Mortgage Design in an Equilibrium Model of the Housing Market

Adam M. Guren*, Arvind Krishnamurthy†, and Timothy J. McQuade‡§

March 29, 2017
First Version: March 20, 2017

Abstract

How can mortgages be redesigned to reduce housing market volatility, consumption volatility, and default? How does mortgage design interact with monetary policy? We answer these questions using a quantitative equilibrium life cycle model with aggregate shocks, realistic and priced long-term mortgages, and a housing market that clears in equilibrium. We begin by comparing ARMs and FRMs to elucidate the core economic tradeoffs. ARMs provide hedging benefits in a crisis by reducing payments when income falls if the central bank lowers interest rates. This stimulates purchases by new homeowners, reduces default, and short circuits a price-default spiral, reducing price declines. The welfare benefits of ARMs in a crisis are large – equivalent to 12.5 percent of a year of consumption over a five year crisis – because ARMs particularly help young, high LTV households who face severe liquidity constraints. The overall benefits of ARMs also depend on the extent to which agents anticipate these hedging benefits and take on more risky debt positions in response. We evaluate several proposed mortgage designs that add state contingency to standard mortgages and find that an FRM that can costlessly be converted to an ARM has the best combination of insurance and macro-prudential benefits.

*Boston University, guren@bu.edu
†Stanford University Graduate School of Business and NBER, akris@stanford.edu
‡Stanford University Graduate School of Business tmcquade@stanford.edu
§The authors would like to thank Chaojun Wang and Xuiyi Song for excellent research assistance and seminar participants at SED, SITE, Kellogg, Queen's, Indiana, and Boston University and Andreas Fuster for useful comments. Guren acknowledges research support from the National Science Foundation under grant #1623801 and from the Boston University Center for Finance, Law, and Policy.

1 Introduction

The design of mortgages is crucial to both household welfare and the macroeconomy. Houses make up a majority of wealth for most households, and mortgages tend to be their dominant source of credit, so the design of mortgages has an outsized effect on household balance sheets (Campbell, 2013). Recent research has shown that these balance sheet effects can dramatically alter households' marginal propensities to consume because mortgages make home equity illiquid (*e.g.*, Kaplan and Violante, 2014). Additionally, in the mid-2000s boom and subsequent bust, housing wealth extraction through the mortgage market boosted consumption in the boom and reduced consumption in the bust (*e.g.*, Mian and Sufi, 2011). Mortgage debt also led to the wave of foreclosures that led to over six million households losing their homes, badly damaging household balance sheets and crippling the housing market (*e.g.*, Guren and McQuade, 2015; Mian et al., 2015). Finally, in the wake of the recession, there has been increased attention paid to the role that mortgages play in the transmission of monetary policy to the real economy through household balance sheets (*e.g.*, Auclert, 2016; Wong, 2015; Di Maggio et al., 2017).

Simple fixed rate amortizing mortgages with no principal indexation, the most common design in the US, result in a pattern of payments that are suboptimal relative to a complete markets Arrow-Debreu benchmark. A number of authors have recently suggested that improved mortgage designs could have blunted the foreclosure crisis (*e.g.*, Caplin, Chan, Freeman and Tracy, 1997; Shiller, 2008; Guren and McQuade, 2015; Mian et al., 2015; Corbae and Quintin, 2015; Eberly and Krishnamurthy, 2014). In this paper, we quantitatively study the link between mortgage design, household choices, monetary policy, and aggregate outcomes through the lens of a heterogenous agents macro model in which loans are priced and the housing market is in equilibrium. We set aside the question of what costs lead to incomplete mortgage contracts or what the optimal fully-state contingent contract would look like and instead quantify the welfare benefits of "simple" and plausibly implementable mortgage designs in a realistic model with equilibrium feedbacks.[1]

Our model features overlapping generations of households subject to both idiosyncratic and aggregate shocks, making endogenous decisions over home purchases, borrowing, consumption, refinancing, and default. We consider different exogenous processes for the interest rate, reflecting different monetary policies. Competitive and risk-neutral lenders set spreads for each mortgage to break even in equilibrium, so lenders charge higher interest rates when a mortgage design hurts their bottom line. Equilibrium in the housing market implies that household decisions, mortgage spreads, and the interest rate process influence the equilibrium home price process. Household expectations regarding equilibrium prices and mortgage rates feed back into household decisions, and we solve this fixed-point problem using using computational methods based on Krusell and Smith (1998).

A key aspect of our analysis is that mortgage design affects household default decisions and hence home prices, which in equilibrium feeds back to household indebtedness. The quantitative implications of our model depend on accurately representing the link between home prices and

[1] We focus on mortgage designs that add state contingency to payments without changing the horizon over which the mortgage amortizes to maintain tractability. Assessing mortgage designs where the amortization schedule is state contingent is left to future research.

default. Consequently, we calibrate our model to match quasi-experimental micro evidence on the effect of LTV and payment size on default from Fuster and Willen (2015). Simulating quasi-experiments in our calibration procedure is an innovation that ensures that our model accurately captures the effects of changes in LTVs and interest rates as we alter mortgage design. We also match standard moments and the empirical distributions of mortgage debt and assets.

The calibrated model provides a laboratory to assess the benefits and costs of different mortgage designs. Our primary application is to a housing crisis, although we also consider the performance of different mortgages in stochastic simulations of "normal times."

To develop intuition, we begin by comparing a world with fixed-rate mortgages to a counterfactual world with all adjustable-rate mortgages. We find that ARMs provide important insurance that ameliorates the welfare impact of a housing crisis over five years by 12.5 percent of one year of consumption. To understand the mechanism, note that as the central bank lowers the short rate in response to the crisis, ARM rates fall dramatically, while FRM rates fall to a lesser extent because they are priced off the long end of the yield curve. Moreover, in a world with fixed-rate mortgages, homeowners with low equity cannot take refinance to take advantage of lower rates due to a minimum LTV constraint. Because the probability of a large and persistent negative income shock also rises in the recession, a fraction of these homeowners become liquidity constrained and default. These defaults increase the supply of homes on the market, further pushing down prices, which in turn leads to more default and prevents more homeowners from refinancing. This phenomenon generates a price-default spiral, amplifying the crisis through equilibrium feedbacks. Conversely, in a world with adjustable-rate mortgages, homeowners do not need to refinance to take advantage of lower interest rates: since the mortgage payment is pegged to the prevailing short rate in the market, payments fall automatically. This leads to less default by underwater homeowners, short-circuiting the default spiral and leading to a less severe housing crisis. Furthermore, consumption falls by less since the decrease in mortgage payments offers a hedge against declining labor income. Finally, since ARM rates fall more than FRM rates, demand for housing by new home buyers rises more under ARMs, further limiting price declines.

Under FRMs, the decline in consumption and defaults are concentrated on a small portion of the population with high mortgage debt and low assets. These constrained households have to cut back on their consumption dramatically to avoid default or end up defaulting, leading to acute welfare losses. Because these households have high marginal utilities of consumption, the welfare benefits of ARMs in a crisis can be large even if the aggregate consumption benefit of ARMs is modest.

There is, however, a countervailing force. Homeowners understand that with an ARM, their payments will fall when the economy enters into a recession. These hedging benefits encourage home purchasers to take on more leverage pre-crisis relative to an all FRM world. While the ARM offers *ex-post* benefits during a housing crisis by lowering mortgage payments and reducing default, *ex-ante* this mortgage design creates greater fragility by increasing household leverage. On net, we find the *ex-post* benefits are stronger, but the benefits of ARMs are substantially reduced by *ex-ante* behavior.

Having established the key economic tradeoffs by studying FRMs and ARMs, we turn to simulating three mortgages proposed in the wake of the crisis. We find that adding state contingency to a standard FRM contract can improve welfare. An FRM mortgage that can costlessly be converted to an ARM in a crisis as proposed by Eberly and Krishnamurthy (2014) provides the best balance of macroprudential and insurance benefits in a crisis in which the central bank reduces interest rates while still limiting downside risk if the central bank raises rates in a recession, for instance to fight inflation. An FRM that can be refinanced under water as proposed by Campbell (2013) provides similar macroprudential benefits but does not do as well at smoothing consumption since it is priced off the long end of the yield curve. Option ARMs, which Piskorski and Tchistyi (2010) argue are optimal in normal times, provide the most insurance but do poorly macroprudentially in a crisis. In future drafts we plan to evaluate more mortgage designs, in particular shared appreciation mortgages, and do more to assess the performance of mortgage designs under different monetary policy responses.

We also investigate how mortgage design interacts with monetary policy. We find that the impact of a more aggressive monetary policy in a crisis depends critically on the extent to which changes in short rates are passed through into long rates and the extent to which the policy is anticipated. If changes in short rates induced by monetary policy barely affect long-term interest rates and in turn FRM rates as is the case under the expectations hypothesis, monetary policy has very little impact on the severity of a housing crisis in an FRM economy because payments are effectively unchanged. On the other hand, with ARMs, lower short rates lead can reduce mortgage payments in bad states, generating less default and a smaller price-default spiral. By contrast, if monetary policy is able to substantially reduce long rates through policies like quantitative easing, it can be effective in ameliorating the effects of a housing crisis in an FRM economy. Rather than lowering the payments of underwater homeowners, though, this policy's main benefit arises through providing new homeowners with cheap financing, which stimulates demand for housing, providing support for house prices and ameliorating the price-default spiral. Regardless of design, the benefits of aggressive monetary monetary policy in a crisis depend on the extent to which households anticipate the policy and take on more mortgage debt in response to the expectation that a crisis will be ameliorated by the monetary authority, creating macro fragility.

The remainder of the paper is structured as follows. Section 2 describes the relationship to the exiting literature. Section 3 presents our model, and Section 4 describes our calibration procedure. Section 5 compares the performance of ARM-only and FRM-only economies to develop economic intuition. Section 6 compares proposed mortgages that add state contingency to an FRM, and Section 7 considers the interaction of mortgage design with monetary policy. Section 8 concludes.

2 Related Literature

This paper is most closely related to papers that analyze the role of mortgages in the macroeconomy through the lens of a heterogeneous agents model. In several such papers, house prices are exogenous. Campbell and Cocco (2015) develop a life-cycle model in which households can borrow using long-term fixed- or adjustable-rate mortgages and face income, house price, inflation, and interest rate risk. They use their framework to study mortgage choice and the decision to default.

In their model, households can choose to pay down their mortgage, refinance, move, or default. Mortgage premia are determined in equilibrium through a lender zero-profit condition. Our modeling of households shares many structural features with this paper, but while they take house prices as an exogenous process, we crucially allow for aggregate shocks and determine equilibrium house prices. This critical feature of our model allows us to study the interaction of mortgage design with endogenous price-default spirals. A prior paper, Campbell and Cocco (2003), use a more rudimentary model without default and with exogenous prices to compare ARMs and FRMs and assess which households benefit most from each design. Similarly, Corbae and Quintin (2015) present a heterogeneous agents model in which mortgages are priced in equilibrium and households select from a set of mortgages with different payment-to-income requirements, but again take house prices as exogenous. They use their model to study the role of leverage in triggering the foreclosure crisis, placing particular emphasis on the differential wealth levels and default propensities of households that enter the housing market when lending standards are relaxed. Conversely, we focus on the impact of mortgage design and monetary policy on housing downturns, allowing for endogenous house price responses.

Other heterogeneous agent models of the housing market have endogenous house prices but lack aggregate shocks or rich mortgage designs. Kung (2015) develops a heterogeneous agents model of the housing market in which house prices are determined in equilibrium. His model, however, lacks aggregate shocks and household saving decisions. He focuses specifically on the equilibrium effects of the disappearance of non-agency mortgages during the crisis. By contrast, we include aggregate shocks and a rich set of household decisions that Kung assumes away. We also study a variety of mortgage designs and analyze how mortgage design interacts with monetary policy. Finally, Kaplan et al. (2016) present a life-cycle model with default, refinancing, and moving in the presence of idiosyncratic and aggregate shocks in which house prices are determined in equilibrium. Their focus, however, is on explaining what types of shocks can explain the dynamics of house prices and consumption in the Great Recession. They simplify many features of the mortgage contract for tractability in order to focus on these issues, while our paper allows for rich variety in mortgage types as well as different monetary policy rules.

Our paper also builds on a largely theoretical literature studying optimal mortgage design. Piskorski and Tchistyi (2010; 2011) consider optimal mortgage design from an optimal contracting perspective, finding that the optimal mortgage looks like an option ARM when interest rates are stochastic and a subprime loan when house prices are stochastic. Brueckner and Lee (2017) focus on optimal risk sharing in the mortgage market. These papers identify important trade-offs inherent in optimal mortgage design in a partial equilibrium settings. Our paper is also related to a literature advocating certain macroprudential polices design to ameliorate the severity of housing crises. Mian and Sufi (2015) advocate for modifications through principal reduction, while Eberly and Krishnamurthy (2014) advocate for monthly payment reductions. Greenwald (2015) advocates for payment-to-income constraints as macroprudential policy to reduce house price volatility.

To calibrate our model, we draw on a set of papers which document empirical facts regarding household leverage and default behavior. Foote et al. (2008) provide evidence "double trigger" theory of mortgage default, whereby most default is accounted for by a combination of negative equity and an income shock as is the case in our model. Bhutta et al. (2010), Elul et al. (2010), and Gerardi et al. (2013) provide further support for illiquidity as the driving source of household default. Fuster and Willen (2015) show that reducing mortgage payments can significantly reduce default. Di Maggio et al. (2017) show that downward rate resets lead to increases in household consumption. This micro evidence motivates our focus on mortgage designs with state-contingent payments, and we calibrate to Fuster and Willen's evidence.

Finally, our research studies how mortgage design interacts with monetary policy and thus relates to a literature examining the transmission of monetary policy through the housing market. Caplin, Freeman, and Tracy (1997) posit that in depressed housing markets where many borrowers owe more than their house is worth, monetary policy is less potent because individuals cannot refinance. Beraja, Fuster, Hurst, and Vavra (2015) provide empirical evidence for this hypothesis by analyzing the impact of monetary policy during the Great Recession. Relatedly, a set of papers have argued that adjustable-rate mortgages allow for stronger transmission of monetary policy since rate changes directly affect household balance sheets (Calza et al., 2013; Garriga et al., 2013; Auclert, 2016). Di Maggio et al. (2017) show empirically that the pass-through of monetary policy to consumption is stronger in regions with more adjustable rate mortgages. Finally, Wong (2015) highlights the role that refinancing by young households plays in the transmission of monetary policy to consumption.

3 Model

This section presents an equilibrium model of the housing market with rich mortgage contracts that we subsequently use as a laboratory to study different mortgage design. Home prices and mortgage spreads are set in equilibrium. Short-term interest rates, on the other hand, are exogenous to the model and depend on an aggregate shock process. We are interested in understanding how the relationship between interest rates and the state of the economy affect the equilibrium. For ease of exposition, we present the model for the case of an FRM, but consider other designs when presenting our quantitative results.

3.1 Setup

Time is discrete and indexed by t. The economy consists of a unit mass of overlapping generations of heterogeneous households of age $a = 1, 2, ..., T$ who make consumption, housing, borrowing, refinancing and default decisions over their lifetime. Household decisions depend both on aggregate state variables Σ_t and agent-specific state variables s_t^j, where j indexes agents. Unless otherwise stated, all variables are agent-specific, and to simplify notation we suppress their dependency on s_t^j.

The driving shock process in the economy is Θ_t, which is part of Σ_t. Θ_t follows a discrete Markov process over three states $\Theta_t \in \{\text{Crisis,Recession,Expansion}\}$ and is governed by a transition matrix Ξ^{Θ}.

Each generation lives for T periods. At the beginning of a period, a new generation is born and shocks are realized. Agents then make decisions, and the housing market clears. Utility is realized and the final generation dies at the end of the period. We make a timing assumption that households enter period t with a state s_t^j and choose next period's state variables s_{t+1}^i in period t given the period t housing price p_t. Utility is based on period t actions. However, agents who take out a new loan start receiving the interest rate prevailing at time t immediately.

Households receive flow utility from housing H_t and non-durable consumption C_t:[2]

$$U\left(C_t, H_t\right) = \frac{C_t^{1-\gamma}}{1-\gamma} + \alpha_a H_t.$$

In the last period of life, age T, a household with terminal wealth b receives utility:

$$\frac{\left(C_t\right)^{1-\gamma}}{1-\gamma} + \alpha_T H_t + \psi \frac{\left(b+\xi\right)^{1-\gamma}}{1-\gamma}.$$

For simplicity, we assume that households use their wealth to finance housing and end-of-life care after their terminal period. Consequently, the wealth b is not distributed to incoming generations, who begin life with no assets.[3]

Households receive an exogenous income stream Y_t:

$$Y_t \equiv \exp\left(y_t^{agg}\left(\Theta_t\right) + y_t^{id}\right).$$

Log income is the sum of an aggregate component that is common across households and a household-specific idiosyncratic component. The aggregate component y_t^{agg} is a function of Θ_t. The idiosyncratic component y_t^{id} is a discrete Markov process over a set $\left\{Y_t^{id}\right\}$ with transition matrix $\Xi^{id}\left(\Theta_t\right)$.

