

Regional Heterogeneity and Monetary Policy*

Martin Beraja Andreas Fuster Erik Hurst Joseph Vavra[†]

December 11, 2016

Abstract

We argue that the time-varying regional distribution of housing equity shapes the aggregate consequences of monetary policy through its influence on mortgage refinancing. Using detailed loan-level data, we begin by showing that: (i) the refinancing response to interest rate cuts is strongly affected by regional differences in housing equity, and (ii) both these differences and overall refinancing vary over time with changes in the regional distribution of house price growth and unemployment. Then, we build a heterogeneous household model of refinancing in order to derive aggregate implications of monetary policy from our regional evidence. We find that the 2008 equity distribution made spending in depressed regions less responsive to interest rate cuts, thus dampening aggregate stimulus and increasing regional consumption inequality, whereas the opposite occurred in some earlier recessions. Taken together, our results strongly suggest that monetary policy makers should track the regional distribution of equity over time.

*First draft: May 2015. We thank Caitlin Gorbach and Karen Shen for excellent research assistance, our discussant Wouter Den Haan and seminar participants at the University of Minnesota, NYU, MIT Sloan, Berkeley Haas, ECB Annual Research Conference, 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

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 interest rate cuts incentivize households to refinance their mortgages and extract housing equity to fund current consumption.¹ However, in this paper we argue that the strength of this *refinancing channel* for stimulating aggregate consumption is time-varying because it is influenced by the distribution of housing equity in the economy which we show varies dramatically across recessions in the US. Furthermore, we show that interest rate cuts likely amplify regional consumption inequality through this refinancing channel 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 significantly and most strongly in regions with the largest declines in economic activity (e.g. Nevada or Spain). Nevertheless, under alternative distributions of house prices, such as those in the 2001 recession, we show that interest rates cuts can both stimulate aggregate consumption and reduce (rather than amplify) regional consumption inequality. Since the distribution of equity both varies across time and changes the consequences monetary policy, we conclude that it is important for policy makers to track this variation when making decisions.

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

In the first part of the paper, we provide empirical evidence about how regional heterogeneity in housing equity affects the refinancing channel of monetary policy. We present two complementary sets of facts that exploit different sources of variation. The first set explores the regional response in the US to the interest rate decline that followed the first round of the Federal Reserve's large-scale asset purchase program—commonly known as the first round of quantitative easing (QE1). Because the US is a monetary union where interest rates are identical across regions and interest-rate-policy is independent of regional variation in economic activity, we could

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

potentially have used other interest rate changes to study the refinancing channel and its interaction with the distribution of housing equity in the economy. However, QE1 provides a unique opportunity to do so both because of its magnitude—the 30-year fixed-rate mortgage rates fell by around 1 percent over the 6 months following the announcement—as well as the large variation in regional housing market conditions at the time of the announcement. The second set moves away from this particular QE1 episode and compares regional house price, economic activity, and refinancing patterns across different recessions in the US.

Using a variety of loan-level data sources, we document three facts regarding the regional response to QE1. (1) There was a boom in mortgage originations that resulted from households refinancing their existing mortgages. (2) Refinancing activity and the amount of equity extracted (conditional on refinancing) increased more in metropolitan statistical areas (MSAs) that had lower unemployment and where homeowners had more equity on the eve of QE1. (3) MSAs where homeowners refinanced the most right after QE1 (early 2009), also experienced the largest increase in car purchases thereafter.

The first two facts strongly suggest that monetary policy affects refinancing and equity extraction in a way that depends on the regional distribution of equity and economic activity in the economy. The third fact suggests that this distribution also affects the transmission from interest rates cuts to household spending, but there are some reasons that one should be cautious when extrapolating our evidence to *aggregate* household spending. The first is that general equilibrium, aggregate effects are differenced out when using regional variation—a point made by Beraja et al (2016). In particular, interest rates declines lower mortgage payments for refinancing borrowers, who then increase spending, but reduces income and spending for lenders. As [Greenwald \(2016\)](#) shows, accounting for offsetting effects by mortgage lenders can substantially dampen aggregate spending responses in general equilibrium. In an environment with a large fraction of fixed rate mortgages like in the US, this offset can be large. The second is that the increase in car purchases in regions with high equity and refinancing after QE1 might be driven by some confounding effect or alternative channel of monetary policy whose effects are also correlated with equity levels. Partly motivated by these concerns, the second part of the paper presents a theoretical analysis of the refinancing channel.

We conclude the empirical part of the paper by documenting three facts regarding the time-series variation in the regional distribution of house price changes,² unemployment changes, and refinancing in the US. (i) There is substantial variation across time in the distribution of house price changes across regions in the US. For example, the average and standard deviation of house price declines were considerably larger during the Great Recession than following the 2001 recession. (ii) There is substantial variation across time in the correlation between house price changes and unemployment changes across US regions. For example, the correlation was considerably more negative during the Great Recession than following the 2001 recession. (iii) Aggregate re-

²We use house price changes to proxy for housing equity changes because our housing equity data is not available prior to 2005.

financing propensity increased both in the Great Recession and in the 2001 recession. However, refinancing propensities increased more in low-unemployment than high-unemployment MSAs during the Great Recession, whereas the opposite was true following the 2001 recession.

The third fact complements the evidence from our QE1 event-study, further showing how the response of refinancing activity to interest rate cuts is affected by regional economic conditions. In particular, when combined with the other two facts, it suggests that the strength of the refinancing channel could be stronger under regional distributions different than the ones we observed during the Great Recession, e.g., the one observed during the 2001 recession. Moreover, the striking variation over time in these regional patterns leads us to believe that if monetary policy can stimulate aggregate spending through the refinancing channel, then both the ability to do so as well as the consequences regarding regional consumption inequality are necessarily time-varying. Analyzing these issues requires a formal model because of the need to perform counterfactuals that go beyond what can be done on purely empirical grounds. Hence, we follow with such theoretical analysis in order to complement our empirical analysis.

In the second part of the paper we build a general equilibrium, incomplete-markets, heterogeneous agents model with both mortgage borrowers and lenders. The key distinguishing features of the model are that mortgage borrower households face both idiosyncratic and regional income and house price risk, and can refinance their mortgage and extract housing equity by paying a fixed cost. The goal of the theory is twofold. First, to clarify how the transmission from interest rate cuts to aggregate spending through the refinancing channel depends on the distribution of housing equity. Second, to understand conditions under which monetary policy makers may face a trade-off between stimulating aggregate spending and increasing regional consumption inequality. Through a series of counterfactual exercises, we show that effects of interest rate changes on both aggregate activity and regional consumption inequality change dramatically with the distribution of housing equity in the economy.

Our first theoretical results focus on the consequences of interest rate cuts in a benchmark economy that matches the joint distribution of housing equity and income that was actually observed in 2008. Moreover, we pick the baseline parameters in our quantitative model in order to match the cross-region effects of QE that we documented in the first part of the paper. Then, we show that a decline in interest rates of the magnitude observed after QE1 modestly raises aggregate spending. This implies that the spending offset coming from mortgage lenders in general equilibrium is not one-for-one in our environment. The main reason is that 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. However, under 2008 economic conditions, this effect is quantitatively small in the aggregate. Next, we show that when interest rates decline under 2008 economic conditions there is a large increase in consumption inequality across regions. High equity homeowners have

stronger refinancing and spending responses to declines in interest rates than low equity homeowners. Hence, the variance of consumption across regions rises because income and equity were highly correlated during the Great Recession. Thus, our model implies that monetary policy in 2008 modestly boosted aggregate spending and did so at the cost of much larger consumption 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. In particular, we look at the consequences of interest rates cuts in counterfactual economies with alternative distributions of equity and income across regions. We find that as the average level of housing equity in the economy rises, both aggregate spending and regional consumption inequality respond more strongly to interest rate cuts. However, the spending response grows more rapidly with house prices and equity than the inequality response. This means that the trade-off between stimulus and inequality is not as bad in recessions with house price increases, such as 2001, as it is in recessions with house price declines, such as 2008. Conversely, a reduction in the variance of house price growth across regions leads to 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, large regional variance in house prices, and a highly positive correlation between house prices and income. 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 equity change the response of the economy to interest rate cuts? 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. How-

ever, 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. Hence, 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 equity 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.

After arguing that the distribution of equity in 2008 hampered the ability of monetary policy to stimulate the economy through the refinancing channel, we show that various policies can be complemented with interest rate cuts in order to do so. In particular, we show that both targeted debt reduction and relaxation of collateral constraints for refinancing can further help monetary policy to stimulate aggregate spending. Policies along these lines were implemented during the Great Recession through the Home Affordable Modification and Refinance Programs (HAMP and HARP), and our results show that such mortgage market interventions can interact importantly with monetary policy. The effects of a combined mortgage market intervention paired with a reduction in rates is larger than the sum of the individual effects. In addition, we explore the role of macroprudential policy and show that time-varying countercyclical leverage requirements have the potential to both dampen the initial depth of house price induced recessions and to modestly strengthen the stimulative effects of monetary policy during such recessions.

Finally, we show that our model implications continue to hold under a variety of alternative assumptions. In these model extensions we account for the presence of adjustable-rate mortgages, calibrated to match the observed regional share in the data; allow for the fact that in this recession large busts were preceded by large booms; allow for cash-out and non-cash-out refinancing; extend our baseline environment with unanticipated one-time interest rate shocks to environments with stochastic, transitory rate changes; show that our results are insensitive to assumptions about short-long interest rate spreads; and explore alternative assumptions on the lender side of the economy about the share of mortgage debt ultimately held by domestic consumers.

The remainder of the paper proceeds as follows: Section 2 discusses the connection to the literature. Section 3 describes the data sources for our empirical work. Section 4 explores the regional distribution of refinancing, cash-out, and spending activity around QE1. Section 5 doc-

uments time-series variation in the distribution of house price changes and refinancing propensities. Section 6 explores the theoretical implications of these patterns in a quantitative model of refinancing. Section 7 presents several robustness exercises and model extensions. Section 8 concludes.

2 Literature Review

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.³ 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 which 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. Together, these features give rise to what we call the "refinancing channel" of monetary policy. Importantly, the strength of this channel depends on the distribution of housing equity in the economy, which exhibits substantial time-variation in the data. This means that policy makers must pay attention to this distribution when determining the rate necessary to achieve a given level of stimulus, 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.⁴ 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.⁵ This channel is complementary but distinct from ours as it arises from monetary policy changing the value of collateral which in turn influences

³See [Woodford \(2005\)](#) and [Galí \(2009\)](#) for canonical expositions.

⁴See 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.

⁵See 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.

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.

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 housing equity 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 refinancing channel of monetary policy. Our focus on the implications of realistic modeling of household borrowing and how it interacts with heterogeneity in the economy also parallels many of the themes in [Auclert \(2015\)](#) who argues that the covariance of the marginal propensity to consume with interest rate exposure across agents matters for aggregate consumption responses to interest rate changes. His analysis abstracts from 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-varying distribution of housing equity.

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.⁶ [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. \(2015\)](#) use detailed data from the Home Affordable Refinancing Program

⁶There 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.

(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. 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.

3 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 details.

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 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.⁷ 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.⁸ 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.

⁷See, for instance, [Avery et al. \(2010\)](#) or [Dell'Ariccia, Igan, and Laeven \(2012\)](#).

⁸There 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.

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. 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.⁹ 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).¹⁰ 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 com-

⁹A 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.

¹⁰We use the zip-code-level index if available, and the MSA-level index if not.

pany 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.¹¹

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. Second, we also show aggregate refinancing trends within the U.S. using published statistics from the Mortgage Bankers Associations (MBA) Refinance Index. Third, as alluded to above, we use house price data from CoreLogic at the zip code or MSA level to measure local house price appreciation. 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.

4 The Refinancing Channel and Regional Heterogeneity: Evidence from QE1

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

1. Mortgage originations increased substantially because households refinanced their existing

¹¹Changes 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/ReportDocuments/May-14-Refi_Report.pdf for statistics on HARP.

mortgages.

2. Refinancing activity and equity extraction (conditional on refinancing) was higher in MSAs where homeowners had more equity and unemployment was lower prior to QE1.
3. Car purchases increased the most in MSAs where homeowners refinanced the most right after QE1.

