\tau_0}^{ij}$. This collateralized borrowing contract captures the primary features of fixed rate mortgages with a refinancing option which are common in the U.S. economy. Finally, we assume that interest rates $\{r_t, r_t^m\}$ 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 $\{r_0^m, \gamma q_0\}$, facing current house price q_t , and current

³²We assume the fixed cost is proportional to current house prices in the region 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, which allows us to eliminate house prices as a state-variable, as shown below.

mortgage rate r_t^m and interest rate r_t can be written recursively as,

$$\begin{aligned}
V(a, y, q_0, r_0^m, q, r^m, r, F) &= \max\{V^{noref i}(a, y, q_0, r_0^m, q, r^m, r, F), V^{ref i}(a, y, q_0, r_0^m, q, r^m, r, F)\} \\
V^{noref i}(a, y, q_0, r_0^m, q, r^m, r, F) &= \max_{\{a', c\}} u(c) + \beta \mathbb{E}[V(a', y', q_0, r_0^m, q', r^{m'}, r', F')] \\
&\text{s.t. } c + a' \leq a(1 + r) + y - \gamma r_0^m q_0 \\
&\quad a' \geq 0, c \geq 0 \\
&\quad \log(q') = \mu_q + \log(q) + v \\
&\quad \log(y') = \mu_y + \log(y) + \varepsilon \\
V^{ref i}(a, y, q_0, r_0^m, q, r^m, r, F) &= \max_{\{a', c\}} u(c) + \beta \mathbb{E}[V(a', y', q, r^m, q', r^{m'}, r', F')] \\
&\text{s.t. } c + a' \leq a(1 + r) + y - \gamma r_0^m q_0 + \gamma(q - q_0) - Fq \\
&\quad a' \geq 0, c \geq 0 \\
&\quad \log(q') = \mu_q + \log(q) + v \\
&\quad \log(y') = \mu_y + \log(y) + \varepsilon
\end{aligned}$$

By inspecting the above recursive formulation, we can readily see household's incentives to refinance. When the current interest rate r^m is below the interest rate at which the household last refinanced r_0^m , the household can secure a lower mortgage payment forever, even if house prices are unchanged. When the current house price q is above the price at which the household last refinanced q_0 , the household can refinance to cash-out accumulated equity even if rates have not changed, but then face increased mortgage payments in the future. In order to make progress in characterizing policy functions, it is convenient to eliminate a state variable. We do so by showing that the value function is homogenous in house prices:

Claim 1 For fixed $\{r, r^m, r^0\}$, the value function is homogenous 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

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

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,

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

■

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 last refinancing, households find it optimal to refinance

and cash-out because equity has become large enough³³. 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, one of which, assets a_t , evolves endogenously even when not refinancing. However, when solving the model numerically, we find that the refinancing decision is indeed a threshold level of equity, where this threshold 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 U.S. 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 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 —e.g., when borrowers' refinance after a decline in interest rates, the dividends accruing to lenders decrease, which should reduce their consumption.

For simplicity, we assume the representative lender is a permanent income consumer and receives certain dividend payments from the mutual fund d (a consequence of a law of large numbers for households). Given a law of motion for dividends $d' = \tilde{G}(\cdot)$,³⁴ their value function is

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

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 dividends $d' = G(S, r^m, r)$, 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 $[\tilde{a}', \tilde{c}, \tilde{x}, \mathbb{I}^{refi}] (\tilde{a}, \tilde{y}, \tilde{x}, r_0^m, 1, r^m, r, F)$ and $[a'_R, c_R] (a_R, d)$ such that

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

³³Alternatively, a sequence of positive house price shocks could also induce refinancing even with no drift

³⁴The law of motion will be determined in equilibrium as a function of the refinancing decisions of all households. Since we do not endogenize interest rates, this is straightforward in practice as we can calculate the present value of the dividend impulse response to a decline in interest rates entirely from the lender side and then separately calculate lender consumption responses.

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 $\tilde{G}(\cdot)$ 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 definition of equilibrium 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 that the mortgage market clears by introducing the representative lender. One can interpret this portion of the model as partial equilibrium or as corresponding to 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 in the international market. 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.

5.2 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 $rm = 0.06$ and then consider the response to a one-time, unanticipated permanent decline to $rm = 0.05$. We begin with this specification since 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.³⁵ The initial value of 0.06 and decline to 0.05 roughly matches 30-year fixed rates before and after quantitative easing. 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.

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 shocks is common to all households in a region. In our baseline results, we assume that the standard deviation of this common region component is 0.025, to match the differences across regions in our empirical work.³⁶ The drift in income and house prices is assumed to be

³⁵It is also consistent with the extremely high persistence of the response of mortgage rates to QE1.

³⁶Note 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 in simulation for calculating measures of regional inequality and has no effect on any aggregate impulse responses.

identical so that real house prices are constant on average.³⁷

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 responses to refinancing across regions in response to Quantitative Easing. Our primary counterfactual analysis then 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.

More specifically, we initialize the model from the stochastic steady-state, but in period t , we assume that the economy is hit with an aggregate decline of house prices of 12.5%, as observed in 2008. Households also get hit with an additional regional house price shock which can take the value -12.5%, 0%, +12.5% so that in 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 decline of 25% is picked to match the empirical decline in house prices for the highest CLTV_80 quartile in our empirical analysis while the 0% change matches that in the lowest CLTV_80 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 our 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 in our economy calibrated to match this period 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 iid fixed cost each period which 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

³⁷Put differently, this implies that in the original problem both house prices and income remain relevant in the long-run. If $\mu_y > \mu_p$, then in the long-run, housing becomes vanishing share of household budgets and the refinancing decision is irrelevant for decisions and if $\mu_y < \mu_p$ then in the long-run housing is the only thing that matters in the budget and income becomes irrelevant.

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 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. We focus our calibration on the 2008 period rather than the stochastic steady-state for several reasons. First, this is the period of our event study and we have strong empirical evidence both on the distribution of house price changes, income and refinancing responses during this period. Second, it is more difficult to construct an empirical counterpart to the stochastic steady-state. 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 is only available 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 off of recession periods, so 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 a diff-in-diff 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.³⁸

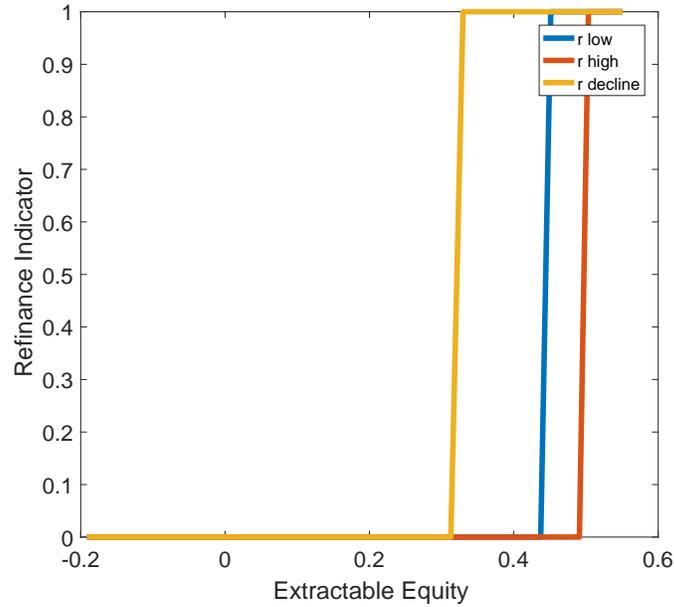
5.3 Model Results

Figure 10 shows the threshold-like property of refinancing policies in a stationary environment with high and 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. 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.

Figure 11 shows how the refinancing threshold changes with assets and earnings. The threshold typically increases with both assets and earnings. Intuitively, the marginal utility of an extra unit of income coming from equity extraction when refinancing is higher when current assets or earnings are lower, making households want to refinance more often at lower equity levels. To

³⁸Nevertheless, alternative strategies for calibrating the fixed cost process produce broadly similar results.