Households retire at age $R < T$. After retirement, households no longer face idiosyncratic income risk and keep the same idiosyncratic income they had at age R, reduced by ρ log points to account for the decline in income in retirement. This can be thought of as a social security benefit that conditions on terminal income rather than average life income for computational tractability as in Guvenen and Smith (2014).

There is a progressive tax system so that individuals' net-of-tax income is $Y_t - \tau\left(Y_t\right)$. The tax system is modeled as in Heathcote et al. (2017) so that:

$$\tau\left(Y_t\right) = Y_t - \tau_0 Y_t^{1-\tau_1}.$$

Houses in the model are of one size, and agents can either own a house ($H_t = 1$) or rent a house ($H_t = 0$). Buying a house at time t costs p_t, and owners must pay a per-period maintenance cost of mp_t. With probability ζ, homeowners experience a life event that makes them lose their match with their house and list it for sale, while with probability $1-\zeta$, owners are able to remain in their

[2]The term α_a describes the utility from homeownership as a function of age. In our calibration, we will assume that α_a is decreasing in age so as to reflect the fact that at older ages the homeownership rate declines slightly.

[3]Including terminal wealth in the utility function is standard in OLG models of the housing market because otherwise households would consume their housing wealth before death. However, in the data the elderly have substantial housing wealth which they do not consume. The functional form for the utility derived from terminal wealth is standard.

house.

The rental housing stock is entirely separate from the owner-occupied housing stock. Rental housing can be produced and destroyed at a variable cost q, so in equilibrium renting costs q per period. Although this assumption is stark, it is meant to capture that while there is some limited conversion of owner-occupied homes to rental homes and vice-versa in practice, the rental and owner-occupied markets are quite segmented (Glaeser and Gyourko, 2009(Glaeser and Gyourko 2009)). This implies that most movements in house prices are accompanied by movements in the price-to-rent ratio. Indeed, in the data, the price-to-rent ratio has been nearly as volatile as price, and the recent boom-bust was almost entirely a movement in the price to rent ratio. Our modeling of the rental market also implies that changes in credit conditions will affect aggregate demand for housing as potential buyers enter or exit the housing market, in contrast to models with substantial conversion between renting and owning such as Kaplan et al. (2016).

A household's date t mortgage balance is $M_t \geq 0$ and carries interest rate i_t. Mortgage interest is tax deductible, so that taxes are $\tau (Y_t - i_t M_t)$. In order to economize on state variables, the mortgage amortizes over its remaining life as in Campbell and Cocco (2003, 2015). This rules out mortgage designs with variable term lengths, but still allows for the analysis of mortgage designs that rely on state-dependent payments. The minimum payment on a mortgage for an agent who does not move or refinance at time t is:

$$M_t \frac{\left(i_t (1 + i_t)^{T-a+1}\right)}{(1 + i_t)^{T-a+1} - 1}.$$

The interest rate on the mortgage at origination is $i_t = i_t^{FRM} (\Theta_t)$, the exogenous FRM rate prevailing at time t, which the borrower keeps until a refinancing occurs. With adjustable rates, the borrower's current interest rate is $i_t^{ARM} (\Theta_t)$, the ARM rate at time t. i_t^{FRM} and i_t^{ARM} are determined based on a yield curve and lending spread for each mortgage type described in the calibration section below.[4] The short interest rate $r_t (\Theta_t)$ is exogenous, stochastic, and a function of the state of the business cycle Θ_t.

At origination, mortgages must satisfy a loan to value constraint :

$$M_{t+1} (a) \leq \phi p_t H_{t+1} (a) , \tag{1}$$

where $t + 1$ is used for M and H because choices of mortgages and housing today determines the entering housing and mortgage balance tomorrow.[5] ϕ parameterizes the maximum loan-to-value ratio.

Mortgages are non-recourse but defaulting carries a utility penalty of d which is drawn each period iid from a uniform distribution over $[d_a, d_b]$.[6] Defaulting households lose their house today

[4] i^{FRM} and i^{ARM} represent the long and short mortgage rates, respectively, and different mortgage designs may have borrowers borrowing at i^{FRM} and i^{ARM} at different times.

[5] Greenwald (2015) and Corbae and Quintin (2015) emphasize the importance of payment-to-income constraints in addition to loan-to-value constraints, which can also easily be added but are currently omitted for parsimony.

[6] The assumption that d is drawn from a distribution rather than a single value helps smooth out the value functions in the numerical implementation, but is not crucial for our results. In practice, d_a and d_b are close and the model is essentially to a single default cost model.

and cannot buy a new house in the period of default due to damaged credit.[7]

Each period, homeowners can take one of four actions in the housing market: take no action with regards to their mortgage and make at least the minimum mortgage payment (N), refinance but stay in their current house (R), move to a new house and take out a new mortgage (M), or default (D). Note that if a household refinances or moves to a new house, they must take out an entirely new mortgage which is subject to the LTV constraint in equation (1). Moving has a cost of $k_m + c_m p_t$ for both buying and selling, while refinancing has a cost of $k_r + c_r M_{t+1}$.

Homeowners occasionally receive a moving shock that forces them to move with probability ζ. In this case, they cannot remain in their current house and either move or default, while agents who do not receive the moving shock are assumed to remain in their house and can either do nothing, refinance, or default.

Finally, regardless of whether they receive a moving shock ζ, renters can either do nothing and pay their rent (N) or move into an owner-occupied house (M) each period.

3.2 Decisions and Value Functions

Consider a household at time t. This household enters the period with housing $H_t \in \{0,1\}$, a mortgage with principal balance M_t, and $S_t(1+r_{t-1}) > 0$ in liquid savings (which has earned the risk free rate r_{t-1}). The state of the economy at time t, Θ_t, is realized. The household receives income Y_t. The agent-specific state $s_t^j = \{S_t, H_t, M_t, i_t, Y_t, a_t\}$ is a vector of the household's assets, liabilities, and income. The vector of aggregate state variables Σ_t includes the state of the economy Θ_t, and $\Omega_t(s_t^j)$, the cumulative distribution of individual states s_t^j in the population. The home price p_t is a function of Σ_t.

The household faces two constraints. The first is a flow budget constraint:

$$Y_t - \tau(Y_t - i_t M_t) + S_t(1+r_{t-1}) + M_{t+1} = C_t + S_{t+1} + (1+i_t) M_t - p_t (H_{t+1} - H_t) \\ + q p_t 1[H_t = 0] + m p_t 1[H_t = 1] + K(\mathit{Action}), \tag{2}$$

where $K(\mathit{Action})$ is the fixed or variable cost of the action the household takes. The left hand side of this expression is the sum of net-of-tax income, liquidated savings, and new borrowings. The right hand side is the sum of consumption, savings for the next period, payments on existing mortgage debt, net expenditures on owner-occupied housing, rental or maintenance costs, and the fixed and variable costs of the action that the household takes.

The second constraint addresses the evolution of a household's mortgage. Given a mortgage balance M_t, implicitly define ΔM_t as the change in the mortgage balance over and above the minimum payment:

$$M_{t+1} = M_t(1+i_t) - M_t \frac{\left(i_t(1+i_t)^{T-a+1}\right)}{(1+i_t)^{T-a+1} - 1} + \Delta M_t. \tag{3}$$

The mortgage balance for the next period is equal to the current mortgage balance inclusive of all

[7] In practice, households that default are locked out of buying for several years depending on the circumstances of their default. We make the assumption that the households must rent today and cannot buy a house tomorrow but is free to purchase thereafter to simplify the household problem and economize on a state variable.

interest costs, minus payments equal to the minimum payment plus ΔM_t. If ΔM_t is positive, the mortgage balance has risen relative to the minimum payment and the homeowner has extracted equity, and if ΔM_t is negative the mortgage balance has prepaid. Thus, households that do not move, refinance, or default face a constraint of $\Delta M_t \leq 0$. If a household moves, it pays off its mortgage balance and chooses a new mortgage balance M_{t+1}, subject to the LTV constraint (1). Finally, a household may also choose to default, in which case it loses its house today and cannot buy, so $M_t = H_t = M_{t+1} = H_{t+1} = 0$.

We write the household's problem recursively. Denote $V_a\left(s_t^j; \Sigma_t\right)$ as the value function for a household, and $V_a^A\left(s_t^j; \Sigma_t\right)$ as the values when following action $A = \{N, R, M, D\}$. Then,

$$V_a\left(s_t^j;\Sigma_t\right) = \begin{cases} \zeta \max\left\{V_a^D\left(s_t^j;\Sigma_t\right), V_a^M\left(s_t^j;\Sigma_t\right)\right\} + \\ (1-\zeta)\max\left\{V_a^D\left(s_t^j;\Sigma_t\right), V_a^R\left(s_t^j;\Sigma_t\right), V_a^N\left(s_t^j;\Sigma_t\right)\right\} & \text{if } H_t > 0 \\ \max\left\{V_a^M\left(s_t^j;\Sigma_t\right), V_a^N\left(s_t^j;\Sigma_t\right)\right\} & \text{if } H_t = 0. \end{cases}$$

On the top line, if the household receives the moving shock with probability ζ, it must decide whether to default on the existing mortgage and be forced to rent, or pay off the mortgage balance, in which case it can freely decide whether to rent or finance the purchase of a new home. On the second line, if the household does not receive the moving shock, it decides between defaulting, refinancing, or paying the minimum mortgage balance. Finally, in the last line, a household that currently has no housing (currently a renter or just born) must decide whether to purchase a house and take on a new mortgage or continue to rent.

We next define the values under each of the actions, $A = \{N, R, M, D\}$. Households who continue to service their mortgage choose their mortgage payment, savings, and consumption to solve:

$$\begin{aligned} V_a^N\left(s_t^j;\Sigma_t\right) = \max_{C_t, S_{t+1}, M_{t+1}} U\left(C_t, H_t\right) + \beta E_t\left[V_{a+1}\left(s_{t+1}^i;\Sigma_{t+1}\right)\right] \text{ s.t. } (2), \\ S_{t+1} \geq 0, \\ H_{t+1} = H_t, \\ i_{t+1} = i_t, \\ \Delta M_t < 0. \end{aligned}$$

Households who refinance make the same choices, but pay the fixed and variable costs of refinancing and face the LTV constraint rather than the $\Delta M_t < 0$ constraint. They have value:

$$V_a^R\left(s_t^j;\Sigma_t\right) = \max_{C_t,S_{t+1},M_{t+1}} \left\{U\left(C_t H_t\left(a\right)\right) + \beta E_t\left[V_{a+1}\left(s_{t+1}^i;\Sigma_{t+1}\right)\right]\right\} \quad \text{s.t. (2)},$$
$$S_{t+1} \geq 0,$$
$$M_{t+1} \leq \phi p_t H_{t+1},$$
$$H_{t+1} = H_t,$$
$$i_{t+1} = i_t^{FRM}.$$

Households who move choose their consumption, savings, and if buying, mortgage balance, as refinancers do, but also get to choose their housing H_{t+1}. They have value:

$$V_a^M\left(s_t^j;\Sigma_t\right) = \max_{C_t,S_{t+1},M_{t+1},H_{t+1}} \left\{U\left(C_t, H_t\right) + \beta E_t\left[V_{a+1}\left(s_{t+1}^a;\Sigma_{t+1}\right)\right]\right\} \quad \text{s.t. (2)},$$
$$S_{t+1} \geq 0,$$
$$M_{t+1} \leq \phi p_t H_{t+1},$$
$$i_{t+1} = i_t^{FRM}.$$

Households who default lose their home but not their savings. The defaulting households choose consumption and savings to solve:

$$V_a^D\left(s_t^j;\Sigma_t\right) = \max_{C_t(a),S_{t+1}(a)} \left\{-d + U\left(C_t\left(a\right), H_t\left(a\right)\right) + \beta E_t\left[V_{a+1}\left(s_{t+1}^a;\Sigma_{t+1}\right)\right]\right\} \quad \text{s.t. (2)}$$
$$S_{t+1} \geq 0,$$
$$H_t = M_t = H_{t+1} = M_{t+1} = 0.$$

In the final period, a household must liquidate its house regardless of whether it gets a moving shock, either through moving or defaulting:

$$V_T\left(s_t^j;\Sigma_t\right) = \max\left\{V_T^N\left(s_t^T;\Sigma_t\right), V_T^D\left(s_t^T;\Sigma_t\right)\right\}.$$

3.3 Mortgage Spread Determination

We assume that mortgages are supplied by competitive, risk-neutral lenders with a one-time origination cost of $\kappa > 0$. In the event of default, the lender forecloses on the home, sells it in the open market, and recovers a fraction Υ of its current value.

Define the net present value of the expected payments made by an age a household with idiosyncratic state s_t^j and aggregate state Σ_t, which is the value of the mortgage to a lender, as

$\Pi_a\left(s_t^j;\Sigma_t\right)$. This can be written recursively as:

$$\Pi_a\left(s_t^j;\Sigma_t\right)=\delta\left(s_t^j;\Sigma_t\right)\Upsilon p_t+\sigma\left(s_t^j;\Sigma_t\right)M_t\left(1+i_t\right)+ \tag{4}$$

$$\left(1-\delta\left(s_t^j;\Sigma_t\right)-\sigma\left(s_t^j;\Sigma_t\right)\right)\left[\begin{array}{c} M_t\frac{\left(i_t(1+i_t)^{T-a+1}\right)}{(1+i_t)^{T-a+1}-1}-\Delta M_t\left(s_{t,a};\Sigma_t\right) \\ +\frac{1}{1+r_t}E_t\left[\Pi_{a+1}\left(s_{t+1}^J,\Sigma_{t+1}\right)\right]\end{array}\right],$$

where $\Delta M_t\left(s_t^j;\Sigma_t\right)$ is the prepayment policy function of the household implicitly defined by (3) and the household policy functions, $\sigma\left(s_t^j,\zeta;\Sigma_t\right)$ is an indicator for whether a household moves or refinances, and $\delta\left(s_t^j,\zeta;\Sigma_t\right)$ is an indicator for whether a household defaults. In the present period, the lender receives the recovered value in the event of a foreclosure, the mortgage principal plus interest in the event the loan is paid off, and the required payment on the mortgage plus any prepayments made by the borrower if the loan continues. The lender also gets the discounted expected continuation value of the loan at the new balance if the loan continues.

We assume that the interest rate paid by the borrower for a given type of loan is a spread over the short end of the yield curve for adjustable rate loans and the long end of the yield curve for fixed rate loans, where the long end is determined by the expectations hypothesis. For mortgages that allow borrowers to choose between an adjustable and fixed rate, we assume the same spread is used over each end of the yield curve.

For now, we assume that lenders determine a single spread π for each type of loan that they charge to all borrowers. This pools risk across borrowers in different states but prices the mortgage so that if a mortgage design shifts risks from borrowers to lenders, the spread rises until the lenders are compensated. In a future draft, we intend to price mortgages differently for each individual based on their state s_t^j.

The condition for the lenders to break even that determines the spread π is:

$$E\left[E_{\Omega_t^{orig}}\left[\frac{1}{1+r_t}\Pi_a\left(s_{t+1}^j;\Sigma_{t+1}\right)-M_t\right]\right]=\kappa, \tag{5}$$

where Ω^{orig} is state distribution of newly originated mortgages. This equates the average value of future loan payments net of the loan principal to the lender's origination cost. We calibrate the model under all FRMs and determine κ from the economy's equilibrium. We then price all other mortgages given this κ.

3.4 Equilibrium

A competitive equilibrium consists of decision rules over actions $A=\{N,R,M,D\}$ and state variables $C_t,S_{t+1},M_{t+1},H_{t+1}$, a price function $p(\Sigma_t)$, a mortgage spread π for each mortgage type, and a law of motion for the aggregate state variable Σ_t. Decisions are optimal given the home price function and the law of motion for the state variable. At these decisions, the housing market clears at price p_t, the risk-neutral lenders break even on average according to (5), and the law of motion for Σ_t is verified.