4.1 Aggregate Trends in Mortgage Activity Around QE1

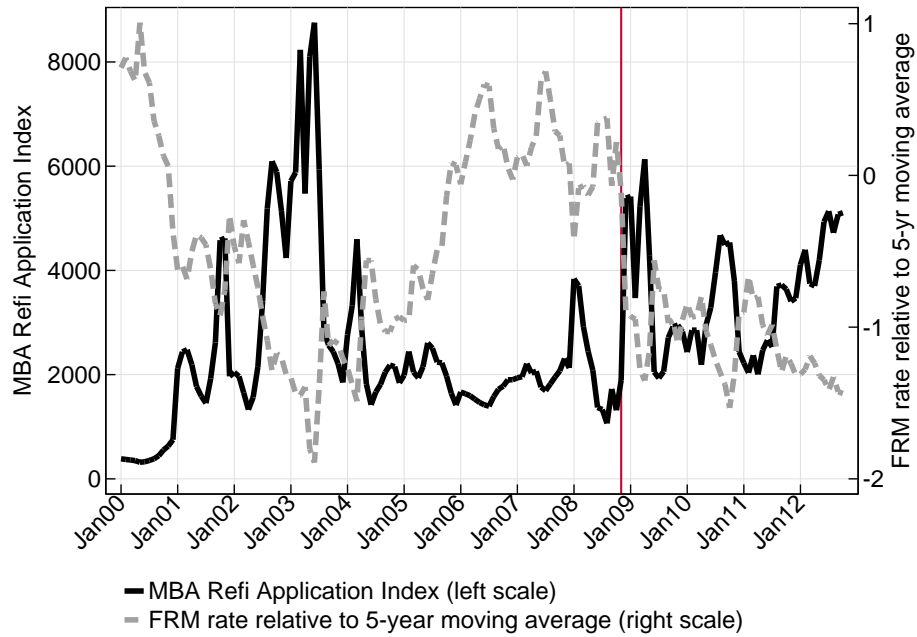
Panel (a) of Figure 1 shows time series patterns in the monthly MBA Refinance Index over 2000 to 2012 (solid line). The figure also includes the difference between the average 30-year fixed-rate mortgage (FRM) rate (also from MBA) in month t and the average of the 30-year mortgage rate over the prior 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 FRM 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 applied for the mortgage (not the month in which the loan was ultimately originated, which is usually 1-3 months later). The dark (light) shaded area measures the national monthly dollar volume of refinancing (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).¹² 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

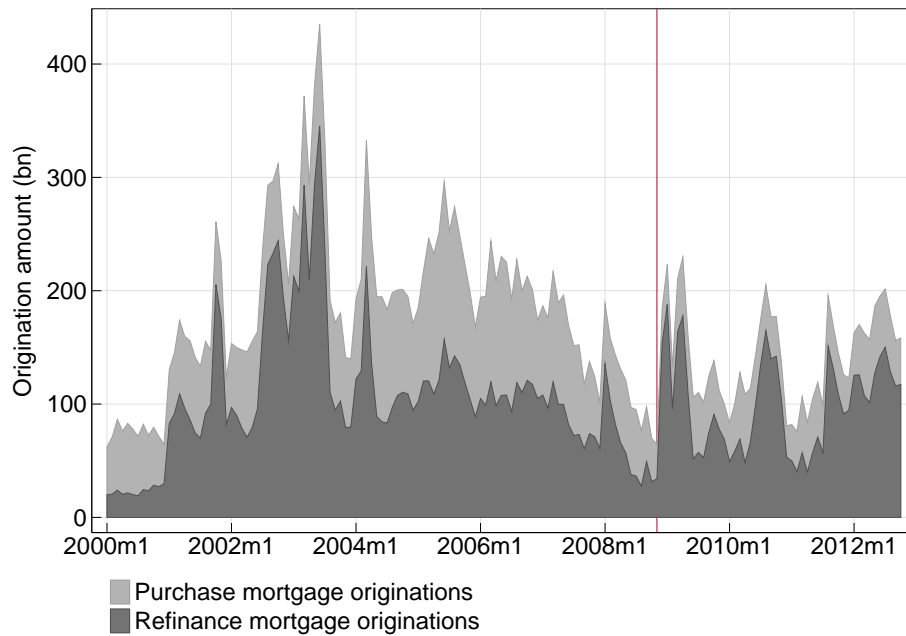
¹²Our HMDA series is a slightly different construct than the MBA Refinance Index. The MBA Refinance Index measures refinance application activity, while in HMDA we only retain applications that ultimately lead to originations. Also, the MBA series does not include broker/correspondent/wholesale origination channels.

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).

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.

4.2 Regional Variation in Loan-To-Value Ratios Prior to QE1

To explore regional variation in loan origination around QE1, we need to define our notion of regions. Throughout our paper, we use metropolitan statistical areas (MSAs) to define regions.¹³ We begin exploring the regional heterogeneity in equity values by documenting the extent to which cumulative loan-to-value ratios (CLTVs) evolved differentially across different MSAs during the 2007 and 2008 period.

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

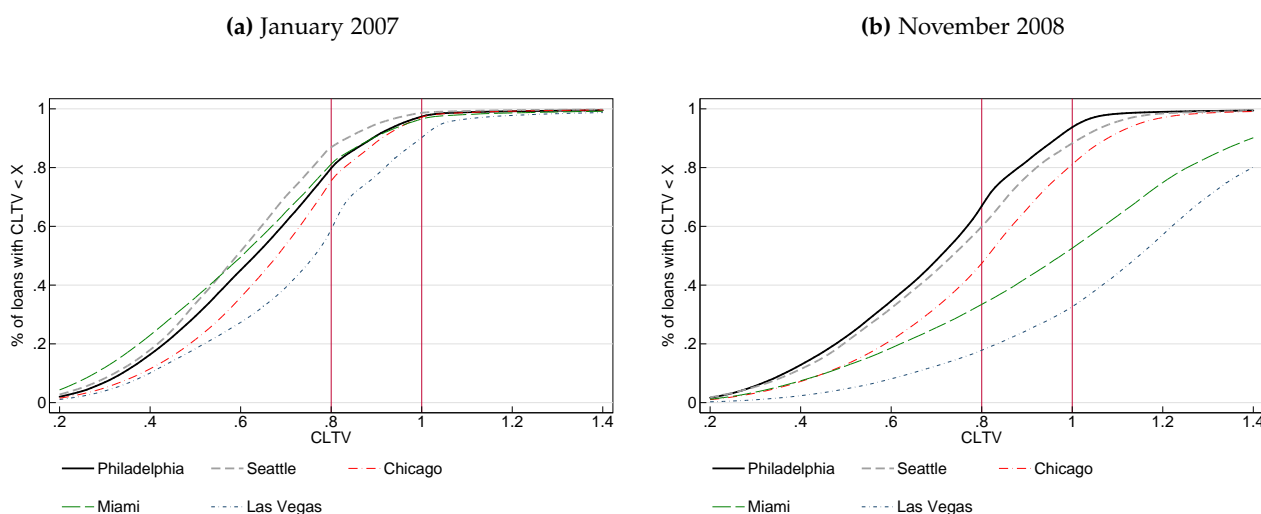


Figure shows the cumulative distribution function of the CLTV distribution of borrowers in five illustrative MSAs in January 2007 (panel a) and November 2008 (panel b). CLTV is measured for each household in an MSA using CRISM data on total current mortgage debt divided by an estimate of their current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

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

¹³For 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 cross-walk available at http://mcdc.missouri.edu/data/georef/zcta_master.Metadata.html.

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 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 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.

For completeness, Table A-1 in the Online Appendix 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. We conclude this section by noting that there was wide variation across MSAs in house price declines during the Great Recession and that these declines were the main driver of high-CLTV shares across the regions prior to QE1. 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. Relatedly, we find that MSAs with the largest increases in local unemployment rates were also associated with having relatively more mortgage owners with a CLTV above 0.8, which is unsurprising given that a large literature has documented that house price declines were associated with weakening labor markets during this period (Charles, Hurst, and Notowidigdo, 2013; Mian and Sufi, 2014).

4.3 Regional Variation in Mortgage Activity Around QE1

We now show that in the months after QE, refinancing activity was much higher in regions where individuals had sufficient equity in their home and the unemployment rate was relatively low. We use our measures of $CLTV_{80}$ and $CLTV_{100}$ (defined above) in order to summarize the amount of equity individuals have in their home within a region. 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 3 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 where the vast majority of mortgage owners had a *CLTV* greater than 0.8. The bottom *CLTV_80* quartile includes MSAs where most mortgage owners had *CLTV*s below 0.8.¹⁴

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 QE1.¹⁵ Refinancing volumes evolved the same between high and low *CLTV_80* MSAs up to 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.¹⁶ 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 a mortgage application is initially made and when the actual mortgage origination takes place.¹⁷ The key point from Figure 3 is that regardless of the metric for measuring local refinancing activity, refinancing activity was significantly higher in MSAs where *CLTV_80* was low.

In Table A-2 in the Online Appendix, we show results from a population-weighted-regression

¹⁴Appendix A.3.1 lists the MSAs within each of the *CLTV_80* quartiles.

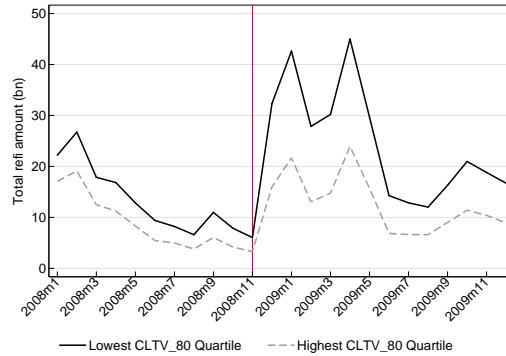
¹⁵Although we show the results monthly, we could have explored weekly refinancing totals. Fuster and Willen (2010) show that refinancing applications in HMDA jumped starting the day of the QE1 announcement. Also, Figure A-2 in the Appendix shows that patterns are nearly identical if we focus on total originations rather just refinancing originations (which is not surprising, given that most mortgage origination activity during this period is refinancing activity, as discussed above).

¹⁶The 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.

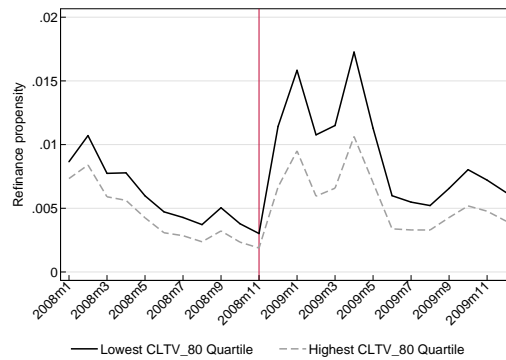
¹⁷During this time both the loan underwriting and closing procedures take place.

Figure 3: Mortgage Refinance Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008

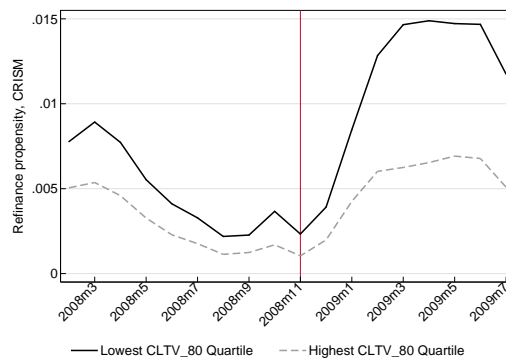
(a) HMDA: Refinance Origination Volumes



(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).

of the change in CRISM refinancing propensity between November 2008 and February of 2009 in an MSA onto either *CLTV_80* or *CLTV_100* within the MSA during November of 2008, the change in an MSA's unemployment rate between January 2007 and November 2008, and a vector of control variables defined from the pooled 2007/2008 American Community Survey. 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. We report the results for a variety of specifications as well as when the dependent variable is the change in the local 3-month refinancing rate between January-March 2009 relative to September-November 2008 instead of the CRISM refinancing propensity.

Collectively, the results from Table A-2 and Figure 3 show that there were large regional differences in refinancing activity in response to QE1. Regions that experienced the largest house price declines and had the largest increases in unemployment were the least responsive to QE1 in terms of their subsequent mortgage refinancing behavior.

4.4 Regional 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.¹⁸ 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.¹⁹

Figure 4 shows the amount of equity removed during the refinancing process for the top and bottom quartiles of the *CLTV_80* distribution around the QE period. Panel (a) we shows estimated dollar amounts, while panel (b) shows equity removed in a given month relative to the total outstanding mortgage balance in the prior month.²⁰ The total amount of equity removed during the refinancing process sums over people who removed no equity during refinancing, people 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

¹⁸See, for example, Brady, Canner, and Maki (2000), Canner, Dynan, and Passmore (2002), Hurst and Stafford (2004) and Bhutta and Keys (2014).

¹⁹See, for example, Mian, Rao, and Sufi (2013).