Figure 10: Refinance Decision Follows a Threshold Policy

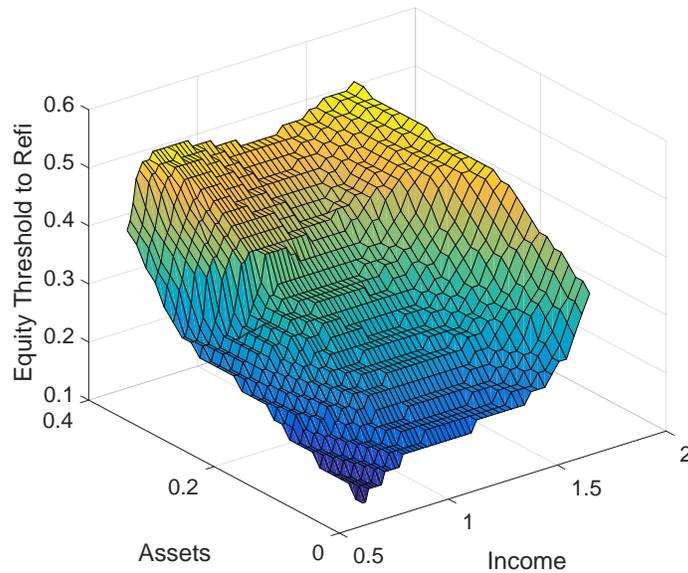


see this, consider the case of a household with no assets and very low income that is borrowing constrained. This household behaves, essentially, as if it was hand-to-mouth. Because an extra unit of income is rather valuable today, the household refinances whenever extractable equity is above the fixed cost. In the other extreme, a household with high assets and income has a larger option value of waiting until equity increases compared to the value of an extra unit of income today.

Figure 12 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 5, since the model is calibrated to hit these numbers. Just as in the data, regions with relatively high equity are much more likely to refinance in response to the decline in rates. The bottom panel shows the response of consumption by region. Unsurprisingly, the high equity regions also increase consumption more. While we do not have empirical measures of broad consumption at the regional level, this is consistent with the empirical patterns for auto spending by region in Figure 7.

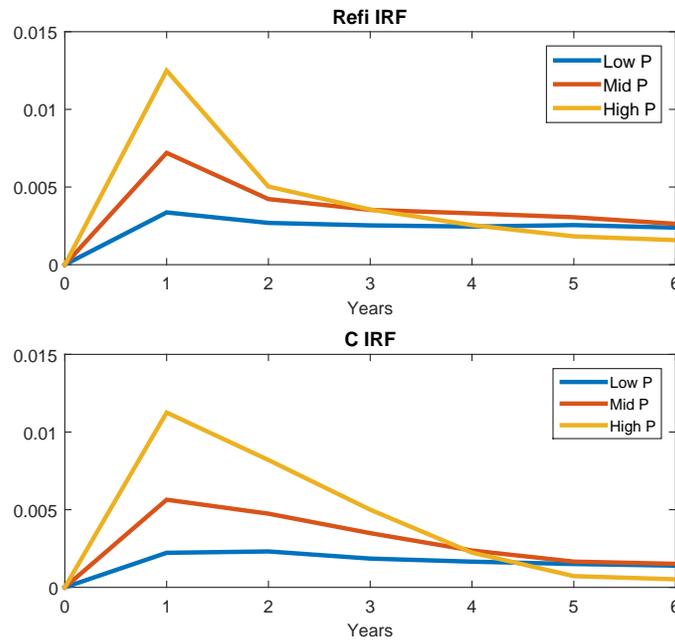
We next compute aggregate consumption in the economy as well as the variance of log consumption across regions. Figure 13 then shows the response of aggregate consumption and regional consumption variance to the decline in mortgage rates. This shows our first important model result: in our baseline economy, a reduction in interest rates increases aggregate output but it also increases the variance of consumption across regions. This increase in regional variance occurs because consumption increases the most in regions with high house prices. Since house prices and income are correlated in this economy, these regions already have the highest

Figure 11: Relationship between refinancing threshold, equity and income



Refinancing Threshold For Each Asset and Income Value for Low Fixed Cost

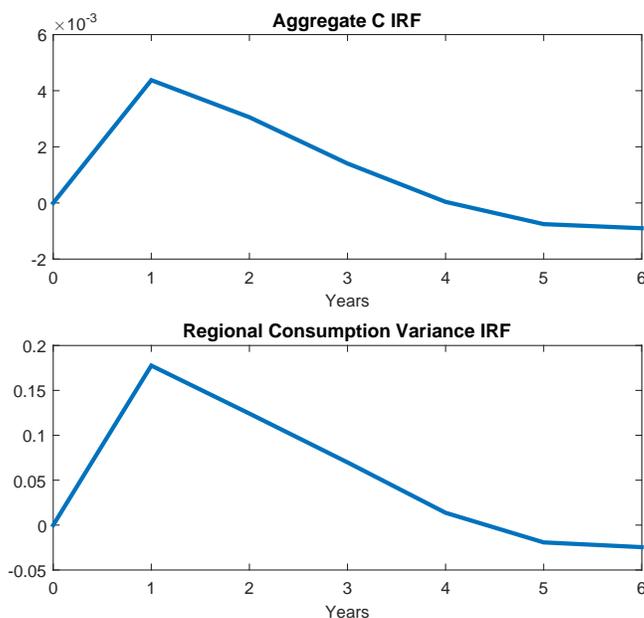
Figure 12: Response to interest decline by region



For the baseline 2008m11 calibration, consumption IRF shows the change in log consumption in response to a one percentage point reduction in mortgage rates. Refi IRF shows the change in the monthly fraction of households refinancing in response to the same one percentage point reduction in mortgage rates.

income and consumption before the reduction in interest rates. Thus, while monetary policy increases overall consumption it mainly does so by stimulating consumption in the locations that are already doing relatively well so that there is an important tradeoff between stimulus and inequality.

Figure 13: Aggregate Stimulus Vs. Regional Inequality



For the baseline 2008m11 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.

This tradeoff 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 diff-in-diff 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 QE, it is possible that high equity regions initially had consumption that was temporarily low, so that QE would actually decrease inequality. In our quantitative model, this is not the case: high equity regions have substantially higher consumption so that a decline in mortgage rates dramatically amplifies inequality. 2) The diff-in-diff empirical analysis measures only differential effects of the policy across regions and cannot measure any aggregate effects. [Greenwald \(2016\)](#) argues that these general equilibrium effects are potentially quite important: while declines in interest rates increase income and consumption for borrowers, one must consider the effects on lenders. When borrowers' payments fall, this reduces the income of lenders so their consumption should fall. In some scenarios this offset can

be one-for-one so that lenders' consumption declines by exactly the amount that borrowers' consumption rises so that the redistributive channel arising from refinancing leads to no aggregate effects. If that were the case, monetary policy would not face a tradeoff between stimulus and inequality, it would just increase inequality with no aggregate stimulus benefit. Our quantitative model explicitly accounts for these lender effects and shows that in our quantitative environment, they do not cancel out. Aggregate consumption rises at least modestly when mortgage rates fall.

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 an increase in aggregate consumption but also an increase in inequality across regions. That our model reproduces the behavior of the economy just before and after QE also gives us some confidence in using our model for more ambitious counterfactual analysis.

We now do just that, moving beyond the analysis of collateral and refinancing in this particular episode to a more general study of the role of collateral conditions in the monetary transmission mechanism. In particular, we show that the theoretical effects of monetary policy depend crucially on the distribution of collateral in the economy. This means that accounting for time-variation in this distribution is crucial for correctly predicting the consequences of monetary policy. As the level of collateral, its variance and its correlation with income moves across time, the effects of interest rate changes on both aggregate economic activity and regional inequality also change. This counterfactual analysis is not of purely theoretical interest: in the empirical analysis, we showed that such time-variation is clearly present in the data.

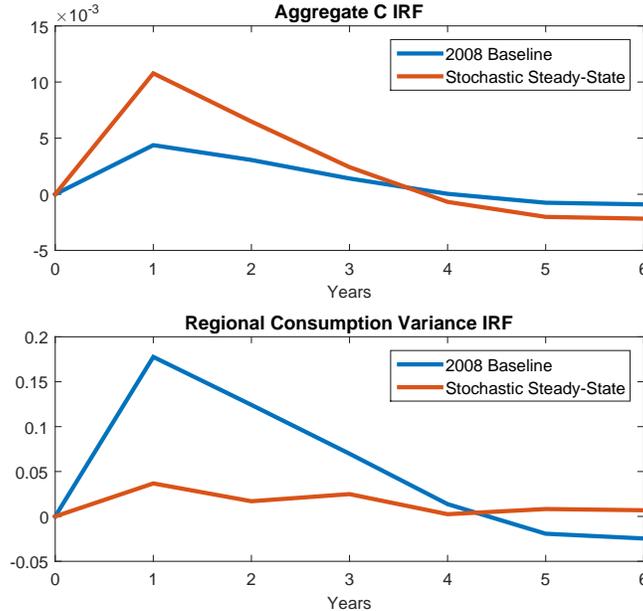
Figure 14 compares the impulse response function of consumption and regional inequality to an interest rate shock in the 2008 baseline to that which would have instead occurred in the stochastic steady-state. In the stochastic steady-state aggregate house prices do not decline, there is no increase in the variance of house prices, and there is no correlation between house prices and income.³⁹ Clearly changing the distribution of equity dramatically changes the consequences of monetary policy: in the stochastic steady-state, the same decline in interest rates raises aggregate consumption by more than twice as much as in 2008. In addition, there is a negligible effect on inequality across regions so that there is essentially no tradeoff between stimulus and inequality.