Given the fixed supply of homes, market clearing simply equates supply from movers, defaulters, and investors who purchased last period with demand from renters, moving homeowners, and investors. Let $\eta\left(s_t^j, \zeta; \Sigma_t\right)$ be an indicator for whether a household moves and $\delta\left(s_t^j, \zeta; \Sigma_t\right)$ be an indicator for whether a household defaults.Movers and defaulters own $H_t\left(s_t^j; \Sigma_t\right)$ housing, while buyers purchase $H_{t+1}\left(s_t^j; \Sigma_t\right)$ housing. The housing market clearing condition satisfied by the pricing function $p\left(\Sigma_t\right)$ is then:

$$\begin{aligned} \int \delta\left(s_t^j, \zeta; \Sigma_t\right) H_t\left(s_t^j; \Sigma_t\right) d\Omega_t + \int \eta\left(s_t^j, \zeta; \Sigma\right) H_t\left(s_t^j; \Sigma_t\right) d\Omega_t \\ = \int \eta\left(s_t^j, \zeta; \Sigma_t\right) H_{t+1}\left(s_t^j; \Sigma_t\right) d\Omega_t, \end{aligned} \tag{6}$$

where the first line side is supply which includes defaulted homes and sales and the second line is demand.[8]

3.5 Solution Method

Solving the model requires that households correctly forecast the law of motion for Σ_t which drives the evolution of home prices. Note that Σ_t is an infinite-dimensional object due to the distribution $\Omega_t(s_t^j)$. In general, this infinite-dimensional object is impossible to handle computationally. To simplify the problem, we follow the implementation of the Krusell and Smith (1998) algorithm in Kaplan et al. (2016). We focus directly on the law of motion for home prices and assume that households use a simple AR(1) forecast rule that conditions on the state of the business cycle today Θ_t and the realization of the state of the business cycle tomorrow Θ_{t+1} for the evolution of p_t:

$$\log p_{t+1} = a_0\left(\Theta_t, \Theta_{t+1}\right) + a_1\left(\Theta_t, \Theta_{t+1}\right) \log p_t. \tag{7}$$

[8] We also introduce investors into our model to rule out unrealistic pathological cases that occur with an extreme sequence of negative shocks. In particular, it is possible that the housing demand and supply functions cross in a manner that housing prices fall near zero before rising back to more normal levels in the next period. Such a price path implies an unrealistically high return on housing for that period. To rule out such prices paths, we add risk neutral investors to the model. These investors receive no flow utility from housing and earn returns purely from a possible capital gain. We assume that in order to purchase investors must expect to earn a one-period return of at least r^{inv} on the housing investment, where r^{inv} is 7.0 percent to match the rate of return cited by arbitrageurs who purchased and rented out foreclosed homes in the crisis.

On the equilibrium path in all of our simulations, the investors sit passively on the sideline. Thus their introduction is largely to rule out rare pathological cases off the equilibrium path. Demand from investors can be written as:

$$d^{inv}\left(\Sigma_t\right) = \begin{cases} 0 & \text{if expected one period return} < r^{inv} \\ \text{unpurchased housing stock} & \text{otherwise} \end{cases}$$

In forming their expectations, investors use the same rational expectations based house price function $p\left(\Sigma_t\right)$ as all other agents. The market clearing condition becomes:

$$\begin{aligned} \int \delta\left(s_t^j, \zeta; \Sigma_t\right) H_t\left(s_t^j; \Sigma_t\right) d\Omega_t + \int \eta\left(s_t^j, \zeta; \Sigma\right) H_t\left(s_t^j; \Sigma_t\right) d\Omega_t + inv_t \\ = \int \eta\left(s_t^j, \zeta; \Sigma_t\right) H_{t+1}\left(s_t^j; \Sigma_t\right) d\Omega_t + d^{inv}\left(\Sigma_t\right), \end{aligned}$$

where the stock of homes owned by investors from last period inv_t is added to supply and investor demand $d^{inv}\left(\Sigma_t\right)$ is added to demand.

Expression (7) can be viewed either as a tool to compute equilibrium in heterogeneous-agent economies, following Krusell and Smith (1998), or as an assumption that households and investors are boundedly rational and formulate simple forecast rules for the aggregate state. To verify that the decision rule is accurate, we both compute the R^2 for each (Θ_t, Θ_{t+1}) realized in simulations and follow Den Haan (2010) by comparing the realized price with the 15, 30, 45, and 100-year ahead forecasts given the realized process of aggregate shocks to verify that the forecast rule does a good job of computing expected prices many periods into the future and that small errors do not accumulate.

The model cannot be solved analytically, so a computational algorithm is used. First, the household problem is solved using the forecast rule by discretization and backward induction. Given the household policy functions, the spread is adjusted so that the lender breaks even on average, and the household problem is resolved. This is repeated until the spread converges. Given the household policy functions and the spread, the model is simulated for many periods with the home price determined by (6) and the AR(1) forecast regression (7) is run in the simulated data for each (Θ_t, Θ_{t+1}). Finally, the forecast rule is updated based on the regression, and the entire procedure is repeated until the forecast rules converges to an approximate solution.

4 Calibration

Our calibration strategy focuses on quantitatively matching micro default behavior in response to changes in interest rates and loan balances and the distribution of assets and mortgage debt in the population. The former is crucial for counterfactuals to asses mortgage design. The later is necessary to accurately reflect the number of individuals who would be affected by different mortgage designs at the margin and how their decisions aggregate to affect the housing market equilibrium.

To capture these features of the world, our calibration proceeds in three steps. First, we select the aggregate and idiosyncratic shocks to reflect modern business cycles in the United States. Second, we exogenously calibrate a number of parameters to standard values in the macro and housing literature or to match moments in the data. Third, we choose the utility benefit of owning a home and the default cost to match a target price level and new quasi-experimental evidence about the effect of debt on default. Importantly, our model does well in matching lifecycle patterns relating to housing and housing debt.

Throughout, we calibrate to the data using a model in which all loans are fixed rate mortgages to reflect the predominant mortgage type in the United States and credit constraints are at their pre-downturn level. Table 1 summarizes the variables and their calibrated values. κ, the fixed origination cost for the lender, is backed out from the FRM equilibrium and imposed in solving for the model's equilibrium for other mortgages. The calibration is annual.

4.1 Aggregate and Idiosyncratic Shocks

We calibrate the Markov transition matrix for the state of the business cycle Θ_t based on the frequency and duration of NBER recessions and expansions. Recall that Θ_t has three values: crisis,

Table 1: Model Parameters in Baseline Parameterization

Param	Description	Value	Param	Description	Value
T	Years in Life	45	c_m	Variable Moving Cost as % of Price	3.0%
R	Retirement	35	k_m	Fixed Moving Cost	0.1
ρ	Log Income Decline in Retirement	0.35	c_r	Variable Refi Cost as % of Mortgage	1.0%
τ_0	Constant in Tax Function	0.8	d_a	Default Cost Dist Lower Bound	15.0
τ_1	Curvature Tax Function	0.18	d_b	Default Cost Dist Upper Bound	25.0
γ	CRRA	3.0	k_r	Fixed Refi Cost	0.04
ξ	Terminal Wealth Multiplier	1.0	q	Rent	0.20
ψ	Terminal Wealth Shifter	500	m	Maint Cost as % of Prices	2.5%
a	Utility From Homeownership	8.0	ζ	Prob of Moving Shock	1/9
β	Discount Factor	0.96	ϕ	LTV Constraint	92.5%
Υ	Foreclosure Sale Recovery %	0.654		Homeownership Rate	65.0%
r	Short Rate	[0.26%, 1.32%, 3.26%]			
i^{ARM}	ARM Interest Rate	[3.01%, 4.07%, 6.01%]			
i^{FRM}	FRM Interest Rate	[4.96%, 5.48%, 5.66%] (expectations hypothesis)			
Y^{agg}	Aggregate Income	[0.0976, 0.1426, 0.1776]			
See text for transition matrix for Θ_t and Y_t^{id}.					

Note: This table shows baseline calibration for a fixed rate mortgage and a maximum LTV of 92.5 percent at origination. Average income is normalized to one. There are three aggregate states, $\Theta_t \in \{\text{Crisis,Recession,Expansion}\}$ and the tuples of interest rates reflect the interest rate in each state.

recession, and expansion. We use the NBER durations and frequencies to determine the probability of a switch between an expansion and crisis or recession, and we assume that crises happen every 75 years and all other NBER recessions are regular recessions. We assume that every time the economy exits a crisis or recession it switches to an expansion and that crises affect idiosyncratic income in the same way as a regular recession but last longer and involve a larger aggregate income drop, with a length calibrated to match the average duration of the Great Depression and Great Recession. A regular recession reduces aggregate income by 3.5 percent, while a crisis reduces it by 8.0 percent, consistent with Guvenen et al.'s (2014) data on the decline in log average earnings per person in recessions since 1980.

We calibrate short rates and mortgages rates during expansions and recessions to historical real rates from 1985-2007. We find that short rates equal 1.32% on average during recessions and 3.26% during expansions. For the crisis state, we assume that the real short rate is 3.0% less than during expansions, or 0.26%. The short rate faced by households that save, r in the model, is one percent above the short rate in the data, reflecting higher rates of returns on savings that are illiquid for the one-year duration of a period in our model. We set the calibrated ARM spread over the short rate in the data to 2.75%, its historical average margin over the short Treasury rate in the Freddie Mac Primary Mortgage Market Survey. The FRM interest rate is set using the expectations hypothesis on the ARM rate with a term of 10 years.[9]

For the idiosyncratic income process, we match the countercyclical left skewness in idiosyncratic income shocks found by Guvenen et al. (2014). Left skewness is crucial to accurately reflect default

[9]In practice, the 30-year fixed rate mortgage is priced off of the 10-year Treasury bond.

in a crisis because the literature on mortgage default has found that large income shocks are crucial drivers of default. To incorporate left skewness, we calibrate log idiosyncratic income to follow a Gaussian AR(1) with an autocorrelation of 0.91 and standard deviation of 0.21 following Floden and Linde (2001) in normal times but to have left skewness in the shock distribution in busts. We discretize the income process in normal times by matching the mean and standard deviation of shocks using the method of Farmer and Toda (2016), which discretizes the distribution and optimally adjusts it to match the mean and variance of the distribution to be discretized. For the bust, we add the standardized skewness of the 2008-9 income change distribution from Guvenen et al. (2014) to moments to be matched, generating a shock distribution with left skewness. This gives a distribution with a negative mean income change in busts and leads to income being too volatile, so we shift the mean of the idiosyncratic shock distribution in busts to match the standard deviation of aggregate log income in the data. In doing so, we choose the income distribution of the newly born generations to match the lifecycle profile of income in Guvenen et al. (2016).[10] We normalize the income process so that aggregate income is equal to one.

4.2 Standard Parameters

Having set the parameters that determine aggregate and idiosyncratic shocks as well as interest rates, we then set a number of parameters to standard values in the literature or directly to match moments in the data.

We assume households live for 45 years, roughly matching ages 25 to 70 in the data. Households retire after 35 years, at which point idiosyncratic income is frozen at its terminal level minus a 0.35 log point retirement decrease. The tax system is calibrated as in Heathcote et al. (2017), with $\tau_0 = 0.80$ and $\tau_1 = 0.18$.

For preference parameters, we use a discount factor of $\beta = 0.96$ and a CRRA of $\gamma = 3.0$. The bequest function parameters are chosen so that consumption is smooth at the end of life.

Moving and refinancing involve fixed and variable costs. We set the fixed cost of moving equal to 10 percent of annual income, or $5,000. The proportional costs, paid by both buyers and sellers, equal 3 percent of the house value to reflect closing costs and realtor fees. Refinancing involves a fixed cost of 4 percent of annual income, or $2,000, as well as variable cost equal to 1 percent of the mortgage amount to roughly match average refinancing costs quoted by the Federal Reserve.

Renters pay a rent of $q = 0.20$ to match a rent-to-income ratio of 20%. Homeowners must pay a maintenance cost equal to 2.5% of the house value every year. We assume that homeowners move an average of every 9 years as in the American Housing Survey. The homeownership rate is set to match the long-run average homeownership rate of 65 percent in the United States.

Υ, the fraction of the price recovered by the bank after foreclosure, is set to 64.5 percent. This combines the 27 percentage point foreclosure discount in Campebell, Giglio and Pathak (2011) with the fixed costs of foreclosing upon, maintaining, and marketing a property, estimated to be 8.5 percent of the sale price according to Andersson and Mayock (2014).[11]

[10] Rather than including a deterministic income profile, we start households at lower incomes and let them stochastically gain income over time as the income distribution converges to its ergodic distribution. This does a good job of matching the age-income profile in the data.

[11] Much of the literature calibrates to the "loss severity rate" defined as the fraction of the mortgage balance

4.3 Matching Quasi-Experimental Evidence on Default

The remaining parameters are a, the utility benefit of owning a home, and $\bar{d}$, the average default cost.[12] We calibrate these two parameters to recent micro evidence on the effect of interest rates and LTV ratios on default as well as a mean price to income ratio for homeowners of 4, which is the mean value in the SCF.[13]

In particular, we focus on quasi-experimental evidence from Fuster and Willen (2015). Fuster and Willen study a sample of homeowners who purchased ALT-A hybrid adjustable-rate mortgages during 2005-2008 period and quickly fell underwater as house prices declined. Under a hybrid ARM, the borrower pays a fixed rate for several years (typically five to ten) and then the ARM "resets" to a spread over the short rate once a year. These borrowers were unable to refinance because they owed the bank more than their house's value, and so when their rates reset to reflect the low short rates after 2008, they received a large and expected reduction in their monthly payment.

Fuster and Willen provide two key facts for our purposes. First, they show that even for ALT-A borrowers – who have low documentation and high LTVs relative to the population – at 135 percent LTV the average default hazard prior to reset was only about 24 percent.[14] The fact that so many households with significant negative equity do not default implies that there are high default costs. It is also consistent with a literature that finds evidence for a "double trigger" model of default whereby both negative equity and a shock are necessary to trigger default, as is the case for most default in our model.

Second, Fuster and Willen use an empirical design that compares households just before and after they get a rate reset and show that these borrowers experience substantial declines in their default rate at reset. In particular, the hazard of default for a borrower receiving a 3.0 percent rate reduction fell by about 55 percentage points, equivalent to going from a 145 percent LTV to 95 percent LTV.

We match our model to Fuster and Willen's estimates by simulating their rate reset experiment within our model. This innovation in our calibration strategy allows our model to accurately capture the effects of changes in LTV and interest rates on default quantitatively. In particular, we compare the default behavior of agents in our model with a 2/1 ARM that will reset next period with the behavior of an agent with a 1/1 ARM that has reset this period. This corresponds to the treatment and control used by Fuster and Willen. We assume that these borrowers are an infinitesimal part of the market, so we can consider them in partial equilibrium, and we compute their default rates at different LTVs with the 2/1 ARM and 1/1 ARM. To deal with the fact that the ALT-A sample used by Fuster and Willen is not representative of the population, we roughly match the assets,

recovered by the lender. We calibrate to a fraction of the price because of a recent empirical literature that finds that in distressed markets, the loss recovery rate is much lower (e.g. Andersson and Mayock, 2014), which is consistent with a discount relative to price rather than a constant loss severity rate.

[12]We choose d_a and d_b, the bounds of the uniform distribution from which d is chosen, to add a small bit of mass around $\bar{d}$. In the calibration, $\bar{d} = 20$, $d_a = 15$, and $d_b = 25$.

[13]Homeowners have an average income of 1.3 times the average income in our model so the price is 5.25 times average income. The SCF calculation Winsorizes the top and bottom 1% to drop extreme outliers on house value and income due to measurement error.

[14]This figure is based on the default hazard in months 30 to 60 in Figure 1b. Fuster and Willen measure "default" as becoming 60 days delinquent rather than an actual foreclosure, so the actual default rate might be slightly lower.