²⁰Since the CRISM data does not cover the whole mortgage market, we scale up dollar amounts in CRISM as explained in Appendix A.1.2.2.

Figure 4: Cash Out Refinancing Activity in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008, CRISM Data.

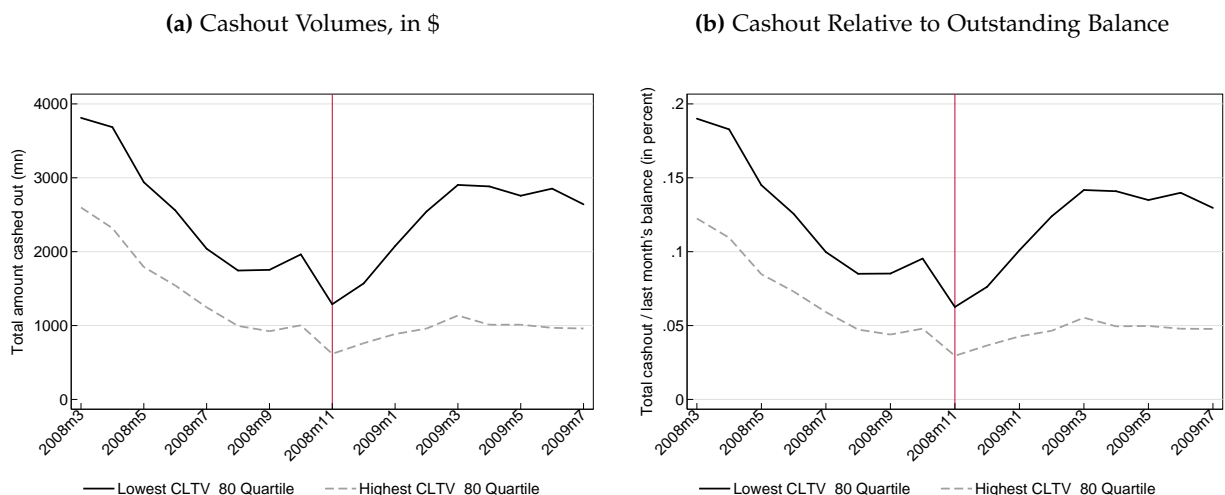


Figure shows total cash removed from the mortgage for the top and bottom quartiles of MSAs with respect to CLTV_80 in November 2008. Panel (a) shows dollar volumes, which are calculated by scaling up amounts from the CRISM data to the market level, as explained in the Appendix. Panel (b) shows cash out 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 3. In both panels, the vertical line indicates the month of the QE1 announcement (November 2008).

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.

In Table A-3 in the Online Appendix, we show results from regressions similar to the ones in Online Appendix Table A-2, but using the change in equity removed (relative to outstanding balance) between November 2008 and February 2009 as the dependent variable. Echoing the results in Figure 4, we find that a one standard deviation change in CLTV_80 is associated with roughly a 0.016 percentage point change in cashout fractions, roughly 32 percent of the mean of that variable. Adding controls only strengthens the coefficient on CLTV_80. Moreover, we ask: how much of the differences in cash-outs 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, we additionally control for the change in refinancing propensity over the same period. We find that the coefficient on the change in refinancing propensity and on CLTV_80 are both strongly significant. Hence, high CLTV_80 MSAs both refinanced less and removed less equity, conditional on refinancing. This is intuitive, since these places have less equity to remove conditional on refinancing.

Figure 5: New Auto Sales in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008

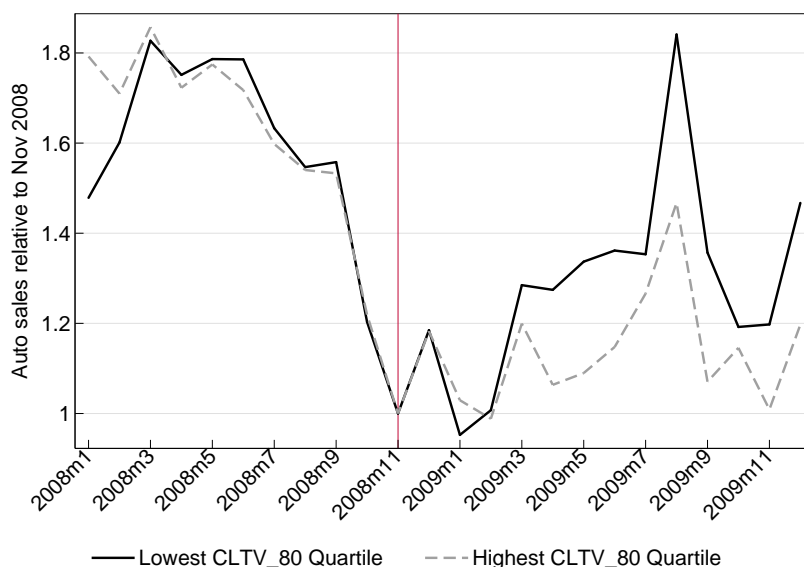


Figure shows the total volume of new car purchases (relative to November 2008) for the top and bottom quartiles of MSAs with respect to CLTV_80 in November 2008. Car purchase volumes come from R.L. Polk. The vertical line indicates the month of the QE1 announcement (November 2008).

Instead of looking at equity removed and implicitly discussing the potential for differential local spending responses to the QE1 induced refinancing, we now explore more direct measures of spending. Figure 5 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 5. 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 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.²¹

Table A-4 in the Online Appendix shows MSA-level regressions, similar to Table A-3, but using the change in the log of new car sales from pre-QE1 (September-November 2008) to post-QE1 (March-May 2009) as our dependent variable. We find that a one standard deviation change in the *CLTV_80* is associated with roughly a one-third of a standard deviation change in the log-change of auto sales. Furthermore, in order to directly test the link between refinancing, equity extraction and auto sales, we add both refinancing and cash-out propensities as independent variables, finding that they are correlated with stronger auto sales over this period.

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.

5 The Time-Varying Refinancing Channel in the US

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. Because these locations also had the largest increases in unemployment, monetary policy stimulated spending the least in the hardest hit locations. We now show that there is substantial variation across time in the distribution of house price changes and refinancing propensities across regions. Hence, when combined with the evidence from QE1, it suggests that the regional distribution of housing equity affects the refinancing channel and produces time-varying effects of monetary policy. In the next part of the paper, we study theoretically how this time-varying refinancing channel arises.

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 A-10 in the Online Appendix 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 6 shows the mean of the annual real house price growth distribution across states calculated using CoreLogic data for each year from 1976-2013.²² 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

²¹In 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.

²²State-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 6: Changes in Distribution of House Price Changes Across Time

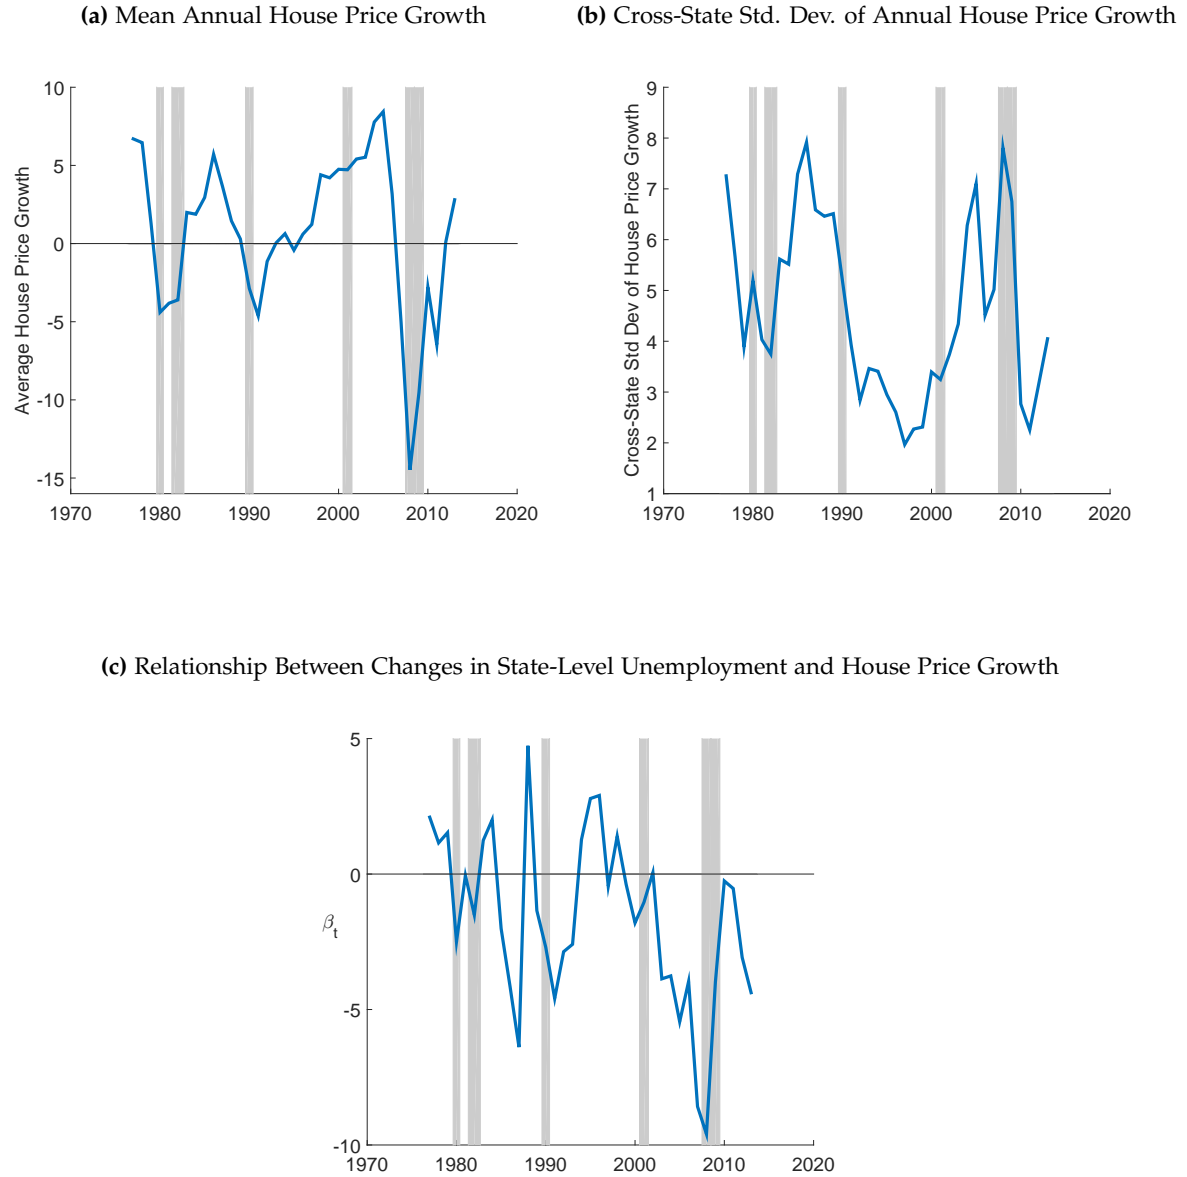


Figure shows the mean, and standard deviation of the state-level house price growth distribution across time as well as the cross-state relationship between house price growth and changes in the unemployment rate. All calculations are population weighted by state.

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 7 illustrates that the pass-through of monetary policy through the mortgage market was much stronger following the 2001 recession than in the 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 period.²³ 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 3.

Panel (b) shows that the time-series patterns for monthly refinancing by unemployment quar-

²³The 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.

Figure 7: Refinancing Propensities for High and Low Unemployment Change Quartiles, 2008 Recession and 2001 Recession.

