

Household Leverage and the Recession*

Virgiliu Midrigan[†]

Thomas Philippon[‡]

January 2016

Abstract

A salient feature of the 2007-2010 Great Recession is that states that experienced the largest declines in household debt also experienced the largest contraction in employment. We study an economy in which household liquidity constraints amplify the response of employment to changes in household debt. We estimate the model using data on consumption, employment, wages and household debt in a panel of U.S. states. The model predicts that the 25% decline in U.S. household debt in this period led to a 1.5% drop in the natural rate of interest, far too small compared to the 4.5% short-term interest rate observed in the U.S. at the onset of the recession. Shocks to household debt, on their own, are thus incapable of explaining the large drop in U.S. employment, since they can be offset by monetary policy. The effect of such shocks is amplified, however, if the zero lower bound on nominal interest rates binds for other reasons.

Keywords: Employment, Household Credit, Zero Lower Bound, Great Recession.

JEL classifications: E2, E4, E5, G0, G01.

*We thank Sonia Gilbukh and Callum Jones for excellent research assistance. We also thank Fernando Alvarez, Andy Atkeson, Dave Backus, Patrick Kehoe, Andrea Ferrero, Mark Gertler, Veronica Guerrieri, Ricardo Lagos, Guido Lorenzoni, Robert Lucas, Atif Mian, Tom Sargent, Robert Shimer, Nancy Stokey, Amir Sufi, Ivan Werning, Mike Woodford for comments on this and earlier drafts of this paper.

[†]New York University, virgiliu.midrigan@nyu.edu

[‡]New York University, tphilipp@stern.nyu.edu

1 Introduction

A striking feature of the Great Recession is that U.S. states that have experienced the largest declines in household borrowing have also experienced the largest declines in employment. Figure 1 illustrates this pattern, originally documented in a series of papers by Mian and Sufi, by plotting the change in employment (excluding employment in the construction sector) of individual states against the change in household debt from 2007 to 2010. States like Arizona or Nevada that have experienced changes in debt-to-income of more than 25% have also experienced changes in the employment to population ratio of more than 10% in the 2007-2010 period. In contrast, states like New York or Pennsylvania that have experienced much milder declines in household debt have also experience much smaller reductions in employment.

One interpretation of the evidence in Figure 1 that has received much attention is the *household leverage* view of the recession. According to this view, declines in household debt, caused by a tightening of credit standards or declines in house prices, forced households in the affected states to reduce consumption. Since a large fraction of household spending in a given state is on that state's non-tradable goods, declines in consumer spending led to a reduction in employment owing to price rigidities in the goods and labor markets.¹

Our goal in this paper is to study the quantitative implications of a model that captures this view and is capable of replicating the cross-state evidence on the comovement between household spending, borrowing, employment and house prices. We use the model to ask: what are the aggregate implications of exogenous fluctuations of household debt limits? By how much does employment fall in the aftermath of an exogenous tightening of credit limits that leads to a 25% reduction in household debt, the magnitude observed in the U.S. during the Great Recession?

We study a model in which consumers are liquidity constrained, that is, unable to convert housing wealth into liquid assets that can be used for consumption. Our focus on liquidity constraints is inspired by the work of Kaplan and Violante (2014) who have shown that a sizable fraction of U.S. households are relatively wealthy yet *hand-to-mouth* in that they hold most of their wealth in illiquid assets, such as housing. Unlike Kaplan and Violante (2014) and more recently Kaplan, Mittman, Violante (2015) and Gorea and Midrigan (2015), we do not explicitly study a rich model of household savings in an incomplete-markets economy with transaction costs and multiple assets that differ in their liquidity properties. Our focus, in contrast to that of these papers cast in a small-open economy setting, is in understanding the general equilibrium implications of liquidity constraints in an environment with an

¹See Mian and Sufi (2011, 2014) and Mian, Rao and Sufi (2013) who provide empirical evidence in support of this view.

occasionally binding zero bound on nominal interest rates. Considerations of computational tractability thus lead us to follow the approach of Lucas (1990) in modeling liquidity constraints as arising due to a timing restriction that stipulates that agents must decide how to allocate their wealth between housing and the liquid asset *before* an idiosyncratic shock to preferences is realized. We assume that such shocks are independent and identically distributed across agents and over time. We follow Lucas (1990) in using, however, a family construct in order to eliminate the distributional consequences of asset market incompleteness. The assumptions we make considerably simplify our analysis and characterization of the economy's responses to various shocks in a rich general equilibrium setup.