Figure 1: Fit to Fuster and Willen (2015) Natural Experiment

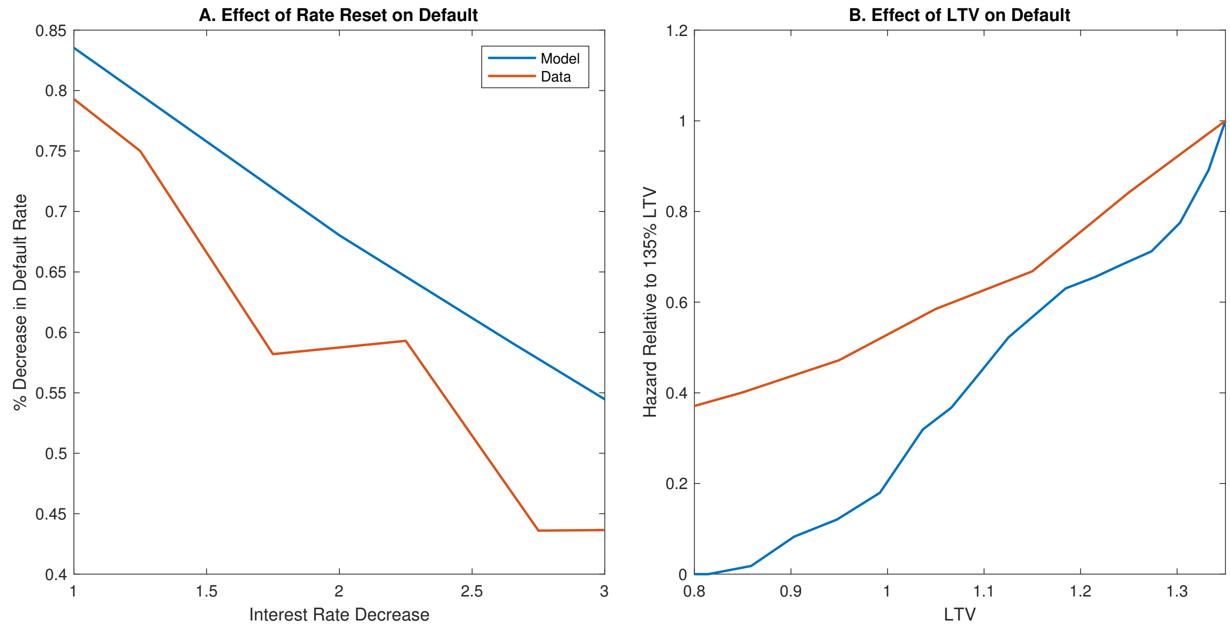

Note: The data from Fuster and Willen come from column 1 of Table A.1 in their paper, which is also used in Figure 2a. The model estimates come form comparing a 2/1 ARM to a 1/1 ARM in our model.

age, and income of the homeowners we consider to households with hybrid ARMs that have yet to reset in the 2007 Survey of Consumer Finances.[15] Finally, we assume that homeowners have a fixed rate corresponding to the FRM rate in the boom and reset to the ARM rate in the crisis for calculating the baseline default hazard. For calculating the effect of the rate reset on default, we consider a one percent, two percent, and three percent interest rate decline by adjusting the initial fixed rate of the hybrid ARM.

Figure 1 compares the calibrated model with the findings of Fuster and Willen (2015). Panel A shows the impact of rate reductions on default in the model and Fuster and Willen's estimates. Overall, the fit is quite good, although at very large rate reductions the model slightly under-predicts the reduction in the default hazard resulting from an interest rate decline. The right panel shows the baseline default rate under the 2/1 ARM at various LTVs relative to the default rate at 135 percent LTV. LTV reductions have only modestly larger effects than in the data until one gets below 100 percent, at which point the default hazard falls off in the model but not in Fuster and Willen's data.[16] The model also does well in terms of the level of default at 135 percent LTV: the

[15] We only roughly match the data because the SCF has few observations that match these criteria. Consequently, we assume a uniform distribution between the 25th and 75th percentiles of age, income, and assets in the SCF data. We find that the ALT-A borrowers have low assets, are young, and have moderate-to-low income, as one would expect.

[16] Fuster and Willen find a substantial default hazard below 100 percent LTV for three likely reasons. First, the combined LTV that Fuster and Willen use is likely measured with error both due to missing liens in the data and due to error in the automatic valuation model. Second, Fuster and Willen measure default as 60 day delinquency and not a final foreclosure. In areas with substantial foreclosure backlogs, borrowers who are above water can become delinquent before they sell. Finally, there are some frictions in terms of time to sell and the fixed costs of sale that may cause above-water households to default. In our model, there is minimal default for above water households

Figure 2: Lifecycle Patterns: SCF vs. Model

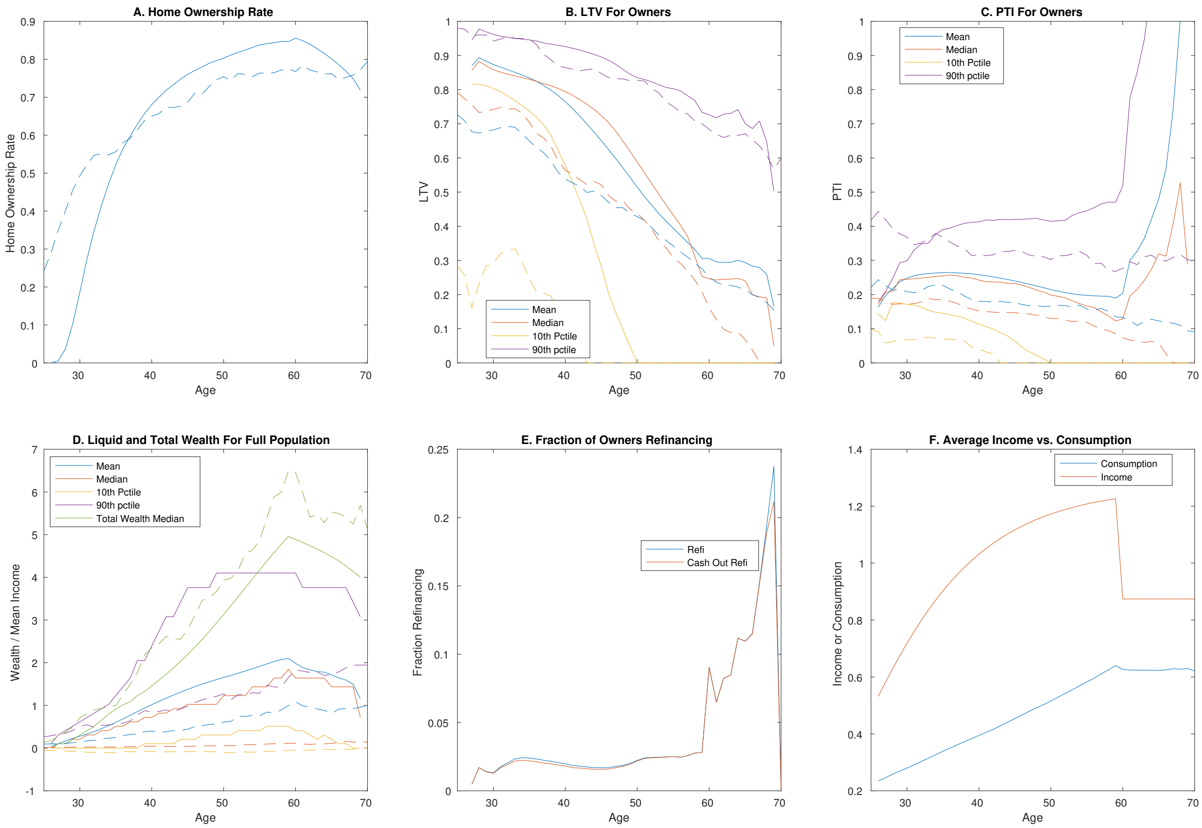

Note: This figure compares the model (solid lines) to SCF data from 2001 to 2007 (dashed lines) in panels A-D. Panels E and F are constructed based only on the model. Panel A shows the homeownership rate. Panel B shows the mean, median, 10th percentile, and 90th percentile of loan to value ratios for homeowners, and panel C shows the same statistics for the payment to income ratio. Panel D shows the mean, median, 10th percentile, and 90th percentile of liquid assets along with median total wealth. Panel E shows the refinance rate, and Panel F shows consumption and investment in the model.

model has a default hazard for the 2/1 ARM of 23.5 percent, falling just short of the roughly 24 percent in Fuster and Willen's sample.

Based on the close fit to the Fuster and Willen quasi-experimental evidence, our model should do a good job of reflecting the benefits of principal and payment reductions.

4.4 Lifecycle Patterns and Distributions Across the Population

The model does a good job matching the lifecycle patterns and overall distribution of debt and assets in the Survey of Consumer Finances for 2001, 2004, and 2007. Figure 2 shows the lifecycle patterns, while Figure 3 shows the distributions across the population. In both figures, the pooled SCF data for 2001 to 2007 is in dashed lines and the model analogues are in solid lines.

Panel A of Figure 2 shows the homeownership rate over the lifecycle. The model slightly underestimates the homeownership rate of the very young and over-estimates the homeownership rate of the middle age.

because some households get a moving shock with very low equity and decide to default.

Figure 3: Distributions Across Population: SCF vs. Model

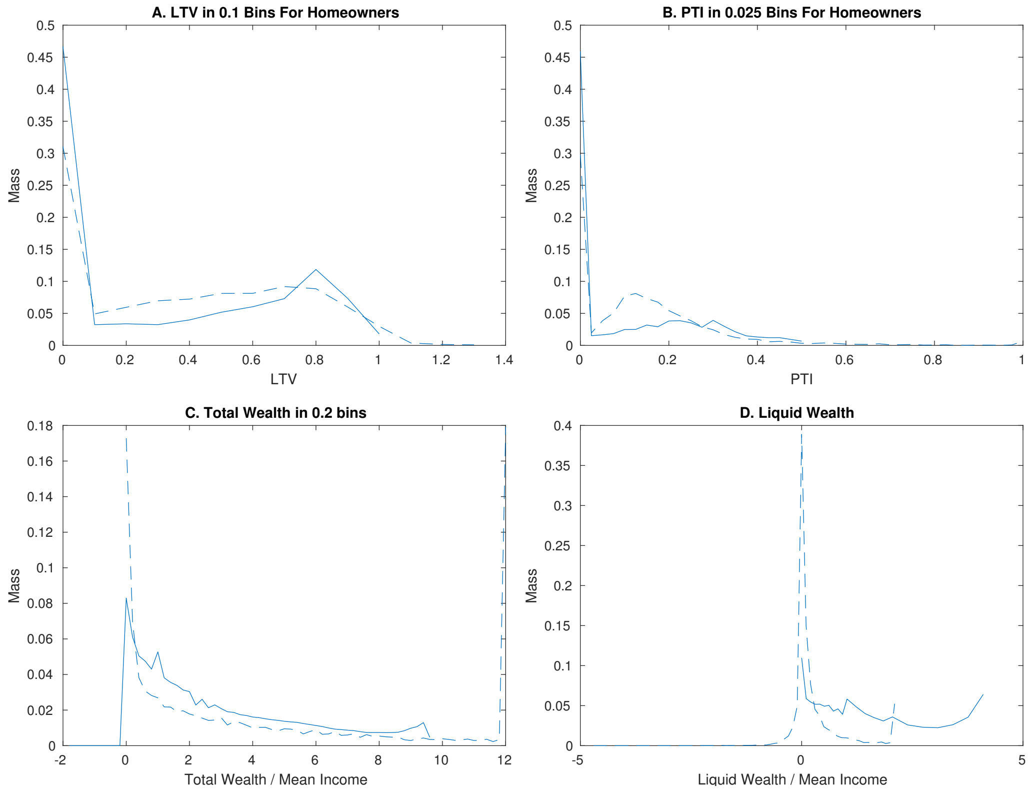

Note: This figure compares the model (solid lines) to SCF data from 2001 to 2007 (dashed lines). Panel A shows loan to value ratios in 10 percentage point bins for homeowners, and panel B shows the payment to income ratio for homeowners in 0.025 bins. Panel C shows total wealth relative to mean income in bins of 0.2, while panel D shows liquid wealth relative to mean income in bins of 0.2. In all figures, the model and data are binned identically.

Panel B of Figure 2 shows the mean, median, 10th percentile, and 90th percentile of the loan to value ratio (LTV) by age, and panel A of Figure 3 shows the distribution of LTV across the population. On the whole, the model does very well at capturing the overall LTV distribution and the rate at which the mean, median, and 90th percentile homeowner pays down their mortgage. This translates into an LTV distribution that closely matches the data, although we have a few too many homeowners who have paid off their mortgage and slightly too few homeowners with high LTVs, reflecting the fact that we impose a 92.5 percent maximum LTV constraint on all agents to reflect 90th percentile LTV of the young in the SCF, while in practice a few homeowners obtain loans with an even higher LTV. The fact that we match the LTV distribution in the data implies that when a crisis hits, our model will have a fraction of underwater and low equity homeowners that is close to the data for the housing bust.

Panel C of Figure 2 shows the mean, median, 10th percentile, and 90th percentile of the payment to income ratio (PTI) by age, and panel B of Figure 3 shows the distribution of PTI across the population. We do reasonably well for the PTIs of the young, although the old have PTIs that are too high. This is the case both because income falls in retirement and because the old who still have a mortgage in the data have a loan that amortizes past age 70. By contrast, our model assumes that they have a loan that must be paid off at age 70, leading to high payments at the end of life. Because most of the equilibrium effects in our model come from the purchase, refinance, and default decisions of the young who have relatively high LTVs (panel B), the fact that our calibration is off for elderly homeowners does not dramatically affect market equilibrium. Overall, our PTI distribution has slightly too much mass at zero PTIs, but the key feature for our analysis of a crisis is that we do a good job matching the upper tail of PTIs.

Panel D of Figure 2 shows the mean, median, 10th percentile, and 90th percentile of the liquid wealth and the median of total wealth by age, and panels C and D of Figure 3 show the distributions of total and liquid wealth in the population. The model does a reasonably good job matching median total wealth over the lifecycle and liquid wealth at young ages. Agents in the model accumulate more liquid assets in retirement than in the data. Again, this is not a significant issue, as the old do not play a crucial role in the housing market in our model. The old accumulate too many assets and so the total and liquid wealth distributions are skewed a bit to the right and there are too few people with zero assets, but we do a good job matching the distribution of liquid and total wealth for the young, who tend to have lower assets. Finally, the data has a thicker right tail of very wealthy individuals. Our model is designed to capture the impact of credit constraints and mortgages on housing markets, so capturing the extremely wealthy is not relevant for our exercise.

Finally, Panels E and F of Figure 2 show the fraction of owners refinancing and income and consumption over the life cycle, respectively. Most refinancing is of the cash out variety because the FRM rate does not fluctuate dramatically due to the expectations hypothesis. Cash out refinancing is low until retirement, at which point it jumps so that agents can smooth their consumption. Again, the refinancing of the old is not crucial to our results, and we plan to address it in a future draft. Income follows a standard lifecycle profile, and consumption is smoother than income and increasing as individuals age, consistent with buffer stock models of consumption.

5 ARM vs. FRM: The Economics of State Contingent Mortgages

Having created and calibrated a laboratory to study mortgage design and its interaction with monetary policy, we now use our model to assess various mortgages. We primarily focus on a crisis scenario that combines a housing bust and a deep recession as in the Great Recession. This allows us to analyze mortgage designs proposed to address the problems revealed by the Great Recession, which is the focus of the recent literature. Additionally, the equilibrium feedbacks that our model features are most interesting and potent in a downturn with a price-default spiral. In the future, we plan to also consider cases where the central bank raises rates in a bust, as in the early 1980s.

To simulate a housing downturn, we combine a deep and persistent recession – which lowers aggregate income and leads to more frequent negative idiosyncratic shocks – with an unexpected tightening of downpayment constraints. Specifically, we consider an unexpected shock which lowers the maximum loan-to-value (LTV) ratio from 92.5% to 80%.[17] This shock lowers demand for housing by the young, who have to save up for longer to afford a down payment and leads to a fall

[17] We choose an unexpected shock because agents did not expect a substantial housing decline in the run up to the crisis, which helps explain the levels of default in the bust. It is difficult to get substantial default if the crisis is anticipated, but in future drafts we plan to consider a case where a crisis is a very low probability event. Despite the unexpected nature of the shock, once the shock occurs the forecast rules are quite accurate.