To understand why the distribution of collateral in 2008 makes consumption less responsive to monetary policy than in the stochastic steady-state, it is useful to first explain why interest rate reductions lead to an increase in consumption. Figure 15 shows the distribution of equity in 2008 and the fraction of households adjusting for a given level of equity when the interest rate is permanently at 0.06, permanently at 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.⁴⁰ Consistent with the threshold policy described above, households are much more likely to refinance when equity is large so that the adjustment

³⁹In the baseline we calibrate the house price shocks across regions to match that of the CLTV groups in the data. In the stochastic steady-state there is not such a natural benchmark for defining this difference, so we choose a house price shocks of $\pm 2SD$ as the difference across regions.

⁴⁰We assume that households have required equity of 20% but can extract equity over 20%.

Figure 14: 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. The baseline is calibrated to match economic conditions in 2008m11 including a 12.5% aggregate house price decline, house price differences between the high and low region of 30% that are correlated with income. Normal times correspond to the stochastic steady-state with no house price decline and house prices that are uncorrelated with income. The low house price region has -2SD price shocks and the high house price region has +2SD price shock.

probability increases with the level of extractable equity. Since $r^m > r$, households who refinance always consume some fraction of their equity in the period of refinancing so that the MPC out of equity is always greater than zero, and typically it is large.⁴¹ This means that consumption increases when there is an increase in the fraction of households refinancing.⁴²

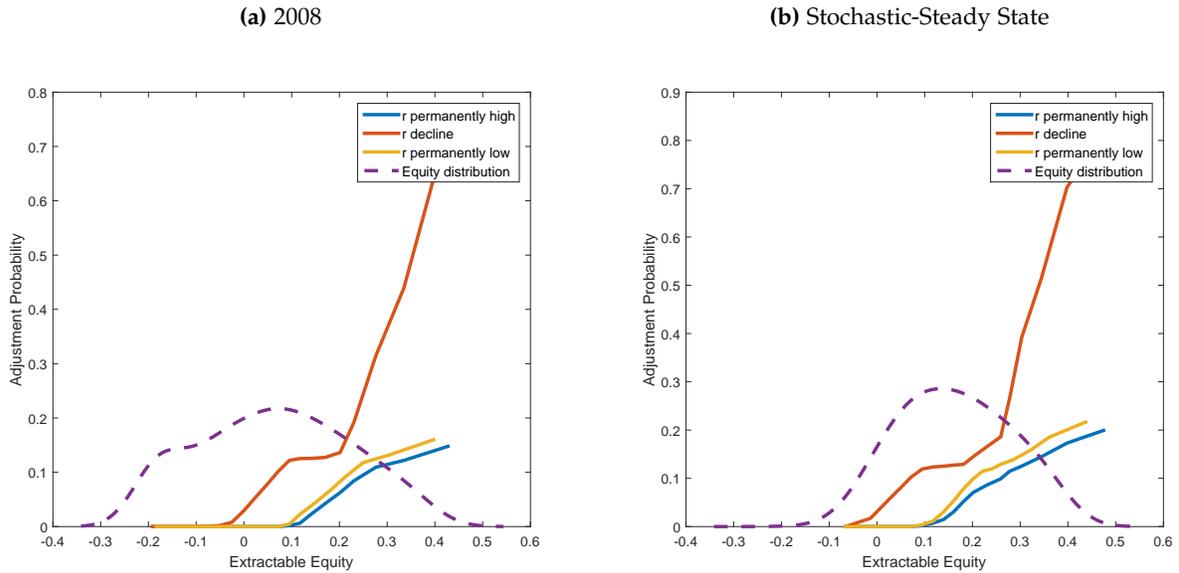
When interest rates decline, households lower their refinancing threshold relative to the case when r^m is permanently low. (The blue line shifts left beyond the yellow line). This is because some households who would not choose to extract equity today if their interest rate was already low choose to 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 refinancing, cash-out and an increase in consumption. This extra refinancing arises from the mass of households between the blue line and the orange line since households to the right of the blue line will typically refinance even if interest rates do not decline and households to the left of the orange line will not refinance even if interest rates do decline.⁴³ It is not the fraction of households who want to refinance and extract

⁴¹On average this MPC is 0.5 so that half of extracted equity is consumed on impact. Unsurprisingly, this number is higher for households with little liquid assets and smaller for those with larger liquid assets.

⁴²As we show later, since the effects of refinancing on borrowers' consumption largely work through the cash-out channel and not through the payment reduction channel, equilibrium effects on the lender side are relatively unimportant.

⁴³The change in the fraction of households refinancing with equity x is given by the vertical difference between the

Figure 15: Distribution of Equity and Adjustment Decisions: 2008 calibration vs stochastic steady-state



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

equity that matters for the strength of the interest rate impulse response, it is the change in this fraction that is the relevant object for assessing the consequences of interest rate declines.⁴⁴

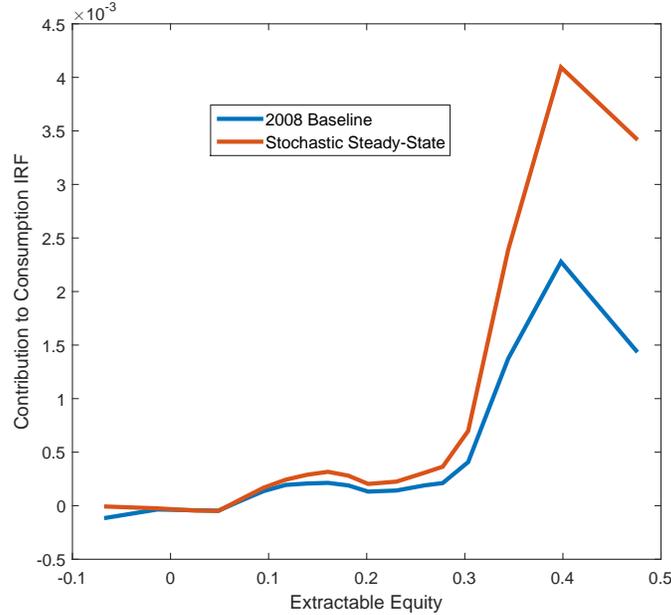
Why does moving from the stochastic steady-state to the distribution of collateral in 2008 reduce aggregate consumption responses? The decline in equity and increase in variance in 2008 shifts the distribution of equity (shown in purple) to the left and fans it out. This means that there are many fewer households in the region between the blue and orange lines, whose decisions to refinance are altered by the decline in the interest rates. This implies that the refinancing response to a decline in interest rates is substantially muted, which in turn reduces the response of consumption.

This can be seen even more precisely by looking at changes in consumption rather than changes in refinancing when rates decline. In order to decompose the consumption response to a decline in interest rates into responses by households of various equity levels, it is useful to compute the contribution of each equity level to the overall consumption impulse response. In order to do this, we first calculate the average level of consumption for each level of extractable equity x , with and without the interest rate decline. The difference between these

orange and blue lines, so the total change in refinancing is given by the difference between these lines integrated over the distribution of equity.

⁴⁴That is, only those households whose decision to refinance are 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 16: $\Delta C(x)$ (Contribution of each equity level to aggregate consumption response): 2008 vs. Stochastic Steady-State



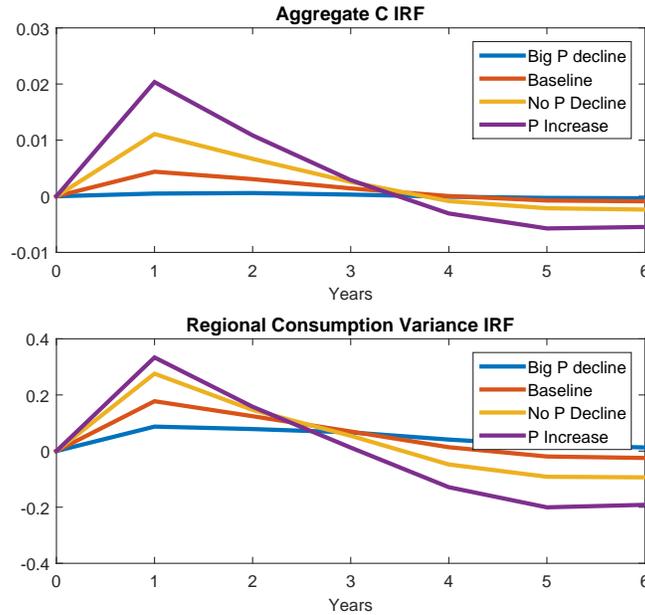
This figure shows the average contribution of each level of x to the aggregate consumption responses on impact shown in Figure 14.

values gives the average consumption response to a decline in interest rates for households with equity x . Finally, we multiply this by the fraction of households with equity x in order to measure the contribution to aggregate consumption responses. More precisely for each x , we compute $\Delta C_t(x) = \int_{a \in A, y \in Y, \chi = x} (c(a, y, \chi, r_{high}) - c(a, y, \chi, r_{decline})) f_t(a, y, \chi) da dy$, where f_t is the joint-density of assets, income and equity at a particular date. Figure 16 plots $\Delta C(x)$ for the distribution of assets, income and equity in 2008, f_{2008} , and in the stochastic-steady state, f_{stoch_ss} . Note that adding $\Delta C_{2008}(x)$ over all x values delivers the impulse response on impact in the 2008 baseline shown in Figure 14, and adding $\Delta C_{stoch_ss}(x)$ over all x values gives the impulse response on impact in the stochastic steady-state.