The Lucas family construct implies that our economy is a representative agent economy. Yet in contrast to economies with complete markets, the quantity of household debt in the economy has important aggregate consequences. Because of the uncertainty about the preference shock of individual members of the family, agents in our environment save for precautionary reasons. In a flexible price variant of the model, the equilibrium interest rate is below the rate of time preference and pinned down by both the strength of this precautionary savings motive as well as by the demand for credit. A tightening of credit leads thus to a reduction in the equilibrium interest rate, yet a negligible drop in consumption or employment. We refer to the equilibrium interest rate in the flexible-price version of our model as the *natural rate*.

In contrast, when prices or wages are sticky, as we assume in our quantitative analysis, the response of real variables to credit shocks depends on the extent to which the monetary authority changes nominal interest rates to ensure that the real interest rate in the economy mimics the dynamics of the natural rate. Absent a lower bound on nominal interest rates, monetary policy in an economy with sticky prices can replicate the dynamics of the flexible-price economy in response to an *aggregate* credit shock, ensuring negligible fluctuations in real variables. Monetary policy cannot react, however, to shocks that are specific to individual states, implying a much greater sensitivity of real variables to credit shocks in a cross-section of states compared to the aggregate time-series.

The zero lower bound on nominal interest rates alters these conclusions, however. If the shocks to credit are sufficiently large, monetary policy may be unable to reduce nominal interest rates sufficiently without violating the zero lower bound. If this is the case, the response of, say, employment to an aggregate credit shock may be much greater than that of an individual state. Intuitively, an individual state can export its way out of the recession, while a closed economy cannot. A key question, therefore, is: how large was the decline in nominal interest rates needed to offset the household credit shocks? Alternatively, by how much did the natural rate of interest decline in response to the tightening of household credit

in the U.S. Great Recession? This paper answers this question by using the cross-sectional evidence from a panel of U.S. states from 2001 to 2012 in order to estimate the key parameters of our model, as well as a Kalman filter to uncover the credit shocks needed to rationalize the dynamics of household debt during this period.

We pin down the model's key parameters by using an indirect inference approach. In particular, we first estimate, in both the model and in the data, auxiliary panel regressions that relate fluctuations of consumption, employment, wages and house prices on one hand, to contemporaneous and lagged changes in household debt on the other hand in the cross-section of U.S. states. As Mian and Sufi have pointed out, changes in household debt are strongly correlated with other state-level variables, so that the explanatory power of these regressions is quite high. We then choose the key parameters, including the persistence of credit shocks, the duration of the long-term securities, the degree of wage rigidity and openness of individual states, by requiring that the coefficients in the auxiliary regressions estimated with data from the model match those estimated using the U.S. state-level data. We show, by bootstrapping our estimates, that the model's parameters are well-identified by the cross-sectional data, with fairly small standard errors around the estimates.

The key parameter in our model is the degree of idiosyncratic uncertainty faced by individual members of the household. This parameter is pinned down by the comovement between consumption and debt in the cross-section of U.S. states. The intuition is as follows. If the amount of idiosyncratic uncertainty is very high, households face severe liquidity constraints and thus save for precautionary reasons. They thus find it costly to reduce the asset side of their balance sheet in order to respond to a tightening of credit. If this is the case, that particular state experiences a *sudden stop*, a sharp increase in its net foreign position that requires a large drop in consumption. If, in contrast, the amount of idiosyncratic uncertainty is low, liquidity constraints and thus the precautionary savings motive is weak and households on the island can simply reduce the asset side of their balance sheet to respond to the tightening of credit. If this is the case, the state's net foreign asset position changes by less, resulting in a milder contraction in consumption. Our model thus captures, in a flexible and parsimonious way, the idea that the sensitivity of the economy to changes in credit limits depends on the extent to which agents are liquidity constrained. In particular, by choosing the degree of idiosyncratic uncertainty appropriately, the model can replicate the comovement between consumption and debt in the state-level data both during the years of the boom as well as the during the bust. In contrast, models that assume large permanent differences in the households' discount factors, such as the patient-impatient model, have a hard time matching the comovement between consumption and debt in the state-level data. In such models impatient agents have no assets and are thus forced to cut their consump-

tion by the full amount of the drop in credit, implying counterfactually large consumption responses.