Figure 4: FRM vs. FRM→ARM: Housing Market Outcomes

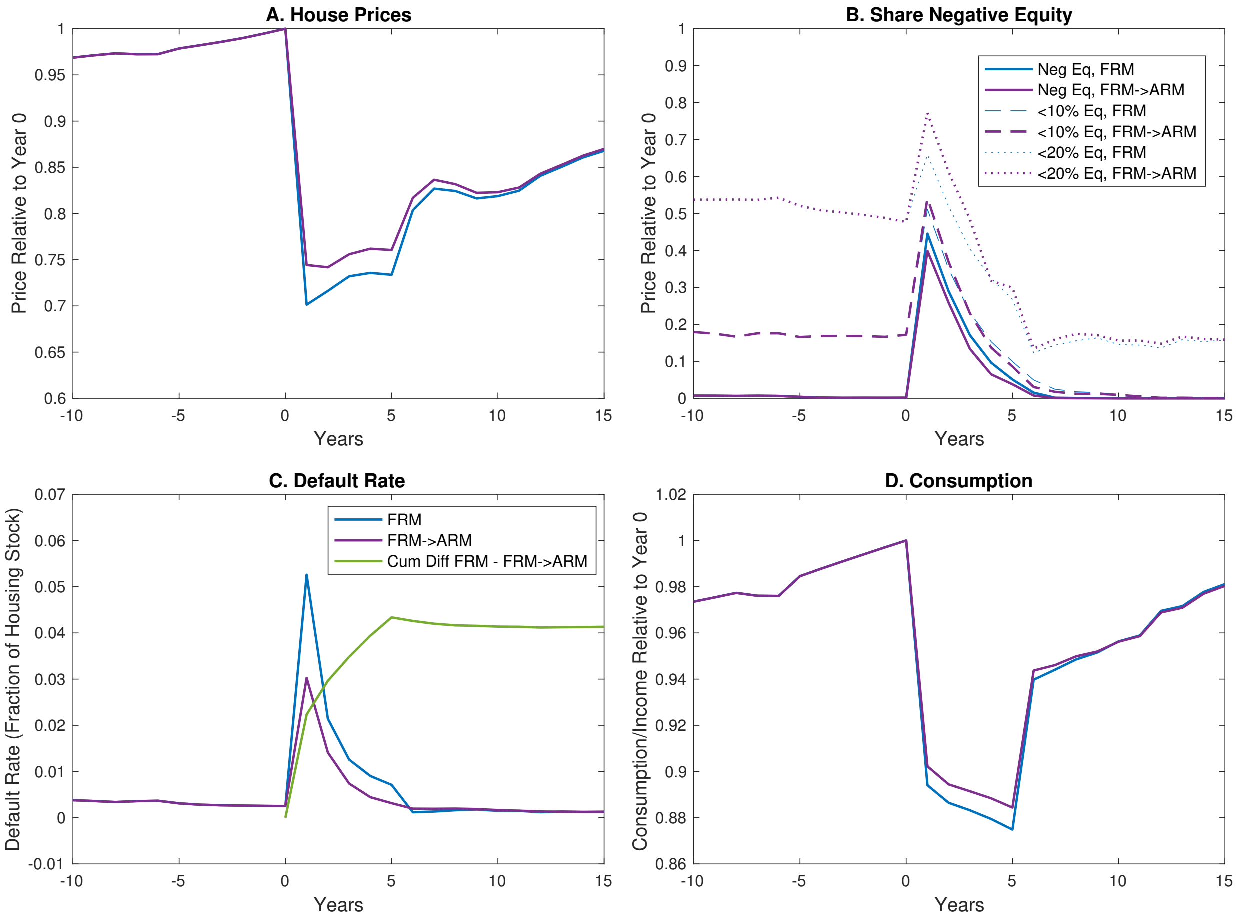

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under FRM and FRM→ARM.

in house prices. We also assume that in the five years prior to the tightening of the downpayment constraint the economy is in an expansion, and coincident with the tightening of the downpayment constraint the economy enters a crisis state for five years. We study the impulse responses of prices, default rates, consumption, and welfare to the resulting downturn.[18]

To analyze the effect of mortgage design in such a crisis, we first compare a world with all fixed rate mortgage borrowers to a world with all adjustable rate mortgage borrowers. This provides us with most of the economic intuition regarding the benefits of adding state contingency to mortgages. We also briefly compare FRMs and ARMs outside of a deep recession. In subsequent sections, we consider different mortgage designs and the interaction of mortgage design and monetary policy.

5.1 Economic Intuition: FRM vs. ARM

Our baseline case is a world in which the only available mortgage to home purchasers is a fixed-rate mortgage. The results are illustrated by the blue lines in Figure 4.

[18] We compute impulse responses by averaging together 100 simulations with random shocks prior to the five year expansion and subsequent to the five year recession.

The model generate a housing crisis in the model of a similar magnitude to the experience in the United States between 2006 and 2012. Panel A of Figure 4 shows that prices fall by about a third, which closely matches the peak to trough decline in national repeat sales house price indices.[19] Panel B shows the fraction of homeowners who have negative equity, under 10 percent equity, and under 20 percent equity. At the depths of the crisis, nearly 40 percent of homeowners are underwater, which roughly matches the data. Approximately 70 percent of homeowners have less than 20 percent equity when prices are at their lowest and are unable to refinance into a lower interest rate mortgage given the tightened LTV constraint. Panel C shows the combination of negative equity and recession leads to substantial default. Cumulatively, 10.6 percent of the housing stock experiences a foreclosure over the first eight years of the crisis, compared with 8 percent in the data from 2006 to 2013. Finally, homeowner consumption dramatically due to the sudden and persistent decline in income and the large number of constrained households.

We examine the differential impacts of adjustable-rate mortgages through two experiments. In the first, we assume that home purchasers take out fixed-rate mortgages pre-crisis, but that when the crisis hits, all mortgages are converted to adjustable-rate mortgages. Because the central bank lowers the short rate in the crisis and this is fully passed through to households under ARM, mortgage payments fall dramatically. This experiment, which we call the FRM→ARM counterfactual (FARM in figure labels), is a useful thought experiment to understand the mechanisms at work in our model because it holds fixed the state of the economy when the crisis hits and isolates the *ex-post* effect of adjustable-rate mortgages on the severity of the crisis. In the second experiments, we consider a world with ARMs both before and after the crisis hits, allowing for *ex-ante* behavior and a different distribution of households across states at the beginning of the crisis.

5.1.1 *Ex-Post* Effect of Switching From FRM to ARM

We present the results of the FRM→ARM counterfactual in the purple lines in Figure 4. Relative to the baseline case, the housing crisis is less severe and defaults fall substantially. Indeed, the green line in panel C shows the cumulative difference in the default rates between the FRM and FRM→ARM worlds over five years is just over four percent, so that 37 percent fewer homeowners default.

The intuition for these effects has to do with the extent to which homeowners can take advantage of lower interest rates in a crisis. In the baseline case, 70 percent of homeowners have less than 20 percent equity and cannot refinance to take advantage of these lower rates. Because of left-skewness of the income shock distribution in the crisis, a significant fraction of these homeowners experience a simultaneous drop in their income. These homeowners become liquidity constrained and those with little savings default because they would have to cut their consumption substantially – and in many cases to zero – to make their mortgage payment, which causes them to be willing to bear the utility cost of default. Default increases the supply of homes on the market, further pushing down prices, which in turn leads to more default. This phenomenon is the canonical price-default spiral.

[19] Our model does not match the time series of house price indices, because prices fall to their lowest level on the impact of the shock, while in the data they decline gradually. This is the case because our Walrasian model does not have any house price momentum. See Guren (2017) for a summary of the literature on momentum.

In the ARM world, homeowners do not need to refinance to take advantage of the lower interest rates. Because the mortgage payment is pegged to the prevailing short rate in the market, payments fall automatically and by more than under FRM because the short end of the yield curve adjusts by more under the expectations hypothesis than the long end. Many defaulters under an FRM find that they would rather cut their consumption and avoid the costs of default. This leads to less default, short-circuiting the default spiral and causing a less severe housing crisis.

Additionally, because FRMs are priced off of the long end of the yield curve, buying becomes more affordable for young first-time homebuyers, increasing their demand in the ARM world more than the FRM world. This limits price declines and further ameliorates the impacts of the housing crisis. Note that the effects are largely concentrated among first-time home buyers since existing homeowners are hedged against the downturn.

These effects are summarized in Figure 5, which plots the mass of homeowners defaulting and renters purchasing at each age in the period in which the crisis begins under FRMs and under the FRM→ARM counterfactual along with the difference. Because the pre-downturn distribution is the same, this only reflects differences in policy functions in the FRM and FRM→ARM worlds. Similar comparisons by age, income, and LTV reveal that the additional default under FRM comes from low income, low savings, and high LTV borrowers, while the additional demand comes from renters with moderate savings and income reflective of first-time homebuyers.

Figure 6 shows the cumulative consumption-equivalent welfare difference between the FRM and ARM worlds as a function of age. Consumption equivalent welfare for each age group is calculated as the amount of consumption one would have to take away in the period before the crisis in order to obtain the welfare in a given period, and is reported as a percentage of annual consumption. A negative value indicates a larger welfare loss under FRMs than in the FRM→ARM counterfactual.

Figure 6 reveals that the welfare benefits of ARMs are concentrated on the young. There is little incidence on the very young who have yet to purchase a home. The old are barely affected because they have significant home equity and liquid savings to smooth income shocks and can refinance if necessary while satisfying their LTV constraint. By contrast, the young have recently purchased a home have high LTVs, cannot refinance and are stuck at a high interest rate, and do not have much liquid savings to help cushion an income shock. Young homeowners also tend to have lower consumption and a higher marginal utility of consumption at baseline because they are borrowing constrained and expect their income to rise later in life. They are thus willing to pay a substantial amount of annual consumption to switch to an ARM, as the lower interest rate translates to a much smaller minimum payment and boosts consumption substantially for these highly constrained households. Young renters also benefit from improved affordability in the crisis.

Figure 7 shows the aggregate decline in consumption equivalent welfare over the simulated crisis. The left panel shows the annual decline in welfare in the FRM world (blue), the FRM→ARM world (orange), the per-year difference (yellow), and the cumulative difference (purple). Recall from panel D of Figure 4 that consumption is about one percent lower under the FRM relative to the FRM→ARM world. The welfare effects are, however, much larger: over five years, the cumulative welfare loss is about 13.6 percent of one year of consumption. This is the case because of the large welfare losses for the high marginal utility of consumption households that are most

Figure 5: FRM vs. FRM→ARM: Default and Renter Demand By Age

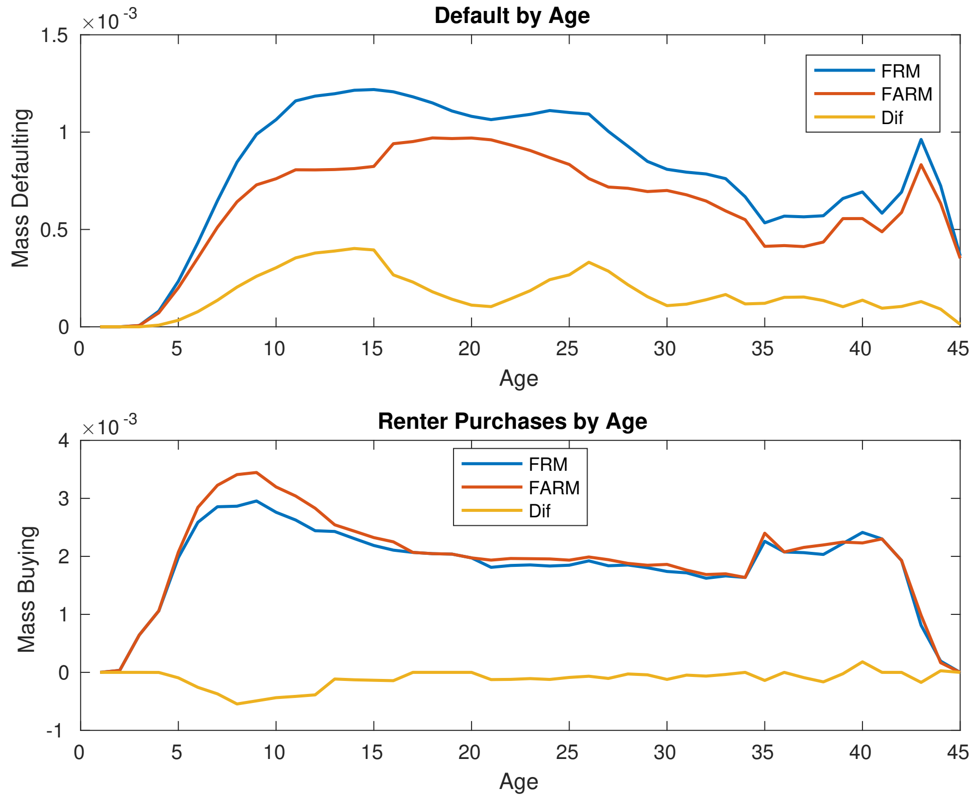

Note: The figure shows the mass of households of each age defaulting and purchasing after renting last period in the first period of the crisis under FRM and FRM→ARM along with the difference between the two lines. To ensure that the differences are entirely due to the policy functions, we compare FRM to FRM→ARM and evaluate the policy functions assuming the price is the realized FRM price.

Figure 6: FRM vs. FRM→ARM: Welfare By Age

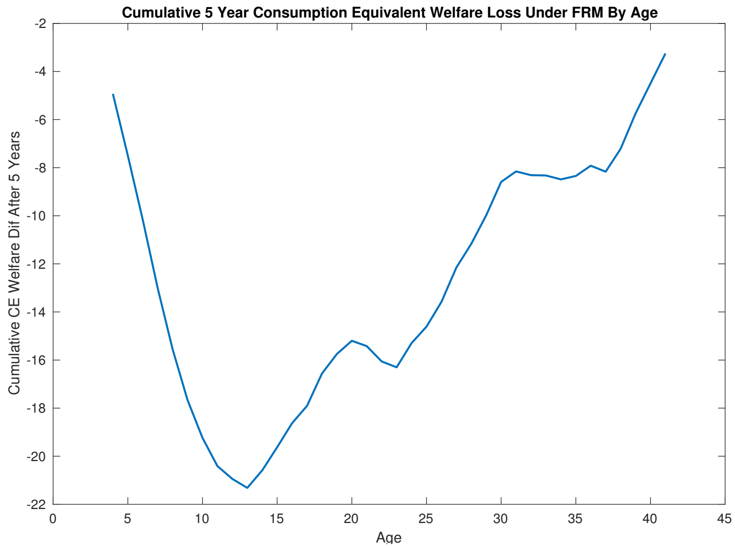

Note: The figure shows the relative consumption equivalent welfare decline by age under FRM relative to FRM→ARM. The line is a moving average of 5 ages centered at the indicated age.

Figure 7: FRM vs. FRM→ARM: Welfare Decomposition

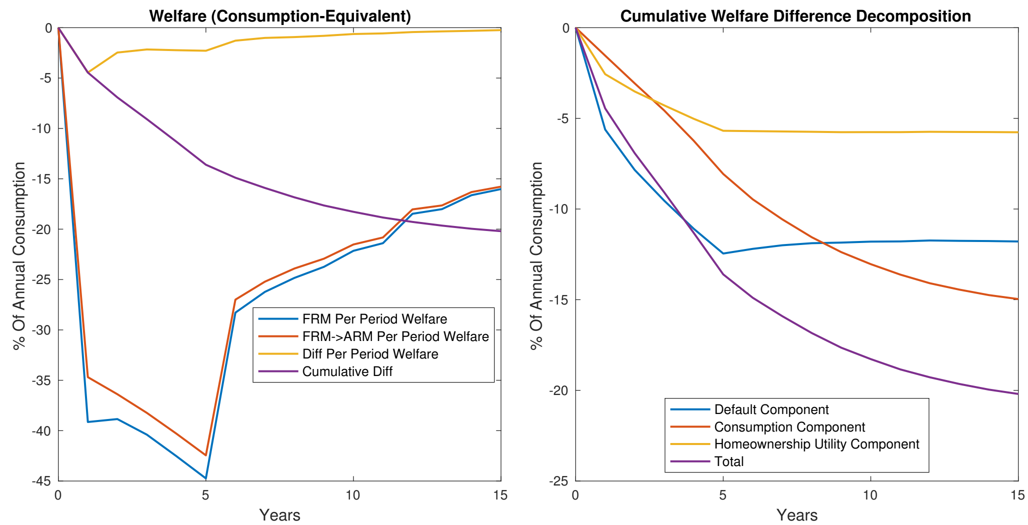

Note: The figure shows the consumption equivalent welfare under FARM→ARM, FRM, the difference, and the cumulative difference in the left panel. The right panel

helped by switching to an ARM.

The right panel of Figure 7 decomposes the cumulative welfare loss of the crisis in purple into three sources. First, consumption falls as households that experience income shocks – or households undertaking precautionary behavior due to potential future income shocks – cut their consumption to make their mortgage payment. Second, those households that are foreclosed on incur default costs, as shown by the blue line. Third, fewer young households own and more older households own, which means that more households with low utility from homeownership own, a small effect effect shown by the yellow line. Consumption is higher under the ARM due to its hedging properties – payments fall when labor income falls – and default is lower under the ARM mortgage since less households become liquidity constrained.

5.1.2 *Ex-Ante* Effect of Switching From FRM to ARM

Our second experiment considers the *ex-ante* effects of adjustable-rate mortgages, in addition to the *ex-post* effects. That is, relative to the baseline case in which all mortgages are fixed-rate, we now consider a world in which all mortgages are adjustable-rate, both pre- and post-crisis. We call this the ARM counterfactual. The results are presented in Figures 8, with the FRM world shown in blue and the ARM world shown in orange. Perhaps surprisingly, the benefits of adjustable-rate mortgages are partially eliminated. Prices fall by less, but not as much as the FRM→ARM counterfactual. *More* households are initially underwater and over eight years, there are 22 percent few defaults than under FRM.