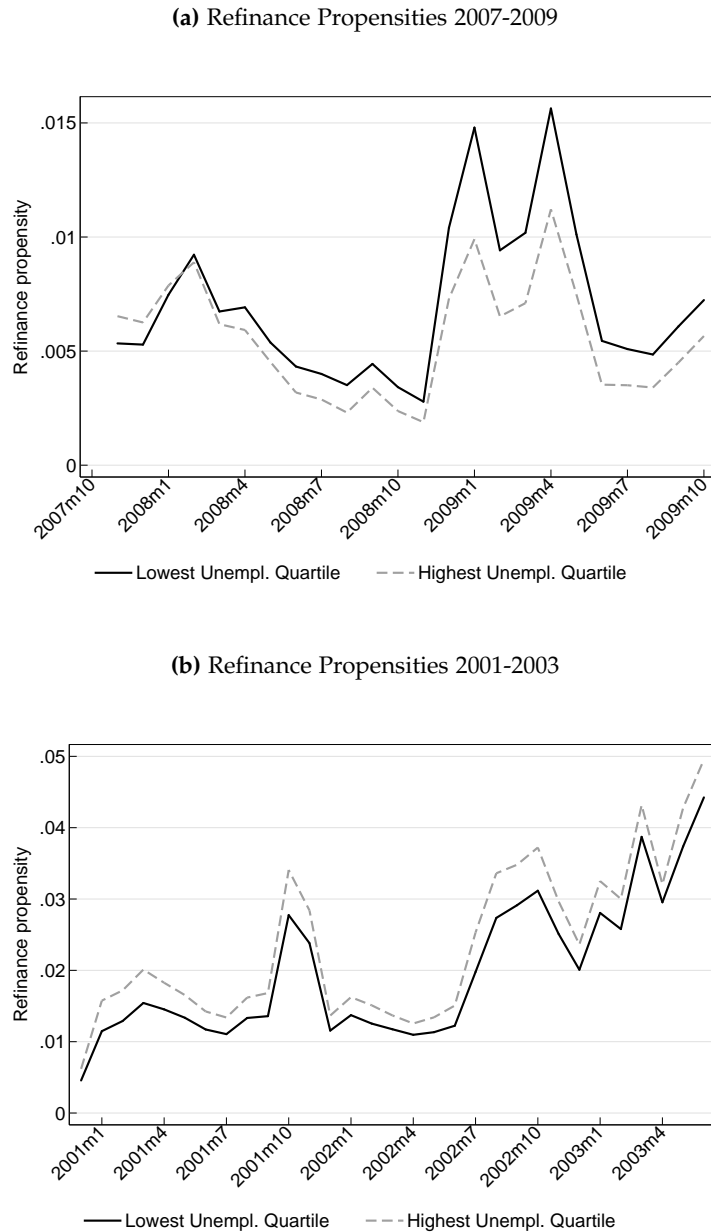


Figure shows monthly refinancing propensities for the top and bottom quartiles of MSAs with respect to their unemployment change over the period shown in the charts. Specifically, for each MSA we compute the change in unemployment between November 2007 and October 2009 (panel a) and December 2000 to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

tiles are quite different during the 2001 recession.²⁴ 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.²⁵ 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.

6 A Model of Regional Heterogeneity and Monetary Policy

In this section, we move to a theoretical analysis of the interaction between monetary policy and regional heterogeneity using a quantitative general equilibrium model of household refinancing. Our goal is both to clarify how the distribution of housing equity affects the transmission from interest rate cuts to aggregate spending through the refinancing channel, and highlight that because such distribution is time-varying, monetary policy makers may at times face a trade-off between stimulating aggregate spending and increasing regional consumption inequality.

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 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. 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 exogenous interest rate changes. We look at these responses both in the aggregate as well as across regions and ask how ex-ante features of regional heterogeneity, such as the regional distribution of housing equity, affect both aggregate consumption stimulus and ex-post regional consumption inequality in response to monetary policy.

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

²⁴To 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.

²⁵[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).

acts importantly with cross-region inequality even if it matters little for understanding within region inequality.²⁶ Furthermore, the regional dimension is relevant because labor markets and housing markets are geographically segmented and this induces a particular joint distribution of equity and income that is very different than at the individual level. For example, within a region, richer households could be buying more expensive houses. But for understanding the consequences of monetary policy for consumption inequality through the refinancing channel, it matters not that average income and house prices at the individual level are highly positively correlated. Instead, it matters that this correlation changes over time and that this change occurs mainly at the regional level, as opposed to the individual level, because of geographic segmentation. In addition, our previous empirical evidence on the heterogeneous response of refinancing and consumer spending to QE1 provides important discipline for our quantitative model. This evidence is inherently regional.

Finally, we decided against endogeneizing interest rates, and goods and house prices because it would make the theoretical insights less transparent, which conflicts with our main goal of clarifying how the regional distribution of equity affects the refinancing channel and the consequences monetary policy. Moreover, we conjecture that doing so would not change our qualitative conclusions and would actually amplifying them quantitatively. Section 7.4 presents a discussion of these issues.²⁷

6.1 Model Description

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

²⁶While there is heterogeneity in within-MSA house prices across zip codes, this is dwarfed by the variance in household level consumption within a region.

²⁷For example, treating 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, endogeneizing house prices would also amplify our conclusions as reductions in interest rates amplify regional house price inequality.

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 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 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 $F_{\tau}^{ij} q_{\tau}^j$.²⁸ 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 which will help us match certain features of the data.

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. We show in [Section 7.2](#) that relaxing this assumption has little effect on our conclusions. This implies that the new mortgage balance is $m_{\tau}^{ij} = \gamma q_{\tau}^j$ 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

²⁸We 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. See Appendix A.2. ■

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 enough.²⁹ 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

²⁹Alternatively, a sequence of positive house price shocks could also induce refinancing even with no drift.

we have a richer state space, with one state variable (assets a_t) evolving endogenously even when not refinancing. However, when solving the model numerically, we find that the refinancing decision is indeed a threshold level of equity, 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 $d' = \tilde{G}(\cdot)$,³⁰ 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.
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

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

³⁰The 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.

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.³¹ Finally, a region in this economy is defined as a collection of households that experience common shocks to house prices and income. Aggregates are simply obtained by adding up all regions.

6.2 Calibration Strategy

The model is calibrated annually and most parameters of the model are calibrated at standard values. As is standard in the risk-sharing literature, we set $\sigma = 2$ to generate an intertemporal elasticity of substitution of $1/2$. We assume the risk-free rate is constant at $r = 0.03$ and set $\beta = 0.93$. For simplicity we set $\beta_R = \frac{1}{1+r}$ so that lenders perfectly smooth consumption. For our baseline results, we assume the mortgage rate is constant at $r^m = 0.06$ and then consider the response to a one-time, unanticipated permanent decline to $r^m = 0.05$. We begin with this specification 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.³² The initial value of 0.06 and decline to 0.05 roughly matches 30-year fixed rates before and after quantitative easing began. In robustness results we instead assume that r^m follows an AR process with annual persistence of 0.88 and standard deviation of 0.0055 to match the behavior of 30-year rates from 1990-2015 and show that this specification delivers extremely similar quantitative conclusions. We set $\gamma = 0.8$, assuming that lenders require a minimum of 20% equity.³³ We set $\theta = 0.5$ in our baseline results following evidence from the NY Fed, but in our robustness results we show that this is not important for our 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.³⁴ The drift in income and house prices is assumed to be identical so that real house prices are constant on average.³⁵

³¹One can interpret the model as capturing a 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.

³²It is also consistent with the high persistence of the response of mortgage rates to QE1.

³³In Section 7.5 we explore the effects of changing this parameter.

³⁴Note 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.

³⁵Put 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.

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 QE1. We call such calibration our *baseline* economy. Then, our primary counterfactual analysis explores whether the same change in interest rates would have had different consequences if it had occurred during a recession with a different distribution of house prices and income.

More specifically, to construct our baseline economy, 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 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 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 for which we have strong empirical evidence both on the distribution of house price changes, income and refinancing responses from our event study around QE1. 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.³⁶

6.3 Theoretical Results in the 2008 Baseline Economy

Figure 8: Refinance Decision Follows a Threshold Policy

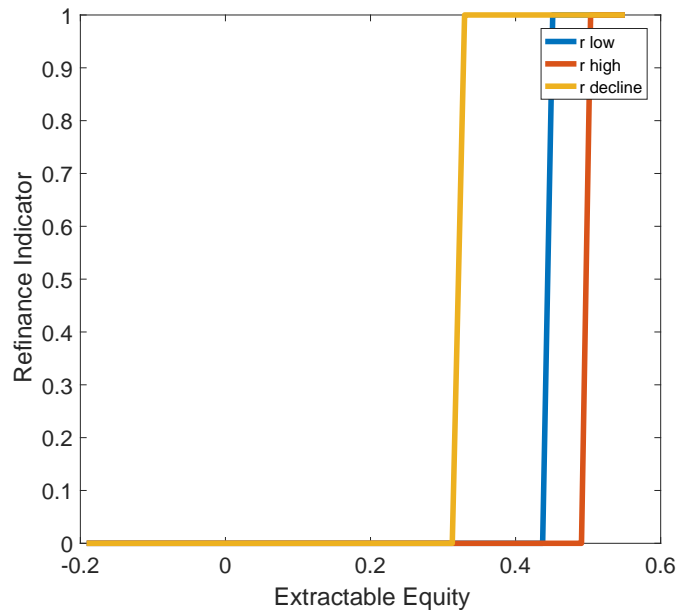


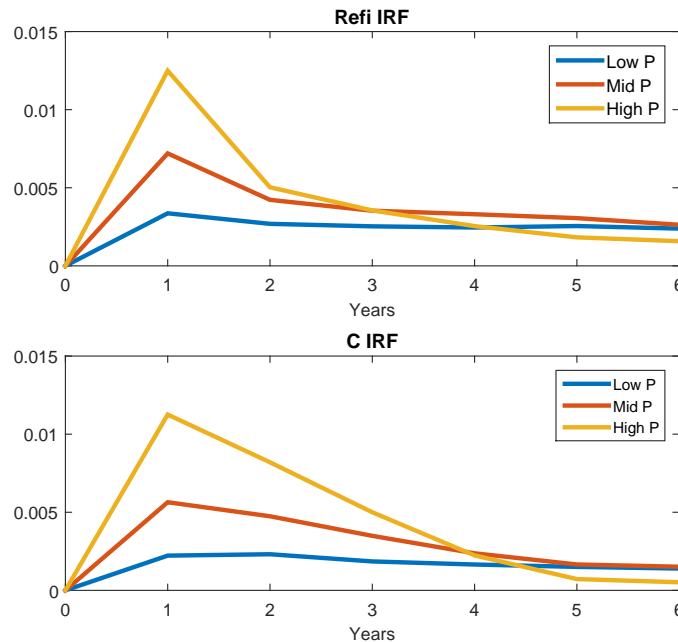
Figure 8 shows the threshold-like property of refinancing policies in a stationary environment

³⁶Nevertheless, alternative strategies for calibrating the fixed cost process produce broadly similar results.

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 A-4 in the Appendix 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 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.

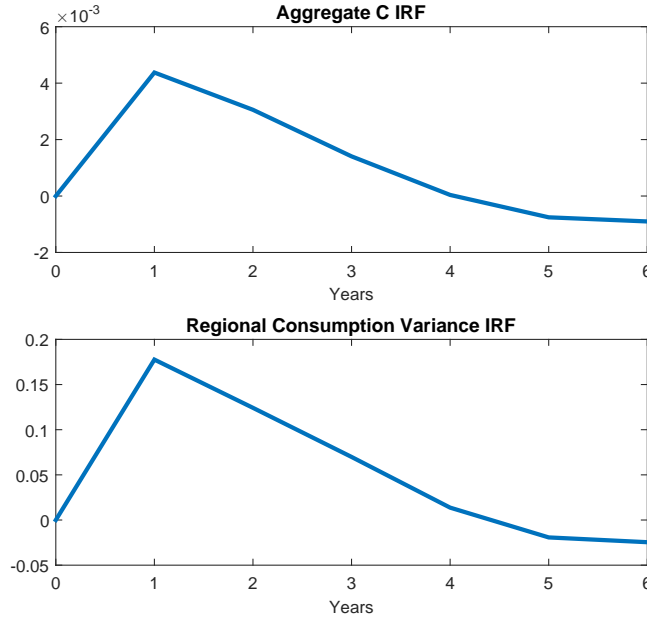
Figure 9: Response to interest decline by region



Note: 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.

Figure 9 shows the impulse response of each region to a decline in mortgage rates from 0.06

Figure 10: Aggregate Stimulus Vs. Regional Inequality



Note: 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.

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 3, 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 5.

We next compute aggregate consumption in the economy as well as the variance of log consumption across regions. Figure 10 then shows the response of aggregate consumption and regional consumption variance to the decline in mortgage rates. This shows our first important theoretical result: in our baseline economy, a reduction in interest rates increases aggregate consumption 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 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.

This tradeoff is strongly suggested by our empirical patterns that show a strong correla-

tion 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. Our quantitative model explicitly accounts for these lender effects and shows they do not cancel out. Aggregate consumption rises at least modestly when mortgage rates fall in the baseline economy.

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 a modest aggregate consumption but also an increase in inequality across regions.