Clearly, high equity values contribute disproportionately to the aggregate impulse response. This is because households with high equity values have refinancing decisions which are most responsive to declines in interest rates (the orange and blue lines are most different for high equity values in Figure 15), and they have the largest consumption responses conditional on refinancing. Moving from the distribution of equity in the stochastic steady-state to the distribution in 2008 reduces the aggregate consumption impulse response by reducing the mass of households with large equity, whose consumption is most responsive to declines in interest rates.

The distribution of collateral 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 variance of equity and the correlation between income and equity for aggregate responses? Furthermore, why does a

Figure 17: Effects of changing average equity



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. The large decline has a 25% decline in house prices and the house price increase economy features a 12.5% house price increase. The variance of equity and its correlation with income is fixed at the 2008 calibration across all simulations.

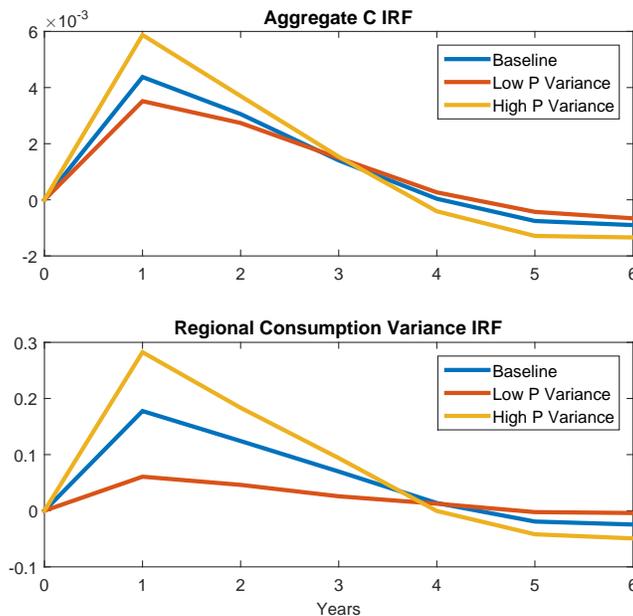
change in interest rates in 2008 increase inequality much more than in the stochastic steady-state? We now go through each component of the distribution of equity and discuss its role in shaping aggregate consumption and inequality responses to interest rate declines.

Figure 17 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 conditional on refinancing.⁴⁵ 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, as shown in Figure 16. 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

⁴⁵ As is typical in fixed cost type models, these responses are non-linear. As the initial equity level becomes extremely large, the consumption response to an interest rate decline begins to fall towards zero. This is because if equity is extremely large, then households all refinance even without the decline in interest rates. However, these effects only become relevant when house prices increase by more than 30-40% in a single year.

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.

Figure 18: Effects of changing equity variance

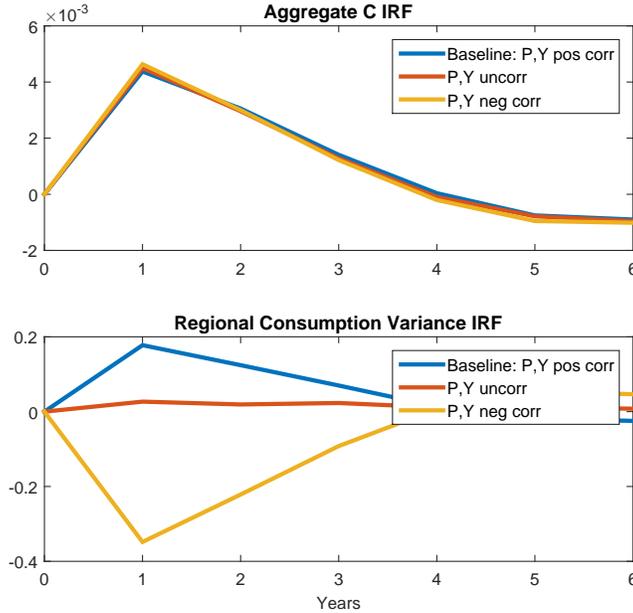


Impulse response functions show the change in log aggregate consumption and in log consumption variance across regions in response to a one percentage point reduction in mortgage rates. The high variance calibration doubles the difference between high and low house price regions while the low variance calibration halves it. All economies feature the same aggregate decline in house prices and correlation with income.

Figure 18 shows the effects of changing the variance of equity across regions. An increase in the variance of equity increases the aggregate consumption response to monetary policy but also worsens the consequences for monetary policy. The intuition is almost identical to the effects of mean shifts and again follows immediately from Figures 15 and 16. 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, Figure 19 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 inequality and when income and equity are negatively correlated, then declines in interest rates substantially reduce inequality. In contrast, the correlation between income and equity has almost no effect on the aggregate impulse response. 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

Figure 19: Effects of changing equity correlation



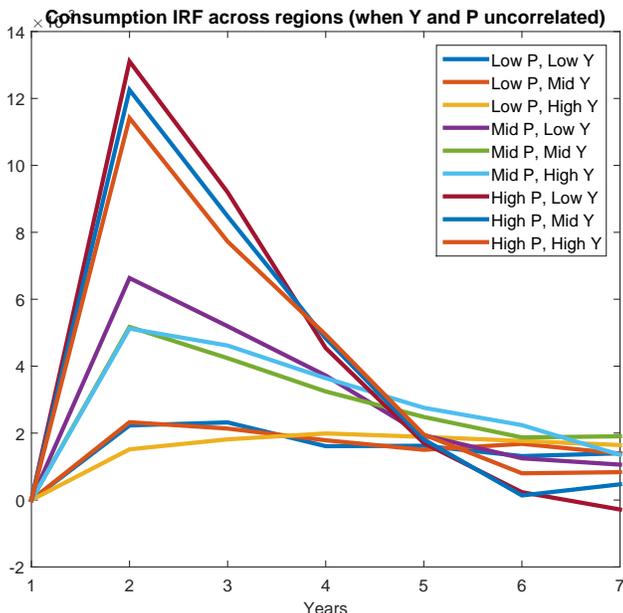
Impulse response functions show the change in log aggregate consumption and in log consumption variance across regions in response to a one percentage point reduction in mortgage rates. In the baseline calibration, income and house prices are positively correlated across regions. In the other two calibrations they are uncorrelated or negatively correlated. All simulations feature the same baseline decline in house prices and variance across regions.

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. Figure 20 shows how consumption growth responses to interest rates vary with equity and income for the 3x3 regions in the stochastic steady-state. Clearly, consumption growth is largest for high house price regions, but in regions with positive equity, consumption growth is higher for those with low income than those with high income. This is because low income regions are more constrained and so respond more strongly to any shocks. This would suggest that when house prices and income are positively correlated, aggregate consumption responses should be reduced since there are no longer any high house price, low income regions which exhibit the greatest consumption growth when interest rates fall. However, the greater response of low income, high house price regions than high income, high house price regions reflects differences in consumption growth not in absolute changes. As noted in the above discussion, low income regions have lower initial consumption levels than high income regions. This means that even though the high income, high house price region has a lower growth rate, it has almost an identical absolute change to that of the low income, high house price region. Since aggregate consumption responses are determined by the average change in levels not the average change

in growth rates, this means that increasing the correlation between income and house prices has a negligible effect on aggregate consumption.

Figure 20: Understanding Correlation Effects: Responses by Region



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 across 9 regions defined by the combination of high, middle and low equity and income. This economy features the baseline decline in house prices and variance across regions but house prices and income are uncorrelated in order to illustrate all possible combinations.

5.4 Robustness

5.4.1 Spatial Variation in Adjustable-Rate Mortgage Share

Our baseline analysis assumes that all mortgages have fixed rates. In reality, a substantial fraction of mortgages have adjustable rates which reset as current mortgage rates decline even if households do not refinance. ARM rates are based on short-term rates (most commonly 6-month or 1-year LIBOR, or 1-year Treasuries), which also fell in late 2008 / early 2009 as the Fed lowered its interest rate target. 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.⁴⁶ 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 means that under ARMs, cash-out based spending is likely to respond less to monetary policy.

⁴⁶See e.g. Auclert (2015).