The reason that the crisis is worse in the ARM world relative to the FRM→ARM world is that homeowners understand the hedging properties of the ARM mortgage and take on more risk by increasing their LTV. That is, homeowners understand that with an ARM mortgage, their mortgage payments will fall when the economy enters into a recession, when they are more likely to experience

Figure 8: FRM vs. ARM: Housing Market Outcomes

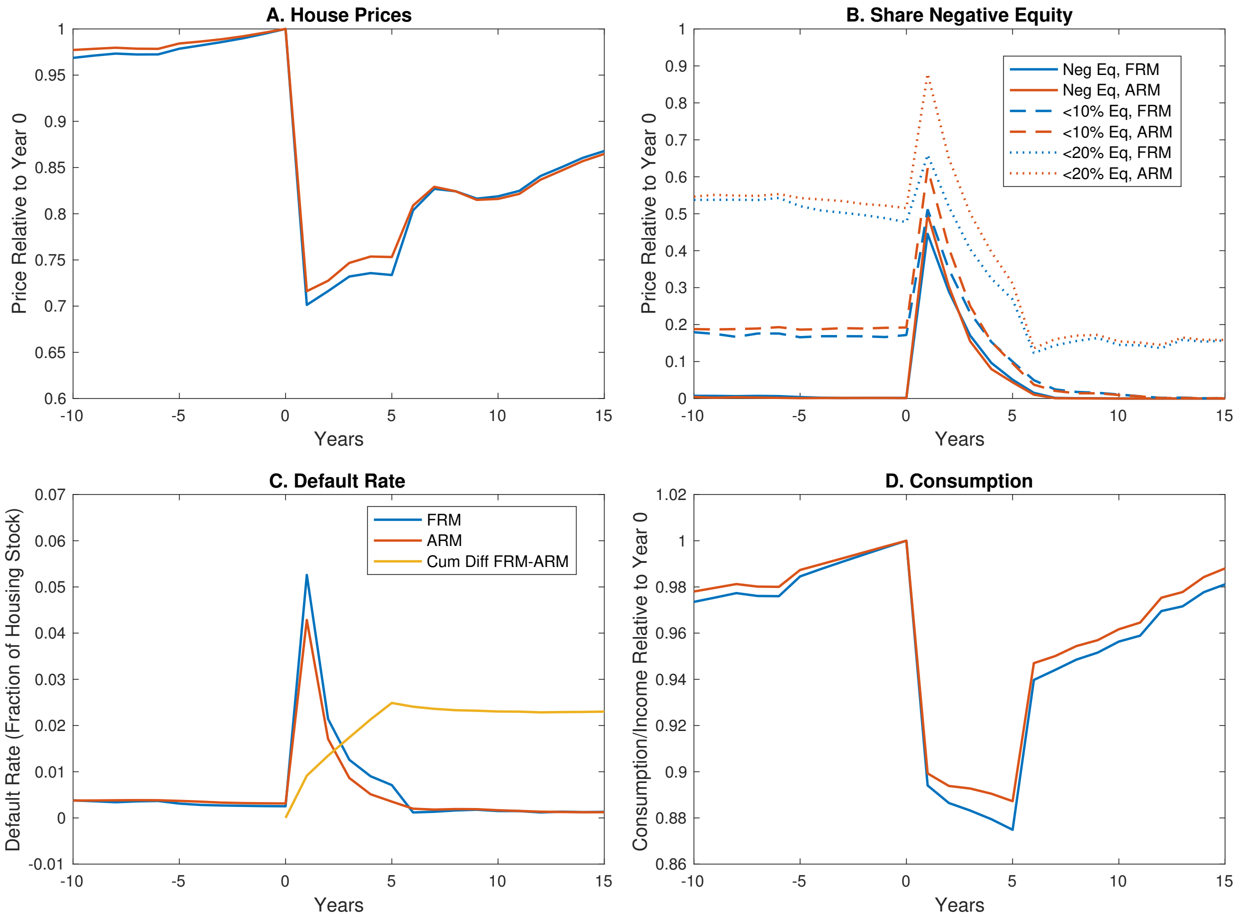

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under FRM and ARM both *ex ante* and *ex post*.

a drop in their labor income. These hedging properties encourage home purchasers to consume more when young relative to the FRM world by taking out loans with higher initial leverage or by cash-out refinancing. This creates macro fragility, which households do not internalize. These effects of switching to ARM *ex ante* on the LTV distribution is shown in Figure 9. In the ARM world in orange, there is more mass at the highest LTV ratios than the FRM world in blue. Consequently, while the ARM offers *ex-post* benefits by lowering mortgage payments and reducing default, given *ex-ante* decisions it leads to a crisis which is similar to the all FRM economy. The degree to which the *ex-post* benefits of ARMs are undone by the *ex-ante* buildup of fragility – at least for house prices and default – is a numerical result. This highlights the need for realistic quantitative models of the type we analyze.

5.2 ARM vs. FRM in Stochastic Simulations

To further elucidate the differences between the ARM and FRM worlds, we compare moments from stochastic simulations. These simulations have a fixed maximum LTV and thus more mild crises relative to the crisis we simulate and should thus be thought of as representing how a mortgage

Figure 9: FRM vs. ARM: Pre-Crisis LTV

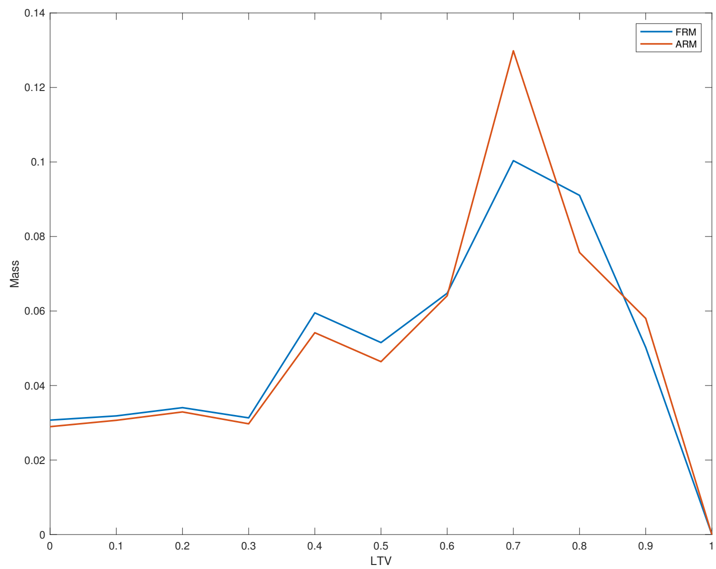

Table 2: FRM vs. ARM: Moments From Stochastic Simulations

Moment	ARM 92.5% Rel to FRM 92.5%	ARM 80% Rel to FRM 80%
Std Dev of Price	88.0%	89.1%
St Dev of Default	38.9%	46.2%
Std Dev of Consumption	86.6%	90.3%
Mean Welfare	101.6%	99.9%
STD Welfare	90.6%	91.0%

Notes: All series are percentages relative to the same statistic for the indicated FRM model. Mean and std of welfare uses a utilitarian social welfare function.

performs in “normal times.”[20]

Table 2 shows the standard deviation of price, default, consumption, and welfare along with the mean of welfare for ARMs relative to an FRM world with the same down payment. The first column shows a 92.5 percent maximum LTV world, while the second shows an 80 percent maximum LTV.

Although mean welfare and prices are roughly similar in the ARM and FRM worlds using a utilitarian welfare criterion, the insurance benefits of ARMs become evident when once examines the volatility of these variables. These results are particularly acute for default, but are also evident for consumption and house prices and consequently welfare. These effects are stronger at a 92.5 percent maximum LTV because more households have a high LTV, and so the insurance benefits are stronger given equilibrium feedbacks.

As we have seen, for particularly severe crises, ARMs can dramatically reduce housing market volatility. Together, then, these results suggest that ARMs provide important insurance.

[20]In future drafts we plan to include crises of the type we simulate in these stochastic simulations.

6 Evaluating New Mortgage Designs

Although we have so far focused on the insurance benefits of ARMs, ARMs do have some drawbacks. In particular, if the central bank raises interest rates to fight inflation, our assumption that in recession short rates fall may be violated. This was the case in the Volcker recessions in the early 1980s. In these cases, ARMs are *worse* from an insurance perspective, the reverse of our main results because the covariance of interest rates and income shocks switches. Furthermore, by revealed preference most Americans prefer fixed rate mortgages, which on average account for roughly 80 percent of mortgage originations.

Given these downsides, in this section we evaluate several mortgage designs that allow for state contingency in a crisis while preserving the benefits of FRMs in normal times. We also evaluate a contract that provides more insurance than a standard ARM. In the future, we plan to evaluate additional designs, including shared-appreciation mortgages.

6.1 Eberly-Krishnamurthy Convertible Mortgage

Eberly and Krishnamurthy (2014) propose a fixed rate mortgage that can at any time be converted to an adjustable rate mortgage, but not back. This is similar to a world in which one can choose between ARM and FRM with two important distinctions. First, homeowners who do not satisfy the LTV constraint can still switch to an ARM. Since in our simulated crisis up to 70 percent of homeowners have under 20 percent equity, this is likely to be significant. Second, homeowners who are so highly constrained that they cannot afford to pay the fixed costs of refinancing can convert to an ARM costlessly.

This mortgage also has three added benefits. First, introducing it in practice is likely to be less disruptive, as it can act exactly like a FRM and does not take away the FRM option. Second, in the event that the covariance of interest rates and monetary policy changes, it performs like an FRM. Third, to the extent that it cushions bank losses in a crisis, the spread banks charge will fall.

The convertible mortgage does, however, come at a cost as borrowers will convert to an ARM when it is beneficial to them and then refinance into a new loan that starts as an FRM when this is no longer the case. Because ARM rates are lower than FRM rates when rates fall, this reduces the NPV of the mortgage to the lender through standard prepayment risk and results in a higher spread. This competes with the lower spread due to lower default losses. On net, in our model the prepayment risk is larger the Eberly-Krishnamurthy mortgage has a 4.8 basis point higher mortgage rate relative to an FRM.[21]

Figure 10 shows a simulated downturn under the Eberly-Krishnamurthy mortgage, both prior to and after the crisis (as will be the case for our other designs). The crisis looks roughly 85 percent of the way to the all-ARM world and realizes many of the consumption-smoothing and macroprudential benefits of ARMs. This is the case because on average the share of mortgages that are functioning as ARMs rises from 15 percent pre crisis to 85 percent during the crisis as

[21] The endogenous pricing response is fairly small here and for the other mortgages we consider in part because banks and households consider the crisis we simulate to be a zero probability event. We plan to relax this assumption in the future so reduced foreclosures will have a more realistic feedback onto spreads.

Figure 10: Eberly-Krishnamurthy Convertible Mortgage: Housing Market Outcomes

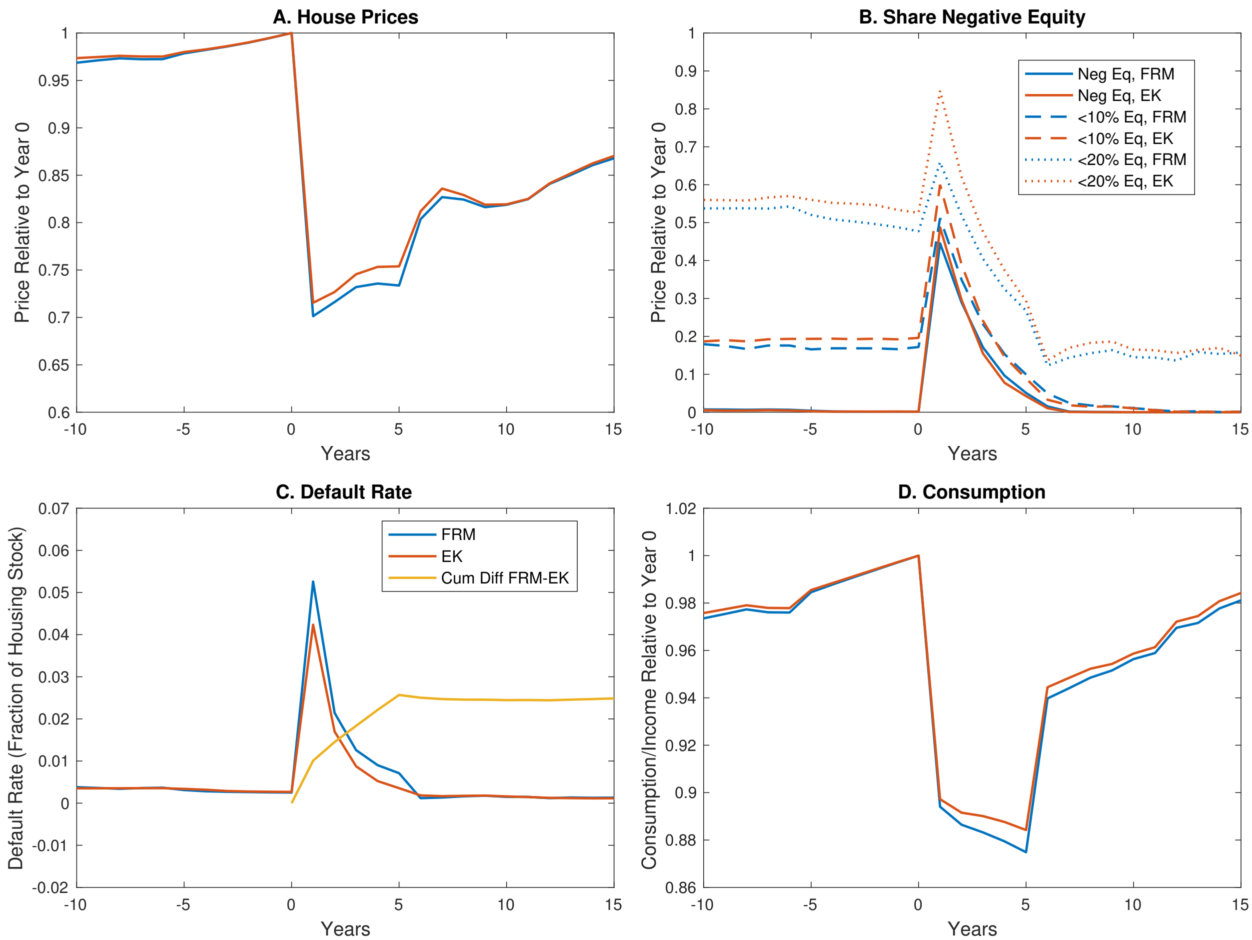

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under FRM and the Eberly-Krishnamurthy mortgage both *ex ante* and *ex post*.

households who need the payments relief the most convert.[22] The convertible mortgage does not look fully like an ARM because older, richer household keep to their FRM.

Table 3 compares the Eberly-Krishnamurthy mortgage with an FRM in stochastic simulations. Again, the convertible mortgage has nearly all of the benefits of the ARM. There is a slight decline in average welfare, mostly because households refinance frequently and spend a larger fraction of their budget paying the adjustable and fixed costs of refinancing, which do not contribute to welfare.

6.2 Fixed Rate Mortgage With Underwater Refinancing

The second mortgage that we consider is a fixed-rate mortgage that can be refinanced even when the LTV constraint at origination is not satisfied, which we call an FRM mortgage with underwater

[22]In our model, households with a choice between an ARM and an FRM, which is what the Eberly-Krishnamurthy mortgage effectively looks like in normal times, will choose to hold about 40 percent ARMs. This is higher than the data, where the average ARM share is about 20 percent. To account for this difference, one would need to subsidize the FRM slightly, as is the case institutionally in the United States. In simulating the Eberly-Krishnamurthy mortgage, we do not put a subsidy on the FRM but not the ARM.

Table 3: FRM vs. ARM: Moments From Stochastic Simulations

Moment	Eberly-Krishnamurthy	FRM Underwater Refi	Option ARM
Std Dev of Price	88.5%	102.1%	83.5%
St Dev of Default	49.2%	101.3%	53.1%
Std Dev of Consumption	90.5%	99.9%	85.8%
Mean Welfare	97.5%	98.4%	99.2%
STD Welfare	90.6%	99.6%	89.0%

Notes: All series are percentages relative to the same statistic for the indicated FRM model. Mean and std of welfare uses a utilitarian social welfare function.

refinancing. This is similar to the Home Affordable Refinance Program (HARP) pursued in the recent downturn. Campbell (2013) has advocated FRMs with underwater refinancing, which are used in Denmark through a system of covered bonds that allow for refinancing when the LTV constraint at purchase is violated.