The degree of idiosyncratic uncertainty in our model has implications not just for an individual state's responses to changes in credit, but also for the response of the natural interest rate to shocks to credit in the aggregate. Intuitively, when idiosyncratic uncertainty is high, agents are strongly liquidity constrained and unwilling to change their savings behavior in response to changes in interest rates. In this case large reductions in the real interest rate are necessary to ensure that the asset market is in equilibrium following a tightening of credit. In this environment monetary policy may not be able to offset credit shocks because of the presence of the zero lower bound constraint. If, in contrast, the degree of demand uncertainty is low, liquidity constraints are weak and agents' savings are quite sensitive to changes in interest rates. In such an environment a mild reduction in real interest rates is necessary to ensure that the asset market is in equilibrium following a tightening of credit, and monetary policy can easily offset a credit shock.

Our main finding is that changes in household debt of the magnitude observed in the Great Recession generate fairly small movements in the natural rate of interest, of about 1.5%, and can, on their own, be easily offset by monetary policy. In our baseline model in which monetary policy follows a Taylor rule and thus imperfectly responds to such shocks, household credit shocks alone generate a fairly mild, 1.4% drop in employment. We stress, however, that our estimates also imply that credit shocks have very persistent effects on the natural rate of interest. For this reason, shocks to household credit would have fairly large effects on aggregate employment if they were accompanied by additional shocks, such as shocks to credit spreads, that trigger the zero lower bound. We conclude, therefore, that shocks to household credit can have sizable effects on real activity in the presence of additional shocks in the economy, but not on their own. Moreover, the persistent nature of household credit shocks can partly account for the slow recovery of U.S. employment in the aftermath of the U.S. Great Recession.

Related Work In addition to the work of Mian and Sufi, most closely related to our paper is the work of Guerrieri and Lorenzoni (2015) and Eggertsson and Krugman (2012) who also study the responses of an economy to a household-level credit crunch. These researchers find, as we do, that a credit crunch has a minor effect on employment if the economy is away from the zero lower bound. Unlike these researchers, who study a closed-economy setting, our model is that of a monetary union composed of a large number of members. Moreover, our main focus is on estimating the model, using cross-sectional evidence, and thus explicitly measuring the mapping from changes in household credit to changes in the natural rate of

interest, which a large literature (see for example Christiano, Eichenbaum and Rebelo (2011)) assumes exogenous.

The view that liquidity constraints can exacerbate the decline in real activity is, of course, not novel to our paper. For example, Lucas and Stokey (2011) have argued that a liquidity crisis “has the effect of reducing the supply available to carry out the normal flow of transactions, leading to a reduction in production and employment.” Our goal in this paper is to evaluate this mechanism using cross-sectional evidence and study its implications for aggregate dynamics.

Methodologically, our emphasis on cross-sectional evidence is also shared by the work of Nakamura and Steinsson (2014). These researchers study the effect of military procurement spending across U.S. regions, and also emphasize the role of cross-sectional evidence in identifying key model parameters. In both our model and theirs differences in the dynamics of employment and other variables *across* states are unaffected by aggregate shocks which are difficult to isolate: for example productivity shocks, changes in monetary policy, or foreign capital flows. As a result, both our and their paper argue, cross-sectional evidence imposes sharp restrictions on the set of parameter values that allow the model to match the data.

Our use of cross-state wage data to estimate the degree of rigidity in the labor market is related to the work of Beraja, Hurst and Ospina (2015) who find, as we do, that wages in individual states comove quite strongly with employment. Matching this evidence implies a fairly steep slope of the Phillips curve in the aggregate, further reinforcing our message that household credit shocks on their own cannot account for the bulk of the recession. If this were the cause, inflation would fall much more in the model than in did in the data.

Our paper is also related to the literature on housing wealth and consumption. An important reference in this literature is Iacoviello (2005) which studies a model in which housing wealth can be used as collateral for loans. In that paper borrowing and lending arise in equilibrium because of differences in the rate of time preference across various agents. In contrast, in our model agents borrow because liquidity constraints give rise to a precautionary savings motive that reduces the interest rate below the rate of time-preference. Thus, in our model individual agents simultaneously save and borrow in equilibrium. In contrast to Iacoviello (2005), in which a tightening of credit forces the impatient agents to reduce consumption one-for-one, in our model agents can respond by reducing the asset side of their balance sheets, thus preventing consumption from falling too much. This feature of our model is important in allowing the model to match the dynamics of consumption in response to the large movements in debt observed in the data.