Figure 21: Adjustable-Rate Mortgage Shares pre-QE1 and Payment Reductions over Nov 2008 - June 2009

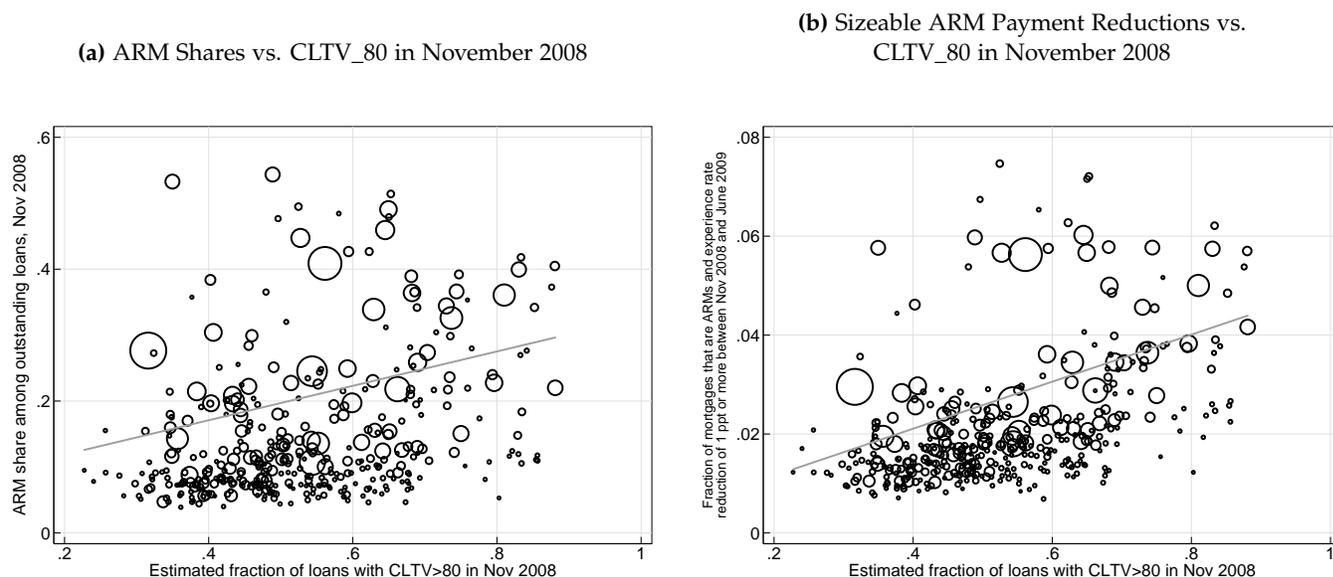


Figure shows scatter plots of the balance-weighted fraction of adjustable-rate mortgages (ARMs) in a given MSA (panel a) or the balance-weighted fraction that are ARMs and experience a rate reduction of 1 percentage point or more over November 2008 – June 2009, as measured in the CRISM data, versus the fraction of loans within the MSA with CLTV greater than 0.8 (CLTV_80) in November 2008. The size of the circle represents the 2008 population of the MSA. The grey lines represent simple regression lines fitting the scatter plot (population weighted).

Accounting for ARMs is also potentially important for the patterns of regional inequality, since ARM shares vary spatially. In particular, the share of ARMs among outstanding loans in November 2008 was higher in MSAs with higher CLTVs, as shown in panel (a) of Figure 21.⁴⁷ This is due to the fact that the ARM share increased more strongly during the boom years in areas with larger price increases (for affordability reasons and potentially also due to speculation—see Barlevy and Fisher 2010), which subsequently experienced the largest bust. This larger share of ARMs in high CLTV regions could potentially mean that a greater share of all mortgages experience rate declines in those regions despite our previous evidence that high CLTV fixed rate mortgages are much less likely to refinance.

However, the extent to which ARM resets offset the heterogeneity in refinancing propensities that we previously documented depends on what fraction of ARM borrowers actually received a payment-reducing rate reset between November 2008 and mid-2009 (our main period of interest). There are a number of reasons why not all ARM borrowers benefited from low short-term rates at that point: (i) in the US, most ARMs are “hybrids” with initial fixed-rate periods of 3, 5, 7 or 10 years. Borrowers with an ongoing fixed-rate period had to refinance to lower their payment, just like FRM borrowers. (ii) The length of the fixed-rate period of hybrid ARMs often coincided

⁴⁷Calza, Monacelli, and Stracca (2013) also show that ARMs are more prevalent in Spain and Italy than in Germany and France.

with the length of an interest-only (IO) period on the loan (during which the borrower only pays interest but does not amortize principal). When the IO period ends, the required payment jumps up, since the borrower now makes payments towards the principal. This could more than offset a decrease in the rate that happens at the same time. (iii) ARMs commonly have “rate floors” which indicate a minimum rate for the loan. For subprime loans, this minimum rate was almost always set at the initial interest rate of the loan (Bhardwaj and Sengupta, 2012), so that borrowers would not benefit from declines in the index rate that would otherwise lead to further decreases.

In Panel (b) of Figure 21 we plot the fraction of loans that are ARMs and experience an economically significant rate reduction of 1 percentage point or more over the November 2008 – June 2009 period against CLTV₈₀.⁴⁸ Since overall, only about 15% of ARMs experienced significant downward resets over our main period of interest, the level is much lower than in panel (a). Differences in ARM resets across MSAs are about half as large as the difference in refi propensities in early 2009. The regression line in the third figure has a slope of 0.047 ($t = 5.3$). Meanwhile, regressing the 6-month refinancing propensity (from January-June 2009) on CLTV₈₀ yields a coefficient of -0.10. Thus, overall, decreases in interest rates were still larger in MSAs with low CLTV₈₀.

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 other households 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 locations to match the variation in Panel (a) of Figure 21. In light of the above discussion and the evidence in panel (b), 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 the behavior of ARMs can change our conclusions. Figure 22 shows that even under this conservative calibration, modeling ARMs has a negligible effect on our conclusions.

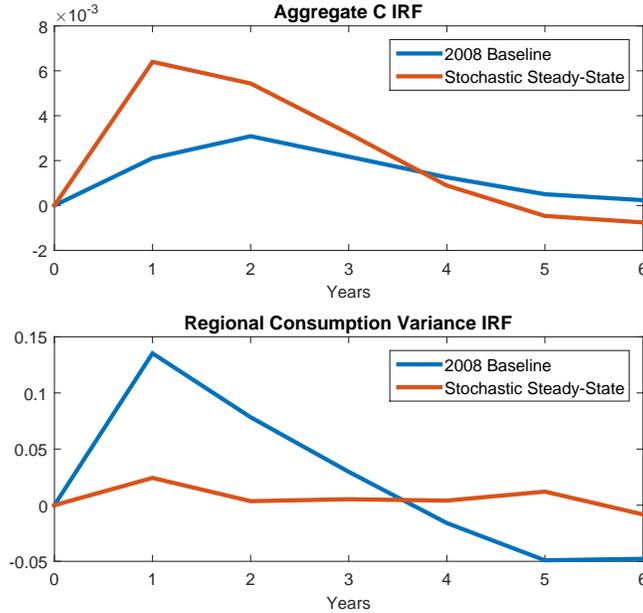
This is because even in the low equity regions, FRM shares remain large, and overall spending responses to interest rate declines are dominated by cash-out effects and not by reductions in interest payments. Low house price regions have more rate resets due to ARMs but less rate resets amongst their FRMs. On net, the FRM effect still dominates so that there are more rate reductions in high equity regions. More importantly, it remains the case that in low equity regions there is no cash-out activity since there is no equity to remove, while in the high equity regions there is a significant cash-out and spending response to interest rate declines.

5.4.2 Additional Model Robustness

In our baseline model, we abstract from the distinction between cash-out and non-cash out refinancing by assuming that 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

⁴⁸We furthermore require that the recorded required monthly payment does not increase over the same period (which may indicate that the loan’s IO period expired).