We implement this in our model by allowing FRM households to refinance – but not originate a mortgage for anew purchase – at a new interest rate as long as they do not reduce their principal.

FRM with underwater refinancing acts like an FRM except it directly addresses the issue of not being able to refinance when one is underwater, undoing the pecuniary default externality in our model that operates through one person's default driving down prices and preventing others from refinancing. Because payments fall in bad states, it should help smooth consumption and reduce default losses for lenders. It will, however, increase lenders prepayment risk. On net, in our model, the FRM with underwater refinancing has a 2.5 basis point higher spread than an FRM.

Figure 11 shows the crisis under FRMs with underwater refinancing. The key to understanding the effects is to note that the FRM with underwater refinancing is priced off the long end of the yield curve, which falls in a crisis but not dramatically. Because the insurance it provides is minimal, households do not take on much more risk, and consumption is roughly the same as FRMs in the crisis. The initial decline in prices is shallower because there is less default as the initial relief kicks in, but more houses stay in a precarious position and are forced to default when an income shock occurs, so some of the initial benefit in terms of the default rate, consumption, and prices is undone as the crisis does not ameliorate as quickly. In normal times, the FRM with underwater refinancing behaves much like an FRM in, as shown in Table 3. Overall, then, allowing for underwater refinancing provides some macroprudential benefits in a crisis but only limited consumption-smoothing benefits because the FRM is still priced off the long end of the yield curve.

6.3 Option ARM

Piskorski and Tchistyi (2010) argue in using an optimal contracting model that the optimal contract looks roughly like an option ARM. This is an ARM mortgage that became popular in the early 2000s boom that allows the borrower to make one of three payments: a fully amortizing payment, an interest only payment, and a potentially negatively amortizing payment equal to the minimum payment based on the interest rate at origination. The negative amortization is allowed up to a ceiling, and after several years the option ARM converts to a fully amortizing ARM. Piskorski and

Figure 11: FRM With Underwater Refinancing: Housing Market Outcomes

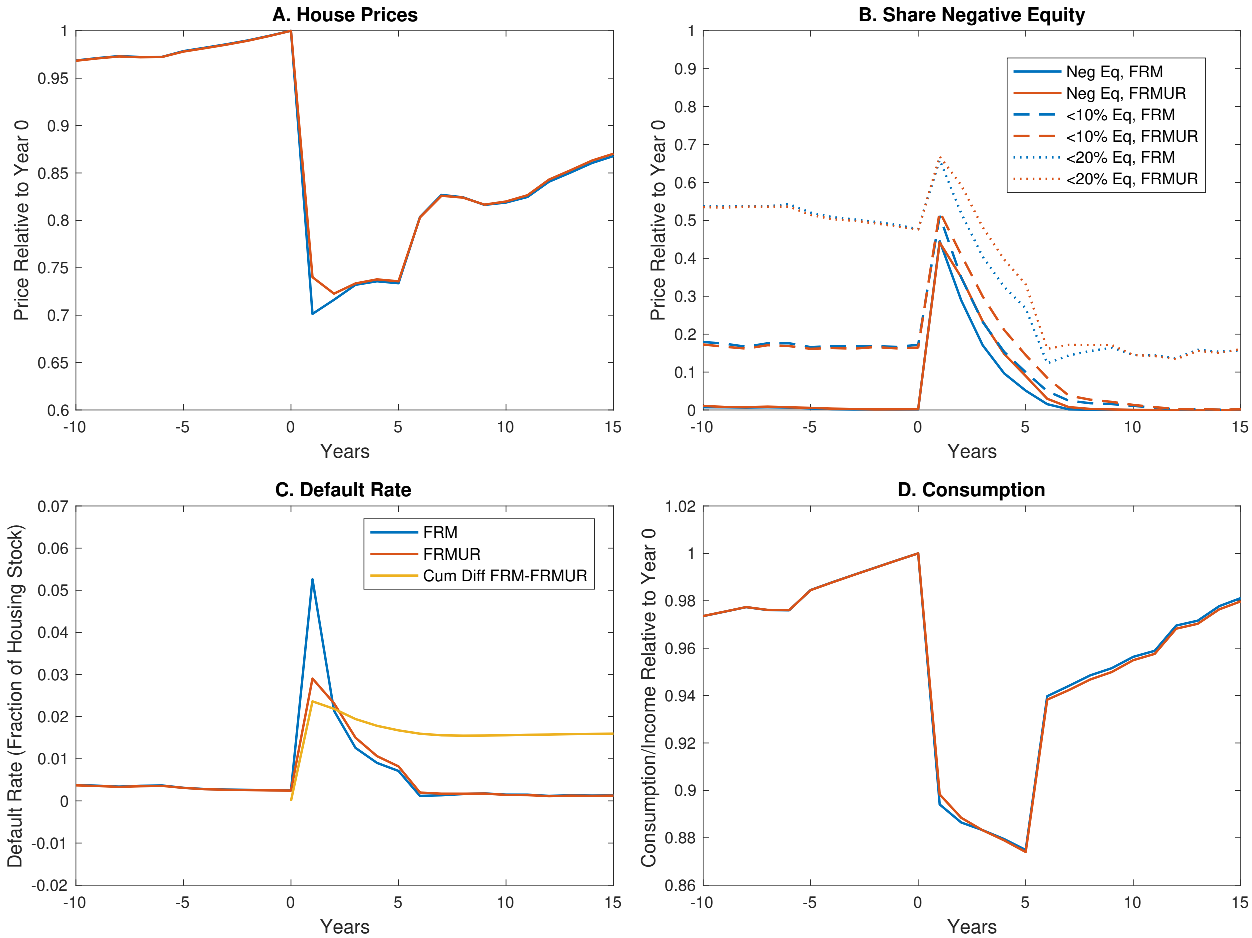

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under FRM and the FRM with underwater refinancing both *ex ante* and *ex post*.

Tchistyi find this is roughly optimal because it provides borrowers with substantial insurance up until a ceiling, which is in place to protect the lender from default.

To mimic the option ARM in our model, we assume that for households in the last 25 years of their life the mortgage behaves like a normal ARM and amortizes. However, households in the first 20 years of life are able to choose a mortgage balance next period equal to the maximum of their current balance and the maximum balance under today's LTV constraint ϕ. This allows for some negative amortization up to a ceiling defined by ϕ, and makes it so that households are not margin called by the bank when the LTV constraint tightens at the beginning of the crisis.

The option ARM adds default risk, but the lender also gets increased interest payments in normal times if the loan does not default. In our model, the later slightly outweighs the former, and the spread is 1.4 basis points below the FRM spread.

Table 3 replicates Piskorski and Tchistyi's findings. In normal times, the option ARM provides higher average welfare and most of the reduced volatility that the ARM provides. This is in part because households can borrow more without refinancing, and thus have smoother consumption.

Figure 12: Option ARM: Housing Market Outcomes

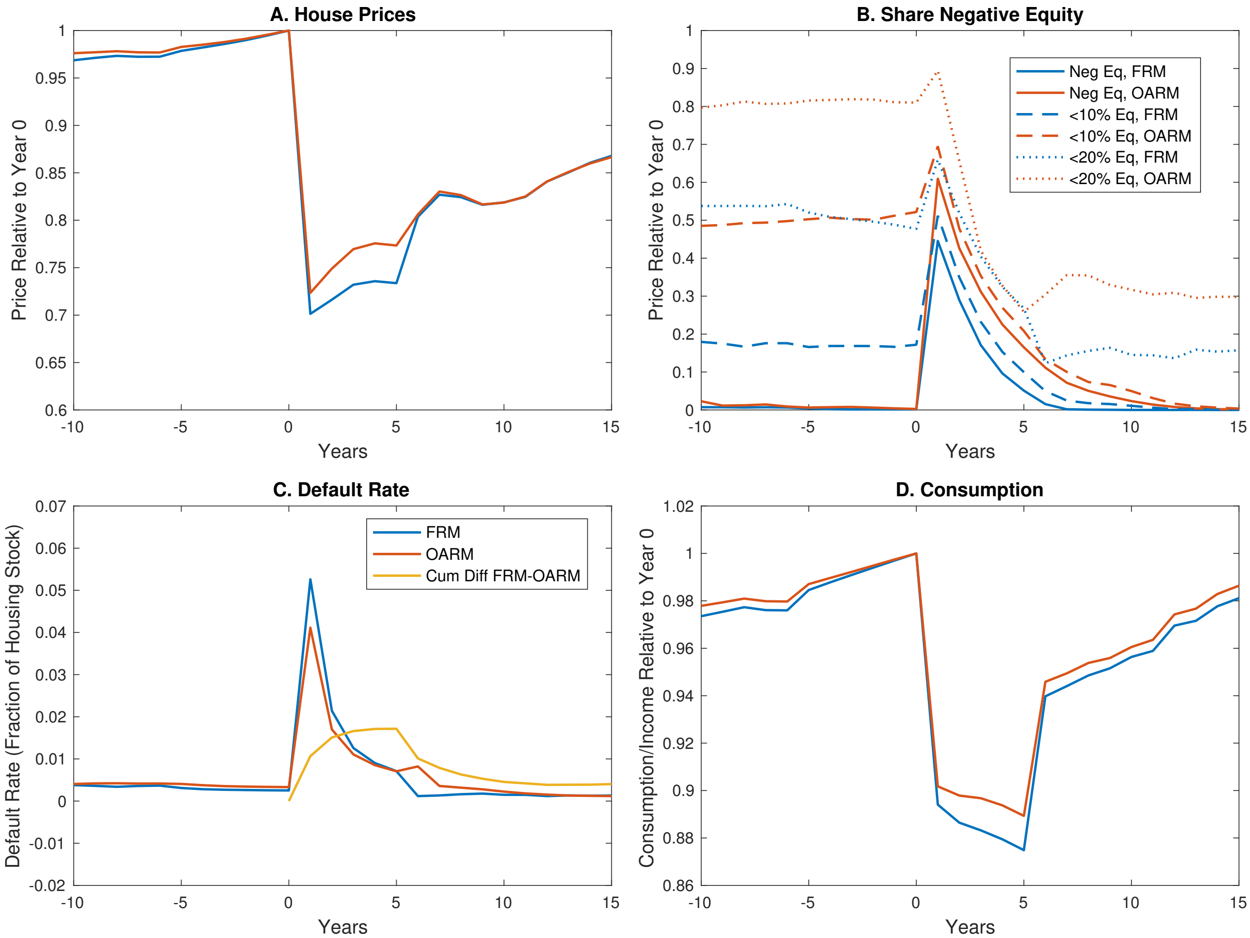

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under FRM and an option ARM both *ex ante* and *ex post*.

However, Figure 12 shows that this comes at the cost of worse performance in the crisis. The option ARM provides more insurance, so households choose more leverage as with the ARM. Additionally, unlucky households who get negative income shocks choose the negative amortization option and increase their LTV further. The increased fragility stemming from more mass at the top of the LTV distribution leads to roughly the same amount of default under FRM over eight years. If the social costs of default are high, the option ARM, which performs well in normal times, may be imprudent from a macro perspective relative to other designs that introduce adjustable rates.

7 The Interaction of Mortgage Design and Monetary Policy

We now turn to examining how monetary policy interacts with mortgage design in a crisis. To keep things simple, we return to comparing the all ARM and all FRM economies and introduce three different modifications to the baseline downturn experiment. In the first modification, we consider expansionary monetary policy in crisis states whereby the central bank lowers the short

Figure 13: Monetary Policy Interaction: Short Rate Falls By More in Crisis

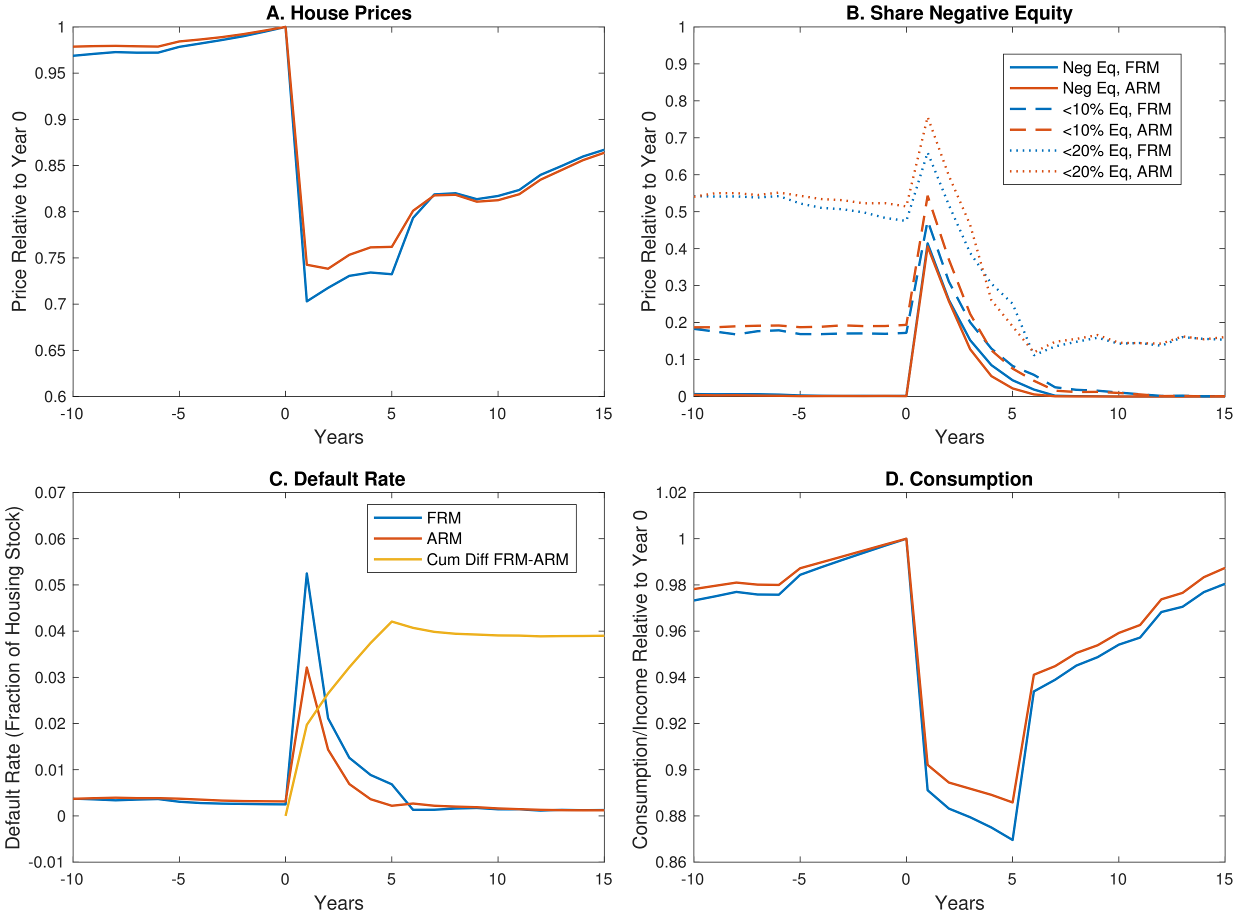

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under an FRM and ARM or the case where the central bank more further reduces the short rate by 50 basis points but the long rate only responds modestly as it is set according to the expectations hypothesis.

rate and consequently ARM mortgage rates by an additional 100 basis points, but this change is only partially passed through to long interest rates, which adjust according to the expectations hypothesis. We think of this as corresponding to more aggressive traditional monetary policy. In the second modification, we assume that in addition to adjusting the short rate, in crisis states only the central bank takes actions that push the long rate down an additional 100 basis points beyond the expectations hypothesis response. We think of this as corresponding to unconventional policies that seek to affect the long end of the yield curve or long-term mortgage rates directly, such as quantitative easing. The third modification assumes that the central bank is doveish and pursues the second policy that combines short rates and QE not only in rare crises but also in more common regular recessions.