6.4 Counterfactual Analysis

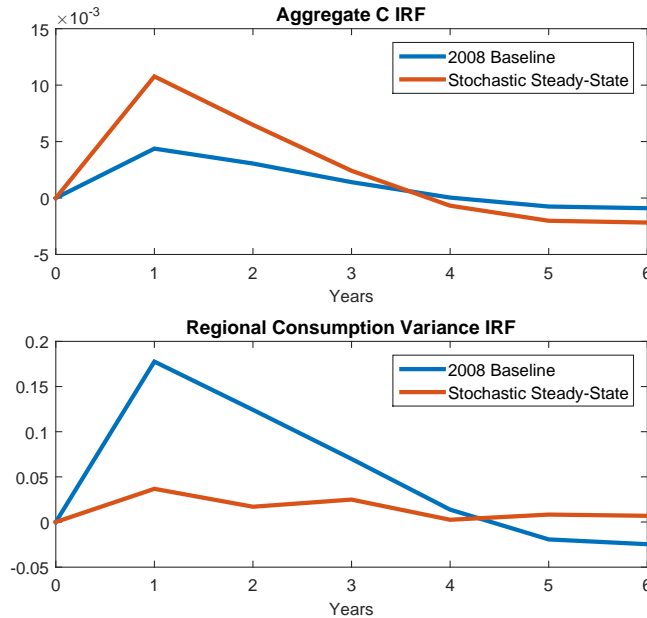
That our model reproduces the behavior of the economy just before and after QE also gives us some confidence in using it for more ambitious counterfactual analysis. We show that the theoretical effects of monetary policy are highly non-linear with respect to the regional distribution of housing equity in the economy. This means that accounting for time-variation in this distribution is crucial for correctly predicting the consequences of monetary policy because the effects of interest rate changes on both aggregate economic activity and regional inequality change as the level of housing equity, its variance and its correlation with income move across time. This counterfactual analysis is not of purely theoretical interest. In Section 5, we showed that such time-variation is present in the data.

6.4.1 2008 Economic Conditions v. Stochastic Steady State

Figure 11 compares the impulse response function of aggregate consumption and regional consumption inequality to an interest rate shock in the 2008 baseline to that which would have instead occurred in the stochastic steady-state. In the stochastic steady-state aggregate house prices do not decline, there is no increase in the variance of house prices, and there is no correlation between house prices and income.³⁷ Clearly, changing the distribution of equity dramatically changes the consequences of monetary policy: in the stochastic steady-state, the same decline in

³⁷In 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.

Figure 11: 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.

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.

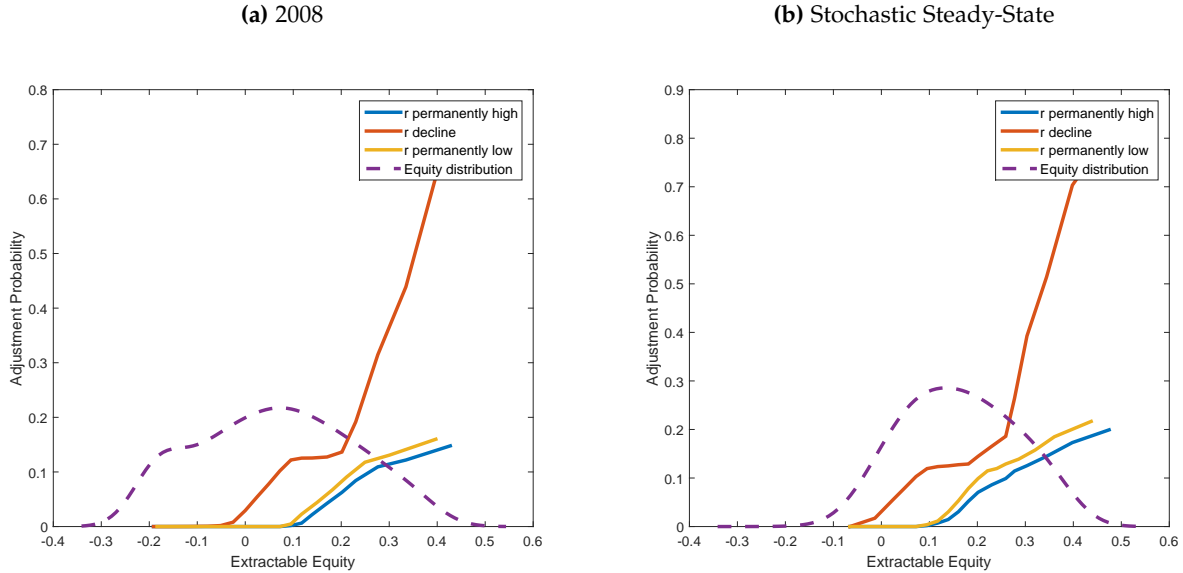
To understand why the distribution of equity in 2008 makes consumption less responsive to monetary policy than in the stochastic steady-state, it is useful to first explain why interest rate reductions lead to an increase in consumption. Figure 12 shows the distribution of equity in 2008 and the fraction 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 8 but averaging over the endogenous joint distribution of assets and income for each value of extractable equity.³⁸ Consistent with the threshold policy described above, households are much more likely to refinance when equity is large so that the adjustment 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 large.³⁹ This means that consumption increases

³⁸We assume that households have required equity of 20% but can extract equity over 20%.

³⁹On 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.

when there is an increase in the fraction of households refinancing.⁴⁰

Figure 12: Distribution of Equity and Refinancing Probability: 2008 calibration vs stochastic steady-state



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

When interest rates decline, households lower their refinancing threshold relative to the case when r^m is permanently low (i.e., the blue line shifts to the left, beyond the yellow line in Figure 12). 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.⁴¹ It is not the fraction of households who want to refinance and extract 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.⁴² Then, the refinancing response to a decline in rates is substantially muted in the

⁴⁰As 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.

⁴¹The change in the fraction of households refinancing with equity x is given by the vertical difference between the orange and blue lines, so the total change in refinancing is given by the difference between these lines integrated over the distribution of equity.

⁴²That 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

2008 calibration vis-a-vis the stochastic steady state because the lower average equity and higher variance in 2008 both shifts the distribution of equity (shown in purple) to the left and fans it out—thus, inducing fewer households in the region between the blue and orange lines to alter their refinancing decisions after the interest rate decline. In turn, the smaller refinancing response implies a smaller aggregate consumption response.

The intuition for the differential response of regional consumption inequality in 2008 vis-a-vis the stochastic steady-state is more subtle because the respective distributions of equity differ in several dimensions that result in off-setting effects of an interest decline. The next section discusses this in detail.

6.4.2 Housing Equity Statistics and the Consequences of Monetary Policy

The distribution of equity in 2008 is different in three ways from the stochastic steady-state. What are the independent roles of changes in the level of equity, the variance of equity and the correlation between income and equity for aggregate responses? Furthermore, why does a 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.

Panel (a) of Figure 13 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.⁴³ 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.

Panel (b) of Figure 13 shows the effects of changing the variance of equity across regions. An increase in the variance of equity amplifies the aggregate consumption response to an interest rate decline but also amplifies the response of consumption inequality. The intuition is almost identical to the effects of mean shifts and again follows immediately from Figures 12. Moving households with low equity to even lower equity has no effect on consumption responses to monetary policy, since these households do not refinance anyway. In contrast, additional equity amplifies the consumption response of those households on the right side of the distribution with substantial equity.

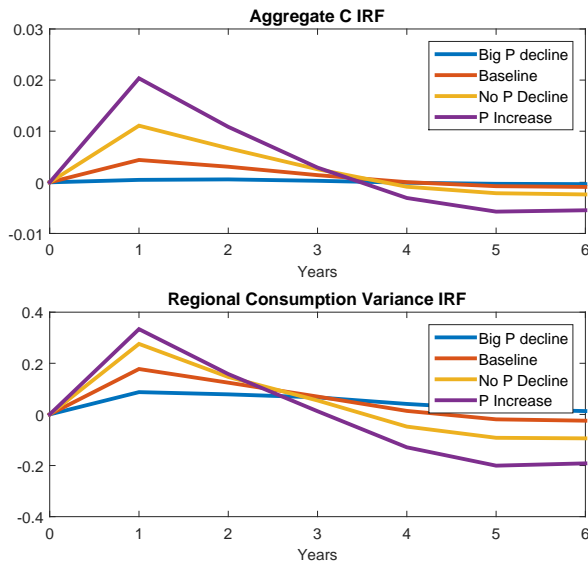
Finally, Panel (c) of Figure 13 shows the effects of changing the correlation between income and equity. If income and equity are uncorrelated then interest rate declines have almost no

and would have little effect on consumption.

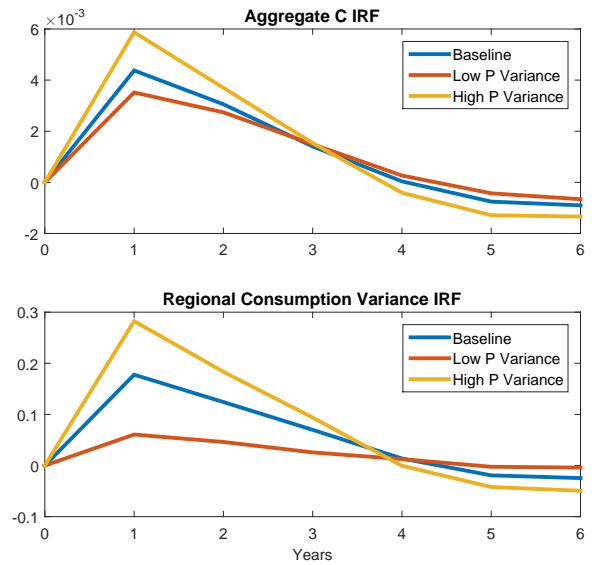
⁴³As is typical in fixed cost type models, these responses are non-linear. As the initial equity level becomes extremely large, the consumption response to an interest rate decline begins to fall towards zero. This is because if equity is extremely large, then households all refinance even without the decline in interest rates. However, these effects only become relevant when house prices increase by more than 30-40% in a single year.

Figure 13: Effects of changing equity distribution

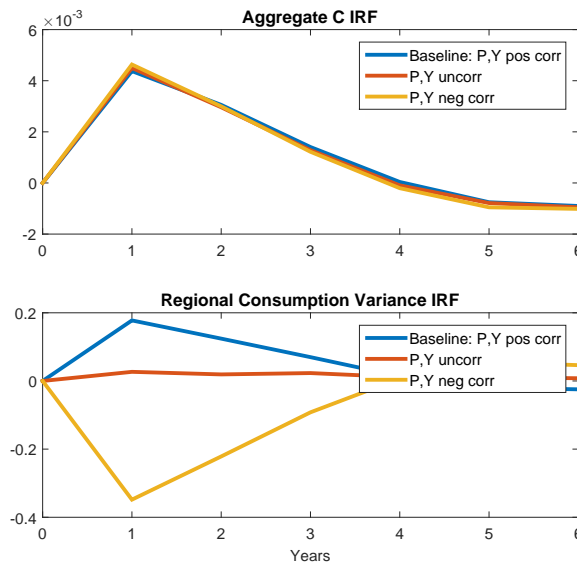
(a) Effects of changing average equity



(b) Effects of changing equity variance



(c) 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. The baseline economy includes a 12.5% aggregate house price decline. In Panel (a) 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. In Panel (b) The high variance calibration doubles the difference between high and low house price regions while the low variance calibration halves it. All economies feature the same baseline decline in house prices and correlation with income. In Panel (c) 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.

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 equity are correlated, this means that interest rate declines exacerbate the initial consumption inequality. If income and house prices are instead uncorrelated, as in the stochastic steady-state, then changes in consumption when interest rates fall are largely uncorrelated with initial levels of consumption, and if they are negatively correlated then inequality is reduced.

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

7 Robustness and Model Extensions

In this section, we consider a variety of extensions of our baseline results. First, our analysis thus far has assumed that regions are identical aside from the particular set of shocks which they experienced during the Great Recession. In reality, the locations with the largest house price declines also had the largest shares of adjustable rate mortgages, which might amplify their sensitivity to monetary policy. In addition, regions with the biggest declines in the Great Recession were those that experienced the largest preceding booms, which might affect our interpretation of inequality effects. We begin by showing that neither of these features overturns our previous results.

We then turn to a number of additional robustness checks exploring different assumptions on households' cash-out constraints, assumptions on the persistence of interest rates, the relationship between short and long-term rates and on the importance of general equilibrium effects. These features do not affect our conclusions qualitatively and actually amplify quantitatively many of our results.

Finally, our baseline results show that the regional distribution of house price growth in 2008 likely reduced the stimulative effects of monetary policy while also leading it to increase inequality. Hence, we explore how various complementary mortgage-related policies help mitigate these adverse effects while strengthening the effectiveness of monetary policy for aggregate stimulus. For example, we study policies that have received attention since the Great Recession, such as

HAMP and HARP.