Figure 22: Robustness to matching regional ARM shares



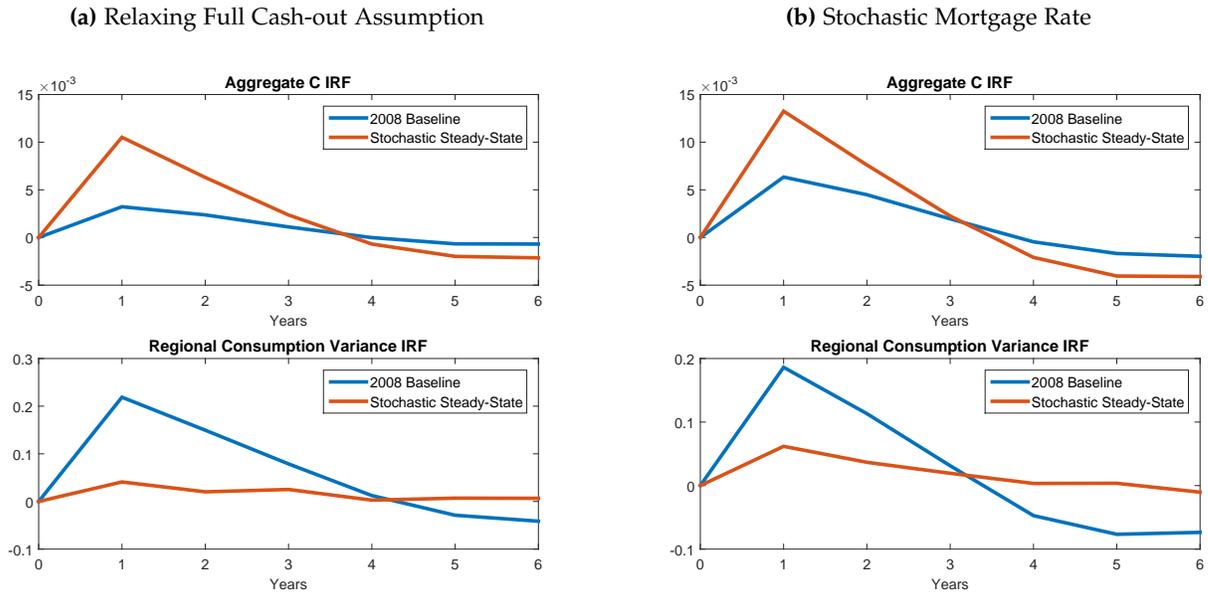
This figure repeats Figure 14 in a model with both FRM and ARMs calibrated to match the regional differences with house prices in the data.

conclusions. In particular, Panel a) of Figure 23 shows that results are extremely similar for a version of the 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 largely choose to extract their equity when refinancing anyway, and households in low equity locations have little equity to extract on average, so that that the distinction between a cash-out and a rate refi is less relevant.⁴⁹

In our baseline model, we assume that mortgage rates are constant across time and explore the responses to a one-time unanticipated decline 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 possible way and increases the computational tractability of the model, which allows some of the above robustness exercises. However, we have explored the robustness of our conclusions under alternative assumptions for the stochastic process on r^m . In particular, Panel b) of Figure 23 shows results in a version of the model where r^m follows an AR process, with persistence of 0.89 and standard deviation of 0.0055, picked to match the behavior of 30-year mortgage rates from 1990-2015. Since this has little effect on our

⁴⁹Overall, in this version of the model, the cash-out share of refinancing when mortgage rates decline is 65% as compared to roughly 50% in the data in the months immediately following 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 associated capital gains when loans are prepaid due to moving as a form of equity extraction would raise the empirical equity share substantially.

Figure 23: Robustness Results



This figure repeats Figure 14 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

conclusions, we use the simpler one-time shock as our baseline experiment.

Finally, our baseline model assumes that 50% of mortgage payments are ultimately made to non-US consumers.⁵⁰ However, large fractions of mortgage debt are held by commercial banks, who are owned in part by foreign shareholders. This makes it difficult to precisely measure the ultimate recipient of mortgage payments. In Figure 24, we show the aggregate response of the economy under alternative extreme assumptions about who holds mortgage debt. This also allows us to assess the importance of general equilibrium for our results. As the foreign share of lenders declines, the importance of GE 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.⁵¹

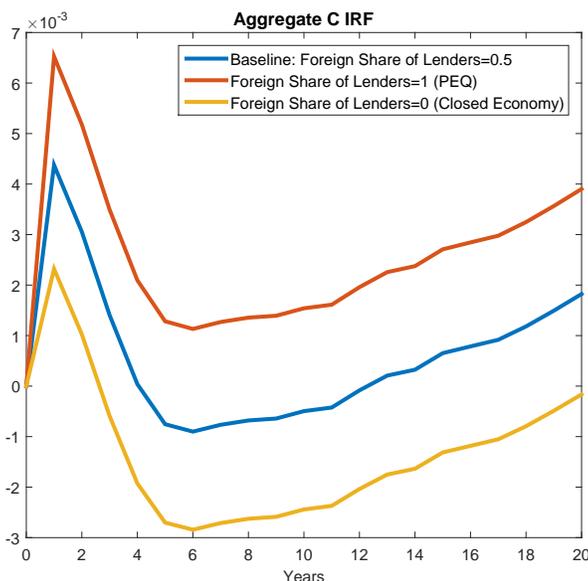
6 Conclusions

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

⁵⁰See e.g. libertystreeteconomics.newyorkfed.org/2012/01/why-mortgage-refinancing-is-not-a-zero-sum-game.html for the distribution of agency MBS ownership.

⁵¹The behavior of lenders is irrelevant for cross-region inequality since we assume that lenders are equally distributed geographically.

Figure 24: Importance of General Equilibrium Effects from Lenders' Consumption



This figure 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.

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. Non-linear interactions between collateral constraints, refinancing and spending mean that the *distribution* of housing equity plays a crucial role in the response of the economy 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 lead to a substantial dampening of this collateralized lending 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 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, monetary policy is much more powerful and can potentially mitigate inequality.

We provide evidence of these collateral effects using novel household level-data which includes comprehensive information on mortgage debt and refinancing. We show that after QE1, there was an aggregate increase in refinancing but that there was very 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 ac-

tivity during the 2001 easing cycle than during the Great Recession, and refinancing was actually more common in regions with high unemployment. Thus, the data confirms that time-varying heterogeneity in the distribution of collateral is important for understanding the consequences of monetary policy across time.

Our data comes from the U.S. 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 U.S. 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 U.S. and Europe and across countries within Europe, Section 5.4.1 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 trade-offs 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.

References

- Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2015a. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence From the Home Affordable Refinance Program." Working Paper 21512, National Bureau of Economic Research.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015b. "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" Tech. rep., National Bureau of Economic Research.
- Auclert, Adrien. 2015. "Monetary Policy and the Redistribution Channel." Working paper, MIT.
- Avery, Robert B., Neil Bhutta, Kenneth P. Brevoort, Glenn B. Canner, and Christa N. Gibbs. 2010. "The 2008 HMDA Data: the Mortgage Market During a Turbulent Year." *Federal Reserve Bulletin* 96 (April):A169–A211. URL <http://www.federalreserve.gov/pubs/bulletin/2010/articles/HMDA/>.
- Barlevy, Gadi and Jonas D.M. Fisher. 2010. "Mortgage Choices and Housing Speculation." Working Paper 2010-12, Federal Reserve Bank of Chicago.
- Baraja, Martin. 2015. "Fiscal Union Rules." *University of Chicago, mimeo* .
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2015. "House Prices and Consumer Spending." Tech. Rep. 21667, National Bureau of Economic Research.
- Berger, David and Joseph Vavra. 2015. "Consumption Dynamics During Recessions." *Econometrica* 83 (1).
- . 2016. "Dynamics of the U.S. Price Distribution." Working Paper 21732, National Bureau of Economic Research.
- 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.
- Bhardwaj, Geetesh and Rajdeep Sengupta. 2012. "Subprime Mortgage Design." *Journal of Banking & Finance* 36 (5):1503–1519.
- Bhutta, Neil and Benjamin J Keys. 2014. "Interest rates and equity extraction during the housing boom." Kreisman Working Papers Series in Housing Law and Policy 3, University of Chicago.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta, and Stephen Terry. 2014. "Really Uncertain Business Cycles." Tech. rep.
- Brady, Peter J., Glenn B. Canner, and Dean M. Maki. 2000. "The Effects of Recent Mortgage Refinancing." *Federal Reserve Bulletin* 86 (July):441–450. URL <http://www.federalreserve.gov/pubs/Bulletin/2000/0700lead.pdf>.

- 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. URL <http://www.federalreserve.gov/pubs/Bulletin/2002/1202lead.pdf>.
- Caplin, Andrew, Charles Freeman, and Joseph Tracy. 1997. "Collateral Damage: Refinancing Constraints and Regional Recessions." *Journal of Money, Credit, and Banking* :496–516.
- Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2013. "Manufacturing Decline, Housing Booms, and Non-Employment." Working Paper 18949, National Bureau of Economic Research.
- 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, National Bureau of Economic Research.
- 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):367–384.
- Di Maggio, Marco, Amir Kermani, and Christopher Palmer. 2016. "Unconventional Monetary Policy and the Allocation of Credit." Working Paper, Columbia Business School and UC Berkeley.
- Di Maggio, Marco, Amir Kermani, and Rodney Ramcharan. 2014. "Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging." Working Paper, Columbia Business School.
- Farhi, Emmanuel and Iván Werning. 2012. "Fiscal Multipliers: Liquidity Traps and Currency Unions." Tech. rep., National Bureau of Economic Research.
- Fratantoni, Michael and Scott Schuh. 2003. "Monetary policy, housing, and heterogeneous regional markets." *Journal of Money, Credit and Banking* :557–589.
- Fuster, Andreas and Paul 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.