The results of the first modification, which reduces only short rates, are shown in Figure 13. Traditional monetary policy has very little impact on the severity of the crisis under FRM. While the return to saving in liquid assets changes, the FRM interest rate does not change. Since few households have liquid assets and most saving in the economy is for retirement, housing demand by

Figure 14: Monetary Policy Interaction: Long Rate Falls More Than Expectations Hypothesis

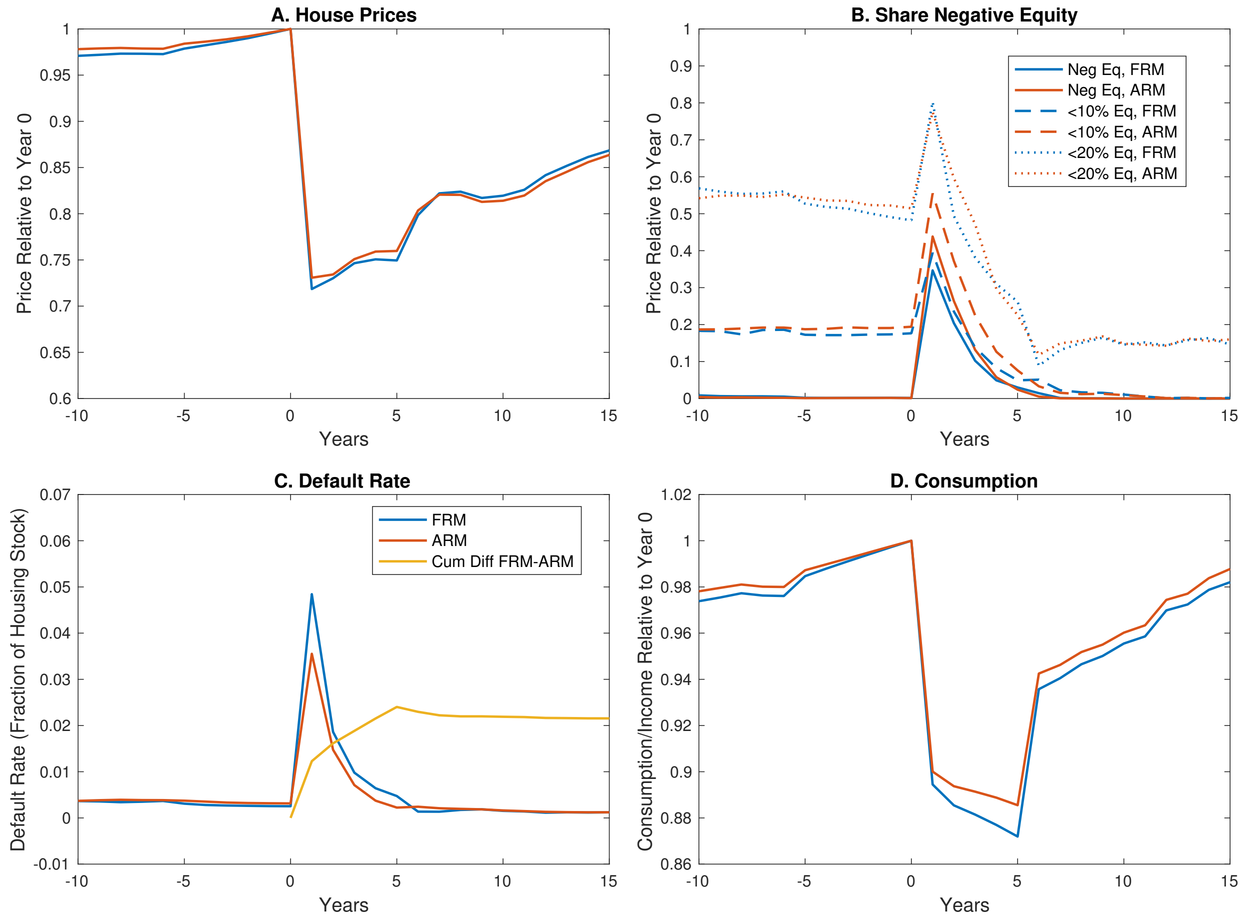

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under an FRM and ARM, for the case where the central bank both further reduces the short rate by 50 basis points and pursues a policy that subsidizes the long rate by 50 basis points in a crisis.

young households and the behavior of high-LTV households that are primarily young is unchanged, and so the housing market equilibrium is not substantially affected.

On the other hand, more aggressive monetary policy that affects short rates is useful when homeowners have adjustable-rate mortgages. Price declines and default are lower and consumption is higher in response to the more aggressive monetary response. This is, of course, not surprising. Lower short-rates leads to lower mortgage payments which leads to less default and a smaller price-default spiral.

When aggressive monetary policy in a crisis is passed through to long rates and FRM interest rates, there are smaller price declines in a world with all fixed-rate mortgages. We demonstrate this in Figure 14. However, the mechanism is quite different from the comparison of ARMs and FRMs. The payments of existing homeowners do not fall, and many remain liquidity constrained. Indeed, there is still significant default. Instead, new homeowners can now lock in cheap financing, which stimulates demand for housing and boosts house prices. Some exiting homeowners are no longer bound by the LTV constraint and are able to refinance due to the rise in prices, undoing some of

Figure 15: Monetary Policy Interaction: Doveish Central Bank

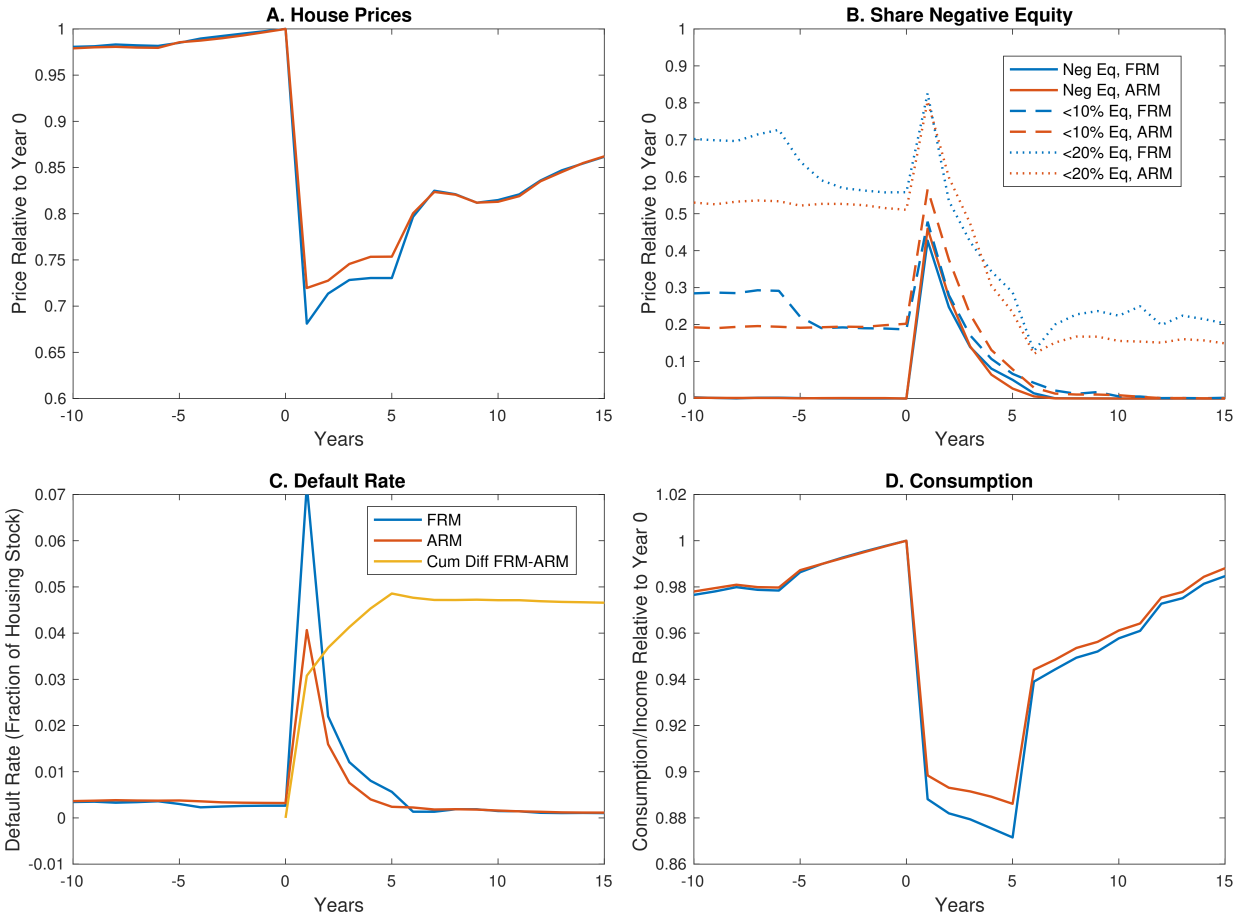

Note: The figure shows the outcomes in a simulated downturn in which the maximum LTV falls from 92.5 percent to 80 percent and there is a five year deep downturn under an FRM and ARM for the case where the central bank both further reduces the short rate by 50 basis points and pursues a policy that subsidizes the long rate by 50 basis points both in a crisis and in regular recessions.

the price-default spiral. In terms of welfare, the crisis still imposes substantial costs on existing homeowners, but aggregate welfare is higher since the monetary policy conveys large benefits to new homebuyers.

Finally, Figure 15 shows the case where the central bank is doveish and pursues a more aggressive monetary policy not only in rare crises but also in more frequent, garden-variety recessions. In this case, the crisis is much worse particularly for FRMs, with larger price declines, more default, and a greater consumption decline for homeowners. Households expect an activist central bank to stabilize the housing market, and they respond by taking on more debt, much like how with ARMs *ex ante* households take on more debt anticipating the hedging benefits of an ARM in a downturn. This highlights that aggressive monetary policy can backfire if it is anticipated and used too frequently.

8 Conclusion

We assess how can mortgages be redesigned or modified in a crisis to reduce housing market volatility, consumption volatility, and default and how mortgage design interacts with monetary policy. To do so, we construct a quantitative equilibrium life cycle model with aggregate shocks in which households have realistic long-term mortgages that are priced by risk-neutral and competitive lenders and household decisions aggregate up to determine house prices. We calibrate the model to match aggregate moments as well as quasi-experimental evidence on the effect of payment size and LTV on default so that our model is tailored to qualitatively assess the benefits of adding simple state contingency to mortgage contracts.

We use the model to assess the performance of various mortgage contracts in a housing crisis that coincides with a recession as well as in normal times. We compare simple adjustable-rate and fixed-rate mortgages to elucidate the economics of adding state contingency to mortgages. We find that adjustable-rate mortgages have hedging benefits because mortgage payments fall if the central bank reduces interest rates in bad states, which relaxes household constraints to help smooth consumption, limits default by relaxing budget constraints in bad states, and stimulates housing demand by new homeowners. These hedging benefits are quite large for constrained, high LTV households who bear the brunt of the housing bust, making the overall welfare benefit to the economy of switching to ARMs large, approximately 12.5 percent of one year of consumption over five years of a housing bust. Crucially, these benefits depend on the extent to which the insurance provided by ARMs is anticipated by households, as households take on more debt when they expect their payments to fall in a crisis, leading to more macro fragility.

We then assess several proposed mortgage designs that add simple state contingency to standard mortgages. These designs are meant to improve on FRMs in the type of crisis that is our main focus while also performing well in normal times and in recessions where the central bank raises rates to fight inflation. We find an FRM with the option to costlessly convert to an ARM provides the best combination of insurance and macroprudential benefits. An FRM that can be refinanced underwater provides the macroprudential benefits but does not help households smooth consumption, while option ARMs provide more insurance but at the cost of macroprudential fragility. We also explore the interaction of mortgage design with monetary policy.

In a future draft we hope to improve the analysis along a number of dimensions. First, we plan to have lenders set spreads separately for each borrower rather than pooling borrower risk. Second, we plan to consider a number of new mortgage designs such as shared appreciation mortgages. Third, we plan to do more to evaluate the interaction of mortgage design with monetary policy by further assessing unconventional policy in a crisis and by assessing the performance of different mortgage designs in a recession where the central bank raises rates to fight inflation.

References

Andersson, F. and T. Mayock (2014). Loss Severities on Residential Real Estate during the Great Recession. *Journal of Banking and Finance 46*, 266–284.

Auclert, A. (2016). Monetary Policy and the Redistribution Channel.

Beraja, M., A. Fuster, E. Hurst, and J. Vavra (2015). Regional Heterogeneity and Monetary Policy.

Bhutta, N., J. Dokko, and H. Shan (2010). The Depth of Negative Equity and Mortgage Default Decisions During the 2007-2009 Housing Bust.

Brueckner, J. K. and K. Lee (2017). Optimal Sharing of Interest-Rate Risk in Mortgage Contracts: The Effects of Pontential Prepayment and Default. *Real Estate Economics*.

Calza, A., L. Stracca, and T. Monacelli (2013). Housing Finance and Monetary Policy. *Journal of the European Economic Association 11*(S1), 101–122.

Campbell, J. Y. (2013). Mortgage Market Design. *Review of Finance 17*(1), 1–33.

Campbell, J. Y. and J. F. Cocco (2003). Household Risk Management and Optimal Mortgage Choice. *Quarterly Journal of Economics 118*(4), 1449–1494.

Campbell, J. Y. and J. F. Cocco (2015). A Model of Mortgage Default. *The Journal of Finance 70*(4), 1495–1554.

Campbell, J. Y., S. Giglio, and P. Pathak (2011). Forced Sales and House Prices. *American Economic Review 101*(5), 2108–31.

Caplin, A., S. Chan, C. Freeman, and J. Tracy (1997). *Housing Partnerships: A New Approach to a Market at a Crossroads*. Cambridge, MA: MIT Press.

Caplin, A., C. Freeman, and J. Tracy (1997). Collateral Damage: Refinancing Constraints and Regional Recessions. *Journal of Money, Credit, and Banking 29*(4), 496–516.

Corbae, D. and E. Quintin (2015). Leverage and the Foreclosure Crisis. *Journal of Political Economy 123*(1), 1–65.

Den Haan, W. J. (2010). Assessing the Accuracy of the Aggregate Law of Motion in Models With Heterogenous Agents. *Journal of Economic Dynamics and Control 34*(1), 79–99.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017). Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging. *Mimeo*.

Eberly, J. and A. Krishnamurthy (2014). Efficient credit policies in a housing debt crisis. *Brookings Papers on Economic Activity*, 73–136.

Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010). What 'Triggers' Mortgage Default? *American Economic Review 100*(2), 490–4.

Farmer, L. and A. A. Toda (2016). Discretizing Nonlinear, Non-Gaussian Markov Processes with Exact Conditional Moments.

Floden, M. and J. Linde (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? *Review of Economic Dynamics 4*, 406–437.

Foote, C. L., K. Gerardi, and P. S. Willen (2008). Negative Equity and Foreclosure: Theory and Evidence. *Journal of Urban Economics 64*(2), 234–245.

Fuster, A. and P. S. Willen (2015). Payment Size, Negative Equity, and Mortgage Default.

Garriga, C., F. Kydland, and R. Sustek (2013). Mortgages and Monetary Policy.

Gerardi, K., K. Herkenhoff, L. Ohanian, and P. S. Willen (2013). Unemployment, Negative Equity, and Strategic Default.

Glaeser, E. L. and J. Gyourko (2009). Arbitrage in Housing Markets. In E. L. Glaeser and J. M. Quigley (Eds.), *Housing Markets and the Economy: Risk, Regulation, and Policy*, pp. 113–148. Cambridge, MA: Lincoln Institute of Land Policy.

Greenwald, D. L. (2015). The Mortgage Credit Channel of Macroeconomic Transmission.

Guren, A. (2017). House Price Momentum and Strategic Complementarity.

Guren, A. and T. McQuade (2015). How Do Foreclosures Exacerbate Housing Downturns?

Guvenen, F. and J. Anthony A. Smith (2014). Inferring Labor Income Risk and Partial Insurance From Economic Choices. *Econometrica 82*(6), 2085–2129.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2016). What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Heathcote, J., K. Storesletten, and G. Violante (2017). Optimal Tax Progressivity: An Analytical Framework. *Quarterly Journal of Economics*.

Kaplan, G., K. Mitman, and G. L. Violante (2015). Consumption and House Prices in the Great Recession: Model Meets Evidence.

Kaplan, G. and G. L. Violante (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica 82*(4), 1199–1239.

Krusell, P. and J. Anthony A. Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Kung, E. (2015). Mortgage Market Institutions and Housing Market Outcomes.

Mian, A. and A. Sufi (2011). House Prices , Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

Mian, A., A. Sufi, and F. Trebbi (2015). Foreclosures, House Prices, and the Real Economy. *Journal of Finance 70*(6), 2587–2634.

Piskorski, T. and A. Tchistyi (2010). Optimal Mortgage Design. *Review of Financial Studies 23*(8), 3098–3140.

Piskorski, T. and A. Tchistyi (2011). Stochastic House Appreciation and Optimal Mortgage Lending. *Review of Financial Studies 24*(5), 1407–1446.

Shiller, R. J. (2008). Understanding Recent Trends in House Prices and Homeownership. In *Housing, Housing Finance, and Monetary Policy: Jackson Hole Conference Series*, pp. 85–123. Federal Reserve of Kansas City.

Wong, A. (2015). Population Aging and the Transmission of Monetary Policy to Consumption.