7.1 Accounting for Additional Heterogeneity: ARM shares and Preceding booms

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.⁴⁴ 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.

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 A-5.⁴⁵ 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 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

⁴⁴See e.g. Auclert (2015).

⁴⁵Calza, Monacelli, and Stracca (2013) also show that ARMs are more prevalent in Spain and Italy than in Germany and France.

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 A-5 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 CLTV_80.⁴⁶ 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 CLTV_80 yields a coefficient of -0.10. Thus, overall, decreases in interest rates were still larger in MSAs with low CLTV_80.

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 A-5. 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. Panel (a) of Figure A-6 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.

Finally, we explore whether the boom-bust nature of the Great Recession substantially affects our conclusions. Since the regions which experienced the largest house price declines previously experienced the largest house price booms, perhaps the Great Recession itself dampens rather than amplifying inequality. This could in turn change the implications of monetary policy for inequality. In Panel (b) of Figure A-6 we repeat our baseline exercise but in a model where the house price bust is preceded by a housing boom.⁴⁷ This figure shows that our conclusions are unchanged. Monetary policy is slightly more effective and the trade-off with inequality is slightly reduced relative to the economy which experiences a bust with no boom, but the differences are

⁴⁶We 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).

⁴⁷In particular, we assume that in the period before the Great Recession, there is an aggregate house price increase of 10% and that the regions with a regional shock of $\pm 7.5\%$ which is perfectly negatively correlated with the shock during the bust. This roughly captures house price movements in the last year of the housing boom.

fairly small.⁴⁸

7.2 Refinancing without Full Cash-out

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 conclusions. In particular, Panel (a) of Figure A-7 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

7.3 Interest Rates

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 A-7 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 conclusions, we use the simpler one-time shock as our baseline experiment.

Our baseline model also assumes that monetary policy lowers the long-term mortgage rate r^m , but has no effect on the short-term rate r . This describes the behavior of rates during quantitative easing when short-term rates were at the zero lower bound. However, more conventional monetary policy would typically result in both rates falling simultaneously. Panel (c) of Figure A-7 shows results when we lower both r and r^m to maintain a constant spread. Again, results are essentially unchanged.

⁴⁸These differences arise because the fraction of households who are underwater during the bust is slightly reduced. Note also that we do not recalibrate parameters in this exercise to match the effects of quantitative easing, but they are quite similar to the baseline.

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.

7.4 General Equilibrium Effects

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 Panel (d) of Figure A-7, 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

7.5 Complementary Policies

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

We explore two forms of mortgage modification policies, which we refer to as “debt-forgiveness” and “relaxed-refinancing requirements.” While we intentionally implement these policies in a very stylized fashion in order to starkly illustrate their interactions with monetary policy, one can think of the first policy as capturing some elements of a modification program such as HAMP (although most loan modifications under HAMP focused on rate reductions and did not feature debt forgiveness) while the second is more similar to the HARP program.⁵² We model debt forgiveness by assuming that a portion of mortgage debt for any household who is underwater during the Great Recession is forgiven. In particular, all households with LTV greater than the maximum allowed for new loans have their loans reset to this maximum value.⁵³

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.

⁵²Existing research has shown that institutional features such as servicer participation and market power matter for the consequences of these policies as actually implemented (Agarwal et al., 2016, 2015). Also, especially with modification programs, which in practice usually focus on delinquent borrowers, moral hazard is an important concern that we do not consider. We are not attempting to provide an evaluation of specific mortgage program implementations in the Great Recession, and our focus is not on evaluating the efficacy or design of any such programs. We are instead interested in the broad ways such programs, independently implemented by the fiscal authority, might affect the consequences of monetary policy.

⁵³Note that we account for the negative effect of this policy on lenders. In reality, lenders would likely be directly compensated by the government, which in turn would raise taxes, but these taxes would likely be borne disproportionately by the richer lenders. Even if taxes were raised lump-sum from all households, borrowers whose debt was forgiven would still receive a net transfer.

Under the relaxed-refinancing requirements policy, we allow underwater households to refinance their mortgage rate without meeting the LTV requirement. To reflect the fact that the practical implementation of these policies explicitly eliminated various appraisal and other fees associated with refinancing, we also assume that under both policies underwater households are able to refinance their mortgages without paying the fixed refinancing cost.

Panel (a) of Figure A-8 shows the response of the economy to these mortgage modification programs, holding interest rates constant. That is, it shows the effects of the programs alone with no simultaneous monetary policy action. The debt-reduction program increases total consumption in the economy as it redistributes resources from less-constrained lenders to more constrained borrowers. It also reduces regional inequality since this debt forgiveness is only available to underwater households. In contrast, the relaxed refinancing requirements have no effect on the economy when interest rates are held constant. This is because if rates are unchanged, then even if underwater households can refinance costlessly, there is no distinction between refinancing and not refinancing for these households. Thus, this policy has no effect unless it is accompanied by a reduction in mortgage rates.

Panel (b) of Figure A-8 shows the response of the economy to *simultaneously* lowering interest rates and implementing these mortgage modification programs. That is, it shows the combined effects of these policies. For comparison, we also show the baseline economy with a decline in interest rates but no mortgage modification. Relative to the baseline economy, the combined policies lead to larger increases in output and smaller increases in inequality. When monetary policy is accompanied by our aggressive implementation of debt forgiveness, the inequality effects on impact are nearly eliminated, and inequality falls in later periods.

Since impulse responses in Panel (b) of Figure A-8 are computed relative to an economy with neither mortgage modification nor monetary policy, they tell us the combined effects of these policies. In contrast, in Panel (c), we isolate the impact of monetary policy from the direct effect of the mortgage programs. In particular, this figure computes the level of output and inequality with a decline in interest rates and mortgage modification relative to an economy with mortgage modification but no decline in rates.

Clearly, the presence of either mortgage modification program increases the effectiveness of monetary policy and reduces its effects of inequality. Interestingly, from the perspective of monetary policy, these two mortgage modification programs work nearly identically. Both policies increase the sensitivity of underwater households to changes in rates and so amplify the response of spending to interest rate declines. The fact that debt-forgiveness has larger effects than relaxed refinancing requirements in Panel (b) reflects the fact that this policy has direct effects on the economy independently of interest rate changes while the relaxed refinancing requirements policies only work through their interaction with interest rates. Since debt-forgiveness has both direct effects as well as interaction effects with monetary policy, the combined effects are ultimately larger. However, the effects of these policies on the efficacy of monetary policy are nearly identical.

Again, it is important to recognize the various caveats in footnote 52. Like any targeted fiscal policy, mortgage modification programs will induce winners and losers and the direct consequences of the policy will depend on the details of implementation. The point of the above exercise is not to evaluate the efficacy of mortgage modification per se but is instead to illustrate the channels through which such policies influence the monetary transmission mechanism.

The experiments thus far explore fiscal and monetary policy interaction in the mortgage market. In our final policy exercise, we explore how a simple form of macroprudential policy can interact with the refinancing channel of monetary policy. In particular, we consider two forms of risk regulation. In the first experiment, we simply lower the LTV cap from 0.8 to 0.7. In the second experiment, we implement a countercyclical LTV cap which is set to 0.7 during normal times but then rises to 0.9 during the Great Recession. That is, the central bank limits risk during good times but then relaxes constraints and increases liquidity in response to bad shocks.⁵⁴

The red line compared to the blue line in Figure A-9 shows the effect of a permanently tighter LTV cap. The first two panels show that this reduces both the aggregate ability of monetary policy to stimulate the economy as well as the increase in inequality generated by monetary policy. This is unsurprising, as a tighter borrowing constraint just reduces the extent to which the refinancing channel of monetary policy matters. The third panel of the figure shows that this risk reduction policy does reduce the depth of the recession, as reduced leverage means that consumption declines less when house prices fall. The more interesting results are shown in yellow, when the LTV cap is lowered ex-ante but is then reduced ex-post when there are large declines in aggregate house prices. This policy mildly increases the stimulative power of monetary policy and reduces its effects on inequality relative to the baseline (yellow vs. blue), and also reduces the depth of the recession. In this sense, the countercyclical LTV cap dominates either of the other policies: it reduces ex-ante risk which reduces the depth of the recession. But it also relaxes LTV requirements when and where they are most binding, which amplifies the effectiveness of monetary policy during recessions and reduces the trade-off between stimulus and inequality.

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

⁵⁴In the results shown here, we assume that this increase in the LTV cap during the Great Recession is completely unanticipated and lasts for a single year. We have solved a version of the model where households are aware of the countercyclical LTV policy ex-ante and recessions which trigger this LTV change occur with some small but non-zero probability and it delivered nearly identical results, so we present results for the simpler environment.

8 Conclusions

The Great Recession led to a prolonged period of monetary stimulus throughout much of the developed world. The effects of these policies are typically studied through the lens of representative agent New Keynesian models, which emphasize the importance of intertemporal substitution. In this paper, we explore a complementary channel of monetary transmission through collateralized lending and show that understanding this channel requires moving beyond a representative borrower. 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 led 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 activity during the 2001 easing cycle than during the Great Recession, and refinancing was actually more common in regions with high unemployment. Thus, the data 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 7.1 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, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. 2016. "Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program." *Journal of Political Economy* forthcoming.
- Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2015. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence From the Home Affordable Refinance Program." Working Paper 21512, 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.
- 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.
- 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.
- Galí, Jordi. 2009. *Monetary Policy, inflation, and the Business Cycle: An introduction to the new Keynesian Framework*. Princeton University Press.
- 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 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.
- Mian, Atif, Kamalesh 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.
- Rubio, Margarita. 2011. "Fixed-and Variable-Rate Mortgages, Business Cycles, and Monetary Policy." *Journal of Money, Credit and Banking* 43 (4):657–688.
- Stroebel, Johannes and John B. Taylor. 2012. "Estimated Impact of the Federal Reserve's Mortgage-Backed Securities Purchase Program." *International Journal of Central Banking* 8 (2):1–42.
- Tracy, Joseph and Joshua Wright. 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 Model Proofs and Description of Solution

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

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

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

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,

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

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

■

In order to solve the transformed value function, we discretize \tilde{x} using 40 grid points evenly spaced with width mu between 0.175 and -0.825, 64 grid points for \tilde{a} between 0 and 1, with more grid points near the lower asset values to account for concavity of the value function, and 46 grid points for \tilde{y} evenly spaced in logs between -0.5 and 0.5. The stochastic shock can take on 3-values: -1SD, 0, and +1SD with probabilities computed using the Tauchen algorithm. The model is then solved using value function iteration. Finally, the model is simulated using 50,000 households and 9 regions.

A.3 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.

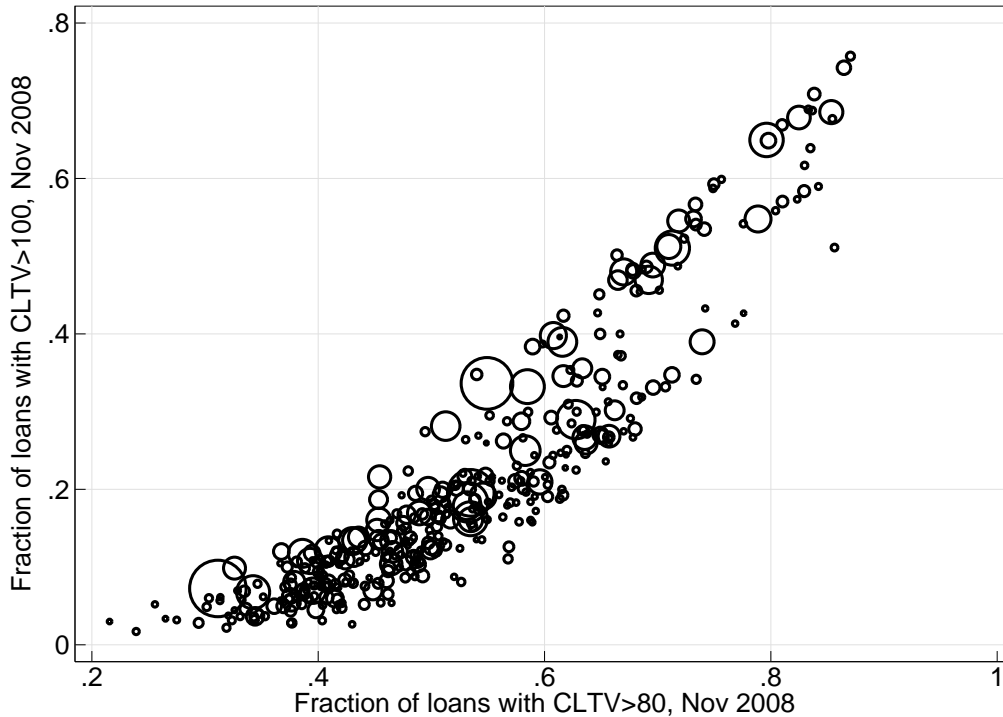


Figure A-2: Total Mortgage Origination Activity 2008-2009 in Top and Bottom Quartile of MSAs Defined by CLTV_80 in November 2008.