- Fuster, Andreas and Paul S. Willen. 2012. "Payment Size, Negative Equity, and Mortgage Default." Staff Report No. 582, Federal Reserve Bank of New York.
- 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):3–43.
- Gali, Jordi. 2009. *Monetary Policy, inflation, and the Business Cycle: An introduction to the new Keynesian Framework*. Princeton University Press.
- Gali, Jordi and Tommaso Monacelli. 2008. "Optimal monetary and fiscal policy in a currency union." *Journal of International Economics* 76 (1):116–132.
- Garriga, Carlos, Finn E. Kydland, and Roman Sustek. 2013. "Mortgages and Monetary Policy." Working Paper 19744, National Bureau of Economic Research.
- Gertler, Mark and Peter Karadi. 2011. "A model of unconventional monetary policy." *Journal of monetary Economics* 58 (1):17–34.
- Goodman, Laurie S. 2012. "Helping Responsible Homeowners Save Money through Refinancing." Testimony to the U.S. Senate Committee on Banking, Housing and Urban Affairs, April 25.
- Greenwald, Daniel. 2016. "The Mortgage Credit Channel of Macroeconomic Transmission." Tech. rep.
- 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. 2015. "Regional Redistribution Through the US Mortgage Market." Tech. rep., National Bureau of Economic Research.
- 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, National Bureau of Economic Research.

- 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* :pp. 215–287.
- . 2013. "The Ins and Outs of LSAPs." Presented at the Federal Reserve Bank of Kansas City's Jackson Hole Symposium.
- Lee, Donghoon, Christopher Mayer, and Joseph Tracy. 2012. "A New Look at Second Liens." In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai. University of Chicago Press, 205–234.
- Lustig, Hanno and Stijn Van Nieuwerburgh. 2010. "How much does household collateral constrain regional risk sharing?" *Review of Economic Dynamics* 13 (2):265–294.
- McKinnon, Ronald I. 1963. "Optimum currency areas." *The American Economic Review* :717–725.
- Mian, Atif, Kamal Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump*." *The 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.
- Mundell, Robert A. 1961. "A theory of optimum currency areas." *American Economic Review* :657–665.
- Nakamura, Emi and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *The American Economic Review* 104 (3):753–792.
- 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.
- Stroebel, Johannes and Joseph Vavra. 2016. "House prices, local demand, and retail prices." Tech. rep., National Bureau of Economic Research.
- Suarez-Serrato, Juan Carlos and Philippe Wingender. 2010. "Estimating local fiscal multipliers." *University of California at Berkeley, mimeo* .
- Tracy, Joseph and Joshua Wright. 2012. "Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses." Staff Report No. 562, Federal Reserve Bank of New York.
- 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." Tech. rep.
- Wong, Arlene. 2015. "Population Aging and the Transmission of Monetary Policy to Consumption." Tech. rep.
- Woodford, Michael. 2005. "Interest and prices: Foundations of a theory of monetary policy."

Appendix

A.1 Data

A.1.1 HMDA

In all our analysis using HMDA data, we only retain applications that led to originations (action code = 1); however, rather than using the action date, we always use the application date. We drop multifamily properties and mortgages with an origination amount $>$ \$3 million (about 0.015% of loans).

A.1.2 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 in between January 2007 and December 2010. The credit records in the dataset cover 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 credits [HELOCs], 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 less loans in each category and never have more than three of any of these types of loans outstanding.¹ This amounts to about 94% of the population of loan IDs, and about 91% of the total loan balance. In creating this loan-level dataset, we assume that the month that 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 zipcode, 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, zipcode, 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. Due 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.

¹This 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.

A.1.2.1 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 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 dataset 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 zipcode 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.1.2.2 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 balance 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 an McDash purpose type of refinance or unknown/other. Our algorithm to identify whether or not 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 dataset 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 zipcode.
- 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.

For each of these cases, we can thus 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 corresponding Equifax record to our refinance loan that has the same Equifax open date and an origination balance of less than 25% of the refinance loan's origination balance, and add the balance of these piggyback second liens to the refinance origination amount when calculating cash out amounts. 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 .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-3, we compare the total estimated quarterly cash-out amounts to those estimated on prime conventional loans by Freddie Mac.² 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.1.2.3 Measuring CLTVs

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 and use this combined balance to compute CLTVs.

²To 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.

We then compute CLTV_X (where X is either 80 or 100) at the MSA level as (balance [first-lien only, from McDash] of loans with CLTV > X%) / (balance of all matched first-liens outstanding in McDash).

A.2 Additional Figures and Tables

Figure A-1: CLTV_100 vs. CLTV_80 in November 2008. Based on CRISM data. Size of circles is proportional to MSA 2008 population.

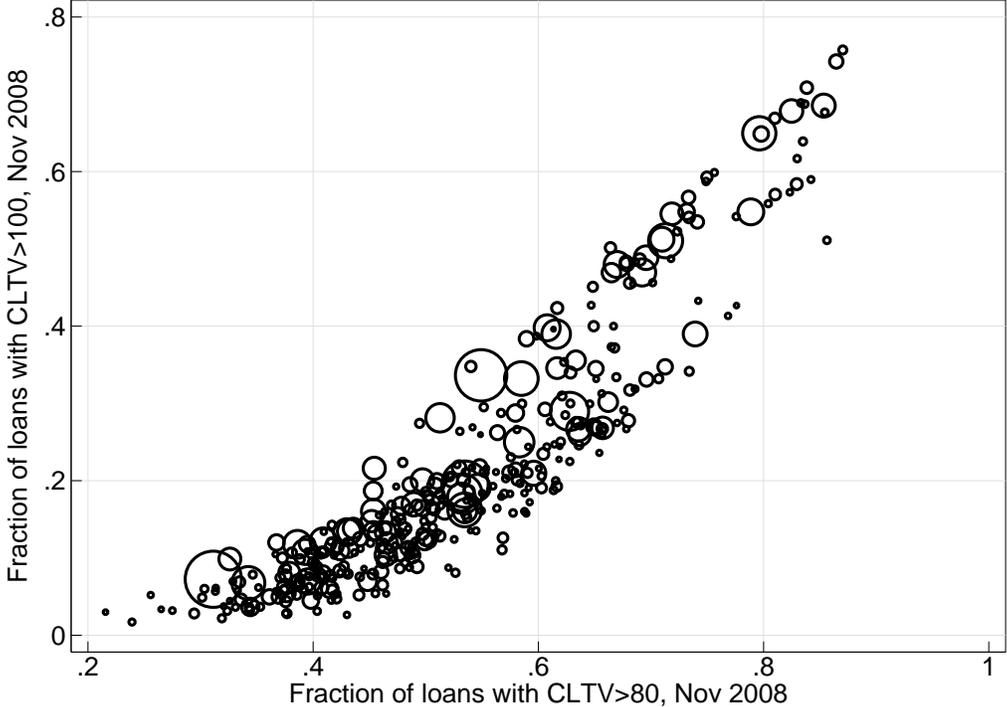


Figure A-2: Total Mortgage Origination Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008.

Figure shows total mortgage origination volume (purchase and refinance) 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. MSA quartile groups are defined the same way as in Figure 5. The vertical line indicates the month of the QE1 announcement (November 2008).

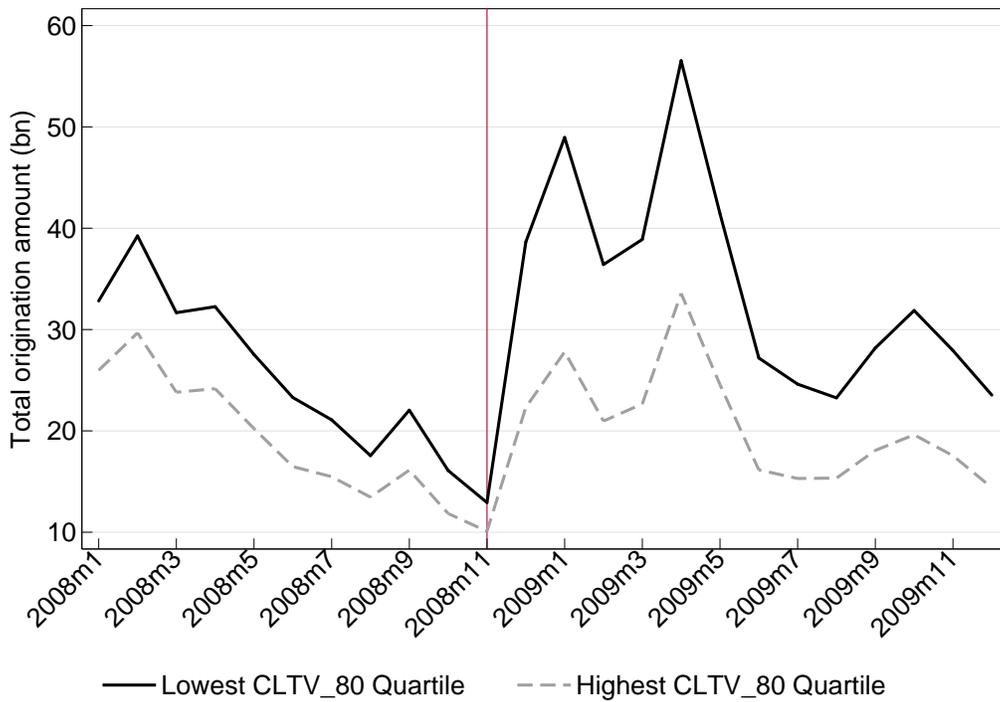
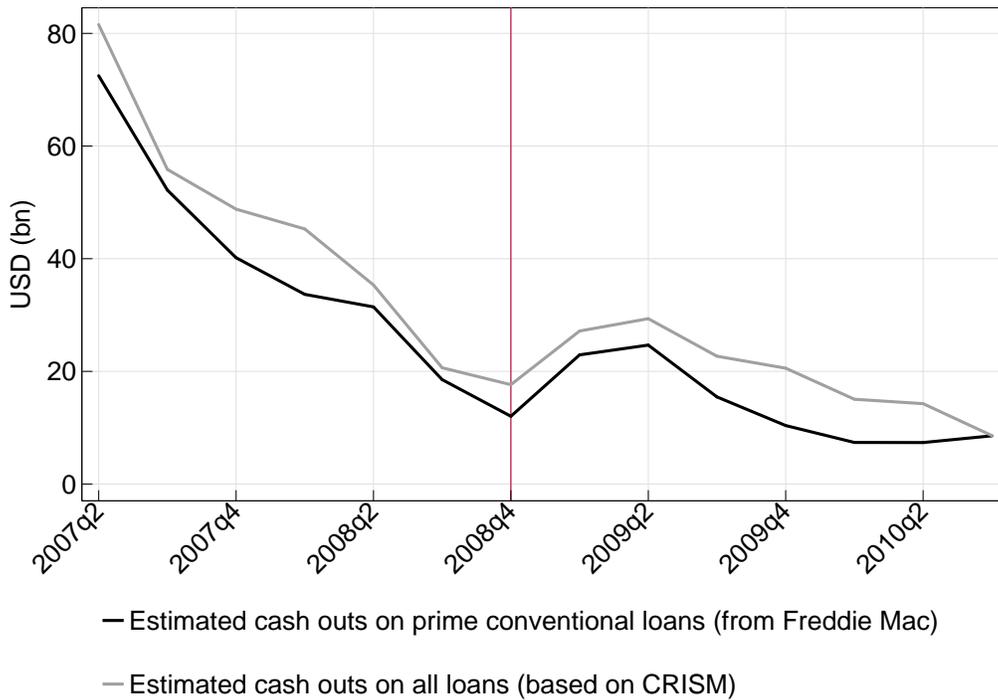


Figure A-3: Estimated Cash-out Amounts from Freddie Mac vs. in our CRISM data

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 includes FHA/VA loans) after scaling up as explained in Section A.1.2.2.



A.2.1 MSA groups used in Figures 5, 6, and 7

MSAs in the quartile with highest CLTV_80 in November 2008:

Akron, OH; Anderson, IN; Ann Arbor, MI; Bakersfield, CA; Bangor, ME; Battle Creek, MI; Bay City, MI; Bradenton-Sarasota-Venice, FL; Canton-Massillon, OH; Cape Coral-Fort Myers, FL; Cleveland-Elyria-Mentor, OH; Columbus, OH; Dalton, GA; Danville, IL; Dayton, OH; Deltona-Daytona Beach-Ormond Beach, FL; Detroit-Livonia-Dearborn, MI (Metro Div.); El Centro, CA; Elizabethtown, KY; Elkhart-Goshen, IN; Fairbanks, AK; Flint, MI; Fort Lauderdale-Pompano Beach-Deerfield Beach, FL (Metro Div.); Fort Walton Beach-Crestview-Destin, FL; Fort Wayne, IN; Fresno, CA; Grand Rapids-Wyoming, MI; Greeley, CO; Gulfport-Biloxi, MS; Hagerstown-Martinsburg, MD-WV; Hanford-Corcoran, CA; Holland-Grand Haven, MI; Indianapolis-Carmel, IN; Jackson, MI; Jacksonville, FL; Jefferson City, MO; Johnstown, PA; Kalamazoo-Portage, MI; Kankakee-Bradley, IL; Kansas City, MO-KS; 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 (Metro Div.); Modesto, CA; Monroe, MI; Muskegon-Norton Shores, MI; Naples-Marco Island, FL; Niles-Benton Harbor, MI; Ocala, FL; Orlando-Kissimmee, FL; Palm Bay-Melbourne-Titusville, FL; Palm Coast, FL; Panama City-Lynn Haven-Panama City Beach, FL; Pascagoula, MS; Pensacola-Ferry Pass-Brent, FL; Phoenix-Mesa-Scottsdale, AZ; Port St. Lucie, FL; Punta Gorda, FL; Reno-Sparks, NV; Riverside-San Bernardino-Ontario, CA; Sacramento-Arden-Arcade-Roseville, CA; Saginaw-Saginaw Township North, MI; Salinas, 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 (Metro Div.); Weirton-Steubenville, WV-OH; West Palm Beach-Boca Raton-Boynton Beach, FL (Metro Div.); Wheeling, WV-OH; Winchester, VA-WV; Worcester, MA; Youngstown-Warren-Boardman, OH-PA; Yuba City, CA; Yuma, AZ.

MSAs in the quartile with lowest CLTV_80 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; Billings, MT; Binghamton, NY; Bismarck, ND; Blacksburg-Christiansburg-Radford, VA; Bloomington-Normal, IL; Boulder, CO; Bridgeport-Stamford-Norwalk, CT; Buffalo-Niagara Falls, NY; Burlington-South Burlington, VT; Cedar Rapids, IA; Charleston, WV; Charlottesville, VA; Cleveland, TN; College Station-Bryan, TX; Corvallis, OR; Cumberland, MD-WV; Davenport-Moline-Rock Island, IA-IL; Dubuque, IA; Duluth, MN-WI; Durham-Chapel Hill, NC; Edison-New Brunswick, NJ (Metro Div.); Elmira, NY; Erie, PA; Eugene-Springfield, OR; Fargo, ND-MN; Florence, SC; Fort Collins-Loveland, CO; Fort Smith, AR-OK; Glens Falls, NY; Grand Forks, ND-MN; Grand Junction, CO; Great Falls, MT; Green Bay, WI; Greenville-Mauldin-Easley, SC; Harrisburg-Carlisle, PA; Harrisonburg, VA; Hartford-West Hartford-East Hartford, CT; Hickory-Lenoir-Morganton, NC; Honolulu, HI; Hot Springs, AR; Houma-Bayou Cane-Thibodaux, LA; Huntsville, AL; Iowa City, IA; Jacksonville, NC; Johnson City, TN; Jonesboro, AR; Kingsport-Bristol-Bristol, TN-VA; Knoxville, TN; La Crosse, WI-MN; Lafayette, LA; Lake Charles, LA; Lancaster, PA; Lawrence, KS; Lebanon, PA; Longview, WA; Lubbock, TX; Lynchburg, VA; Madison, WI; Midland, TX; Missoula, MT; Mobile, AL; Monroe, LA; Mount Vernon-Anacortes, WA; Nassau-Suffolk, NY (Metro Div.); New Orleans-Metairie-Kenner, LA; New York-White Plains-Wayne, NY-NJ (Metro Div.); Ocean City, NJ; Oshkosh-Neenah, WI; Philadelphia, PA (Metro Div.); Pittsburgh, PA; Pittsfield, MA; Pocatello, ID; Raleigh-Cary, NC; Rapid City, SD; Reading, PA; Roanoke, VA; Rochester, NY; Salem, OR; San Angelo, TX; San Francisco-San Mateo-Redwood City, CA (Metro Div.); Scranton-Wilkes-Barre, PA; Seattle-Bellevue-Everett, WA (Metro Div.); Sheboygan, WI; Sioux City, IA-NE-SD; Sioux Falls, SD; Spokane, WA; State College, PA; Syracuse, NY; Trenton-Ewing, NJ; Tulsa, OK; Victoria, TX; Waterloo-Cedar Falls, IA; Wenatchee-East Wenatchee, WA; Williamsport, PA; Wilmington, NC; Yakima, WA; York-Hanover, PA.