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 3. The vertical line indicates the month of the QE1 announcement (November 2008).

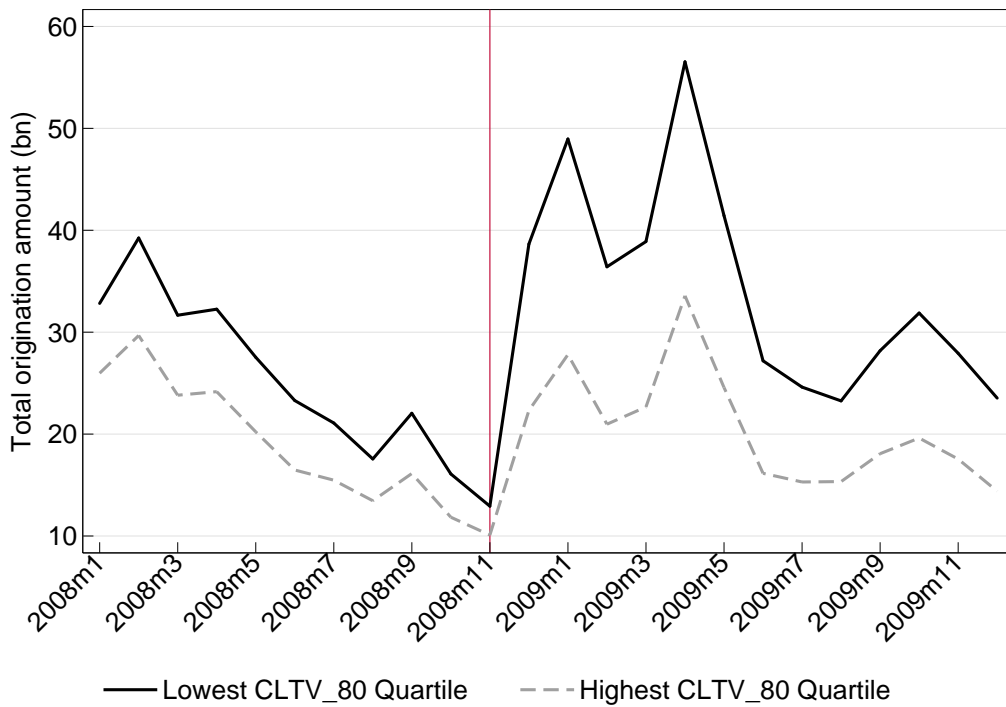


Figure A-3: Estimated Cash-out Amounts from Freddie Mac vs. in our CRISM data

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.

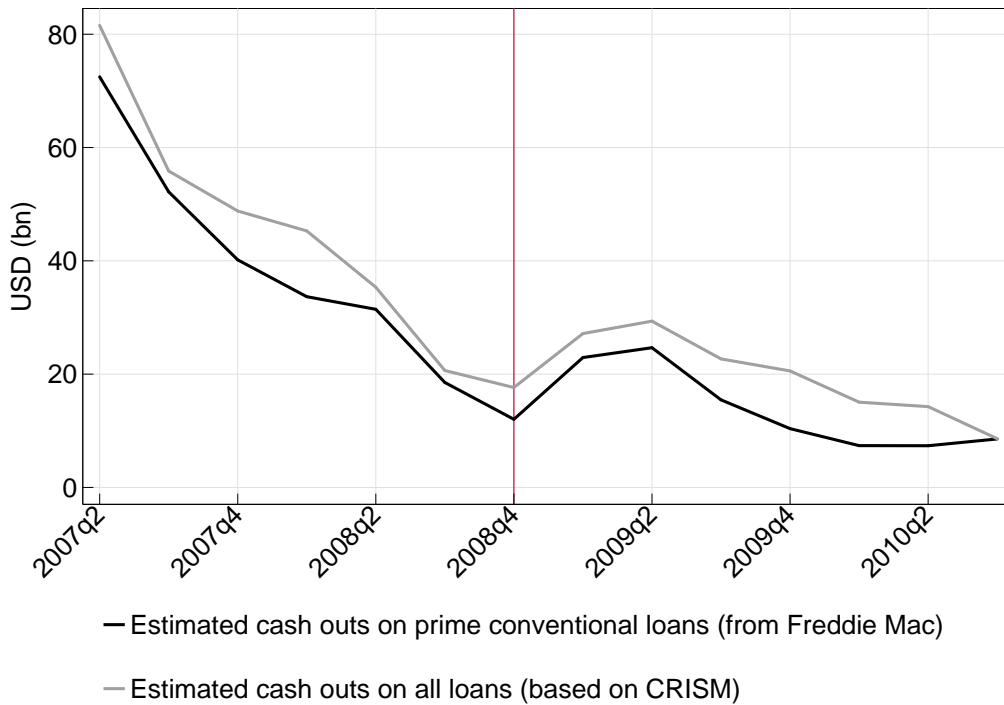


Figure A-4: Relationship between refinancing threshold, equity and income

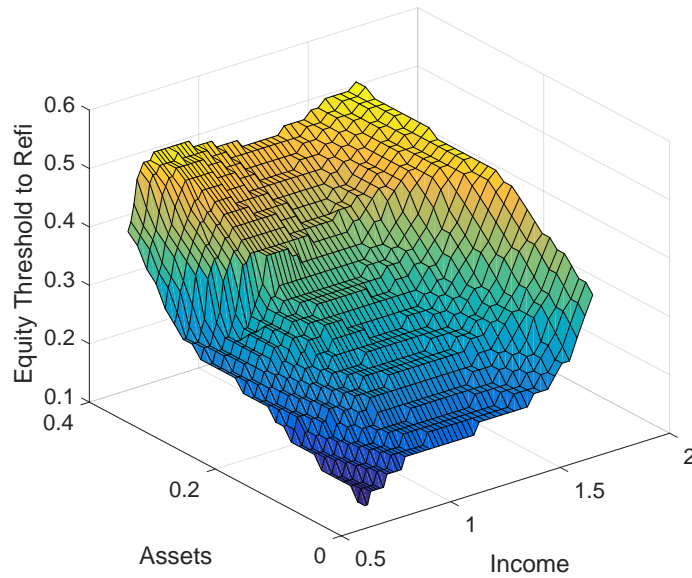


Figure A-5: Adjustable-Rate Mortgage Shares pre-QE1 and Payment Reductions over Nov 2008 - June 2009

(a) ARM Shares vs. CLTV₈₀

(b) Sizeable ARM Payment Reductions vs. CLTV₈₀

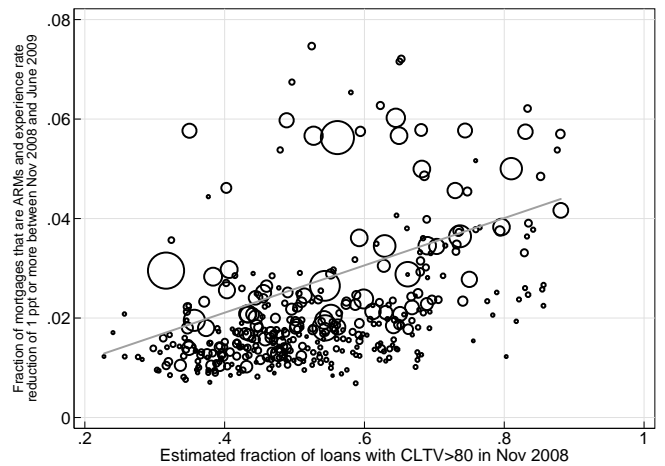
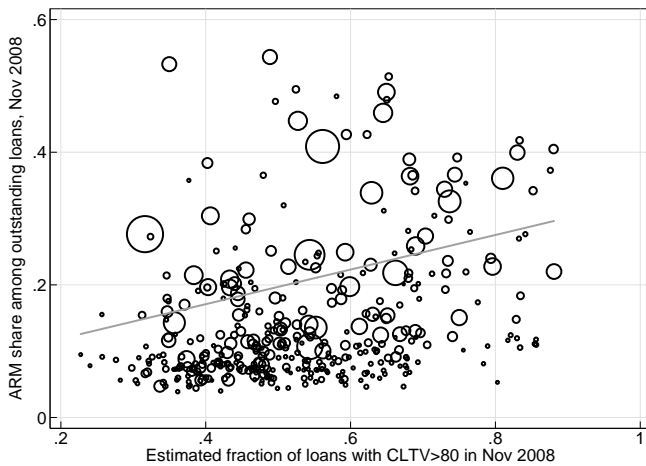
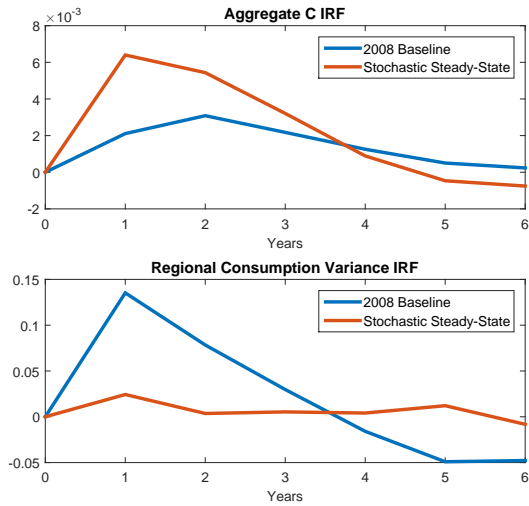


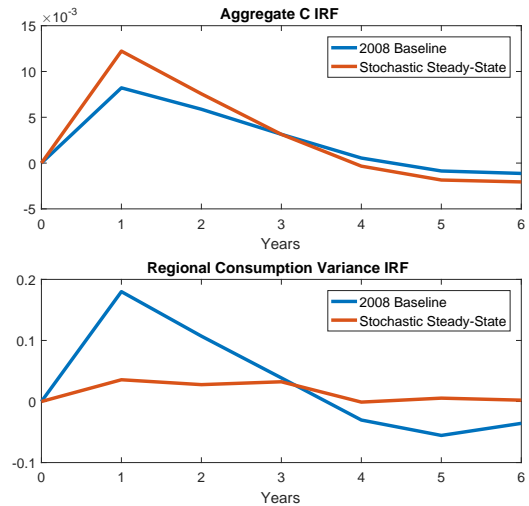
Figure shows scatter plots of the balance-weighted fraction of adjustable-rate mortgages (ARMs) in a given MSA (panel a) or the balance-weighted fraction that are ARMs and experience a rate reduction of 1 percentage point or more over November 2008 – June 2009, as measured in the CRISM data, versus the fraction of loans within the MSA with CLTV greater than 0.8 (CLTV₈₀) 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).

Figure A-6: Accounting for Ex-Ante Heterogeneity

(a) Match ARM-Share



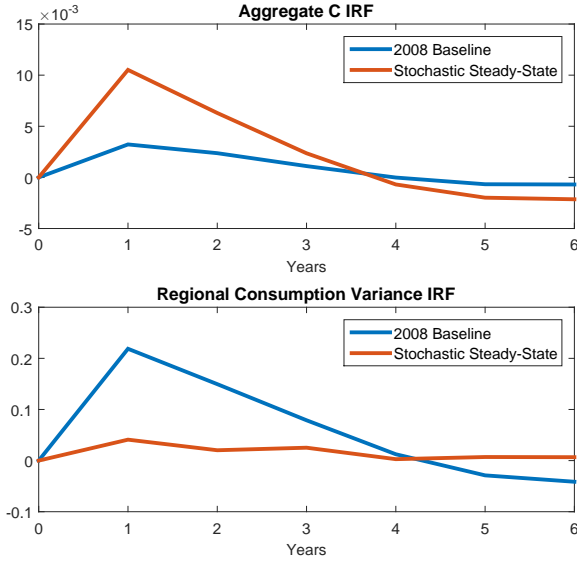
(b) Boom-Bust



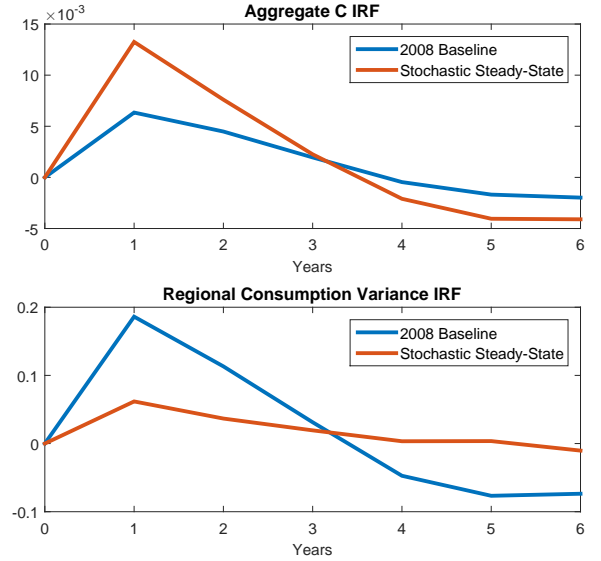
Panel (a) repeats Figure 11 in a model with both FRM and ARMs calibrated to match the regional differences with house prices in the data. Panel (b) repeats Figure 11 in a model where regions with the largest house price declines experience a prior house price boom.

Figure A-7: Robustness Results

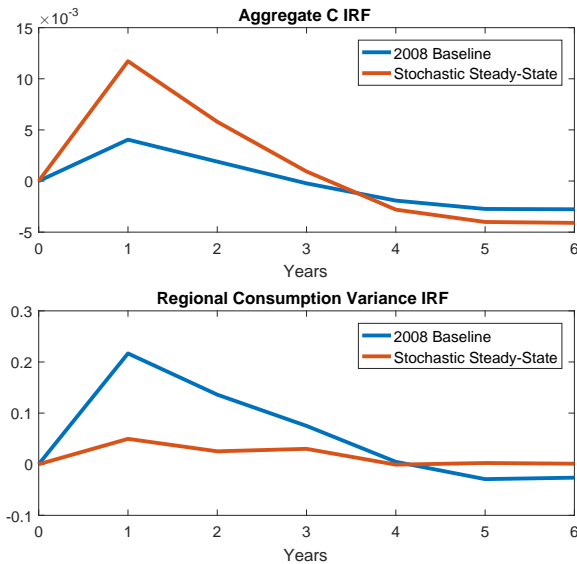
(a) Relaxing Full Cash-out Assumption



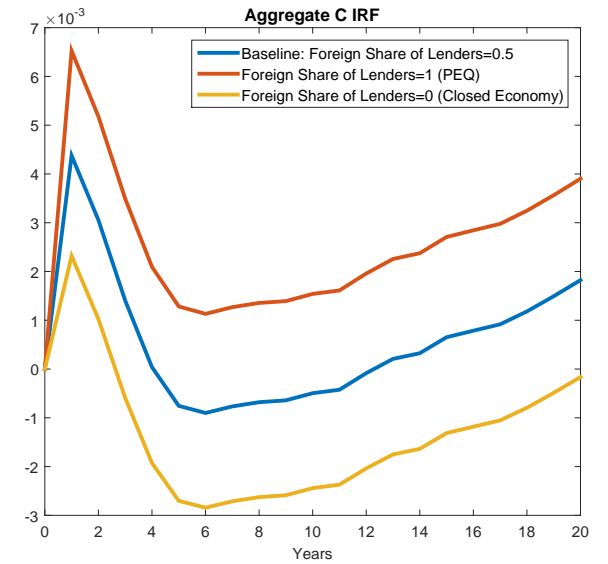
(b) Stochastic Mortgage Rate



(c) Reduce r and r^m



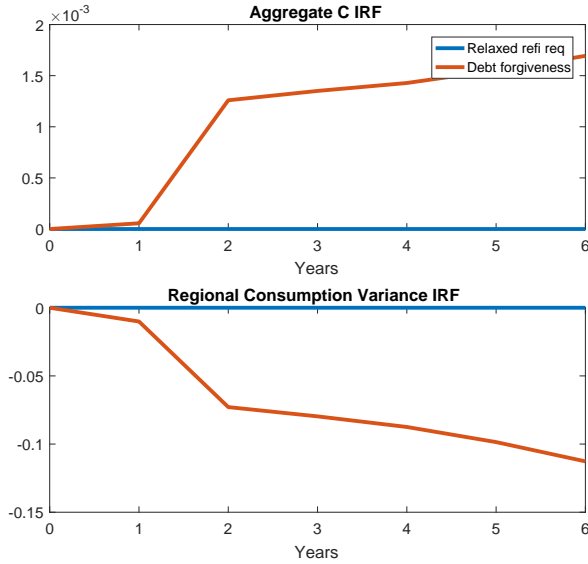
(d) Importance of GE Effects from Lenders' Consumption



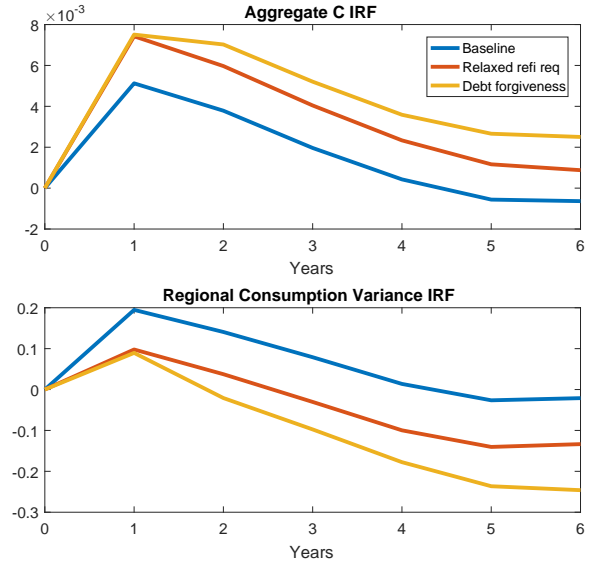
This figure repeats Figure 11 in two alternative models. In Panel (a) households can choose between full cash-out and no cash-out refinancing. In Panel (b) interest rate movements are stochastic instead of a one-time shock. In Panel (c) r and r^m both decline, with a constant spread. Panel (d) shows aggregate consumption IRFs for alternative assumptions on what share of mortgage debt is held by domestic households, whose income will fall when interest rates decline.

Figure A-8: Mortgage Modification Effects

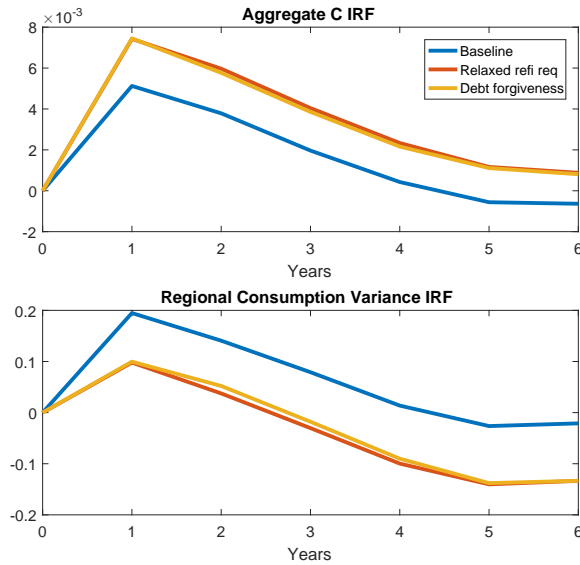
(a) Response to Mortgage Modification w/ constant r^m



(b) Response to Mortgage Modification + r^m Decline

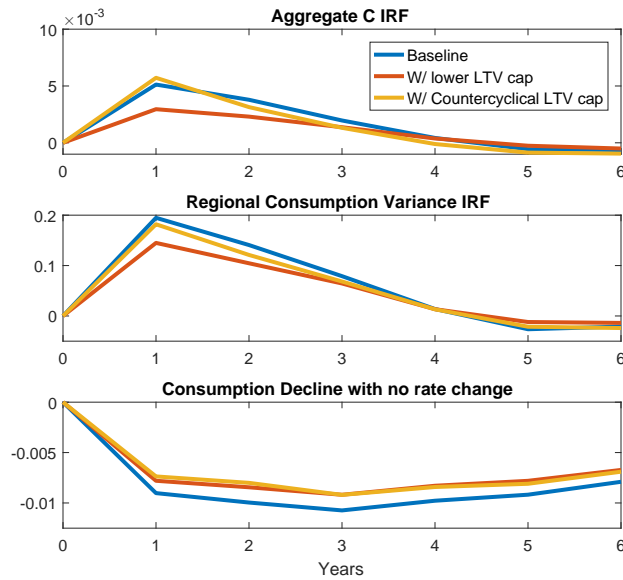


(c) Response to r^m decline, taking Mortgage Modification Programs as Given



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

Figure A-9: Effects of Countercyclical LTV Cap



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

A.3.1 MSA groups used in Figures 3, 4, and 5

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-Stuebenville, 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.

Online Appendix

A.4 Descriptive statistics prior to QE1

Table A-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

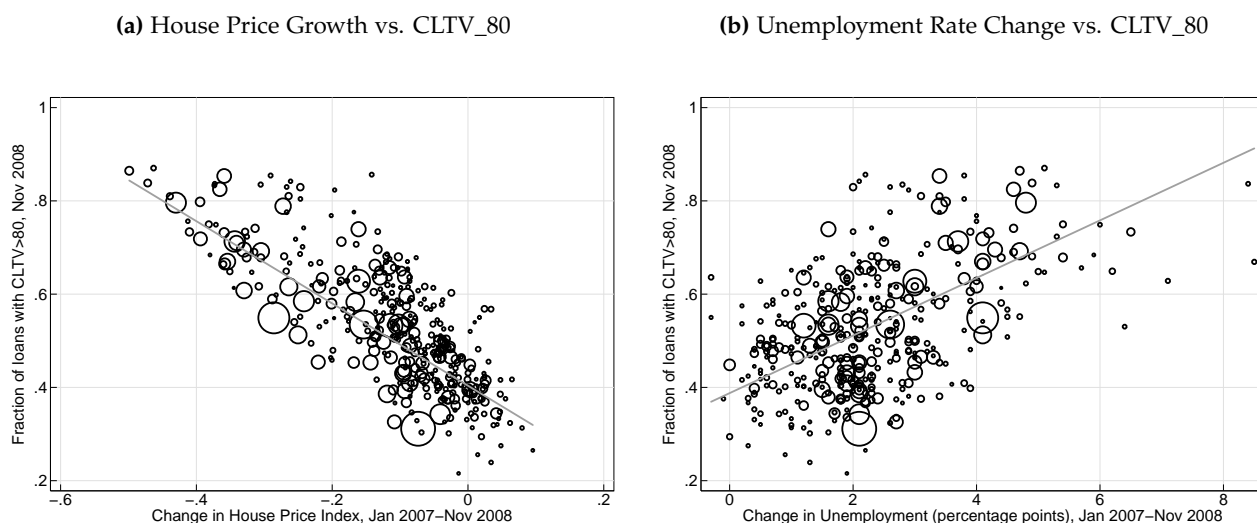
Table A-1 shows descriptive statistics for all 381 MSAs in our analysis sample, including the distribution of *CLTV_80* and *CLTV_100* in January 2007 and November 2008, and the distribution of house price and unemployment rate changes (in percent and percentage points, respectively) between January 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 2008.³ In January of 2007, essentially all mortgage owners had a CLTV below 1.0 and the vast majority 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 Panel (a) of Figure A-10 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). Panel (b) of Figure A-10 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.

³To 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.

Figure A-10: Relationship between CLTV and Other Measures of Economic Activity.



Panel (a) 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. Panel b) 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.

A.5 Mortgage refinance activity around QE1

Table A-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

Table A-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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

(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 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 A-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 3 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.

A.6 Cash-out Refinancing and Auto Spending around QE1

Table A-3 shows results from regressions using the change in equity removed (relative to outstanding balance) between November 2008 and February 2009 as the dependent variable and the same independent variables as in the regressions in Table A-3.

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. 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. In column 4, we 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 strongly significant. This implies that high $CLTV_{80}$ MSAs both refinanced less and removed less equity during the refinancing process conditional on refinancing.

Table A-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 results of a regression of the change in cash out amount relative to outstanding balance between Nov 2008 and Feb 2009 on the fraction of local borrowers with CLTV greater than 0.8 in Nov 2008. We measure cash out amounts using CRISM data and divide it by the previous month's mortgage balance outstanding (see Appendix for details). Column 2 adds 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 adds the change in local refinancing propensity between November 2008 and February 2009. 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 A-2). Robust standard errors in parentheses. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A-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
<i>N</i>	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 A-2) as an additional control. Column 4 includes the change in the refinancing propensity (dependent variable in Table A-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: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.