Our work is also related to a literature that tries to account for the dynamics of house prices in the data. Lustig and van Nieuwerburgh (2005) show that the collateral value of

housing plays an important role in shaping asset returns because a decline in house prices undermines risk sharing and increases the market price of risk. Favilukis, Ludvigson and van Nieuwerburgh (2015) emphasize the role of time-varying risk premia in the recent increase and declines in housing prices. Burnside, Eichenbaum and Rebelo (2015) emphasize heterogeneous expectations about long-run fundamentals and "social dynamics". Compared to these papers, our paper is less concerned with the exact source of house price movements, but rather with their effects on real activity in the aggregate and in the cross-section. An important mechanism in our model is the feedback from lending standards to house prices. Landvoigt, Piazzesi and Schneider (2015) provide evidence consistent with this feedback in a detailed analysis of the housing market of San-Diego. They find that easier access to credit for poor households led to higher house prices at the low end of the housing market. Similarly, Garriga, Manuelli and Peralta-Alva (2014) focus on the role of asset market segmentation in accounting for the large swings in house prices in the data.

Finally, our work is related to a large literature on financial intermediation in both closed and open economies, originating with Bernanke and Gertler (1989), Kiyotaki and Moore (1997), Bernanke, Gertler and Gilchrist (1999) and more recently Mendoza (2010), Gertler and Karadi (2011) and Gertler and Kiyotaki (2010). This literature focuses on understanding the role of shocks that disrupt financial intermediation as well as shocks to the firms' ability to borrow, which we argue must accompany household credit shocks for the model to be able to replicate the large decline in U.S. employment.

2 A Baseline Closed Economy Real Model

We first describe our baseline model of liquidity constraints in a closed-economy without nominal rigidities. We describe how the strength of the precautionary savings motive and household's ability to borrow against the value of their homes interact to determine the equilibrium interest rate in the economy. We also study the dynamics of the key variables in response to a one-time tightening of household credit.

2.1 Setup

We first describe the assumptions we make on technology and preferences, then the nature of securities agents trade and finally the frictions we impose.

2.1.1 Technology and Preferences

The production side of the economy is simple. We assume that competitive firms produce output y_t using labor n_t subject to a constant returns production function

$$y_t = n_t. \quad (1)$$

Competition pins down the real wage: $w_t = 1$. The supply of housing is fixed and normalized to 1 and we let e_t denote the price of housing. The consumption good is the numeraire and we normalize its price to 1.

The representative household has preferences of the form

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int_0^1 v_{it} \log(c_{it}) di + \bar{\eta} \log(h_t) - \frac{1}{1+\nu} n_t^{1+\nu} \right] \quad (2)$$

where h_t is the amount of housing the household owns, which enters preferences with a weight $\bar{\eta}$, n_t is the amount of labor it supplies and c_{it} is the consumption of an individual member i . The term $v_{it} \geq 1$ represents a taste shifter, an i.i.d random variable, which we assume is drawn from a Pareto distribution

$$F(v) = 1 - v^{-\alpha}. \quad (3)$$

Here $\alpha > 1$ is a parameter that determines the amount of uncertainty about v . In particular, the standard deviation of the natural logarithm of v is equal to $1/\alpha$. A lower α implies fatter tails and thus more uncertainty about the taste shifters.

2.1.2 Securities

We assume that the only security traded in this economy is a long-term perpetuity with geometrically decaying coupon payments. The duration of the security is determined by a parameter γ that governs the rate at which the coupon payments decay. A seller of such a security issues each unit of the security at a price q_t in period t and is obligated to repay 1 unit of the consumption good in period $t+1$, γ units in $t+2$, γ^2 in $t+3$ and so on in perpetuity. As we show below, the representative household both borrows and lends using this security.² The household trades this security with perfectly competitive financial intermediaries.

It is convenient to describe a household's financial position by keeping track of the amount of coupon payments b_t that the household must make in period t . Letting l_t denote the amount of securities the household sells in period t , the date $t+1$ coupon payments are

$$b_{t+1} = \sum_{i=0}^{\infty} \gamma^i l_{t-i} = l_t + \gamma b_t \quad (4)$$

²See Hatchondo and Martinez (2009) and Arellano and Ramanarayanan (2012) who describe the properties of these securities in more detail.

Similarly, we let a_t denote the amount of coupon payments the household is entitled to receive in period t .

2.1.3 Budget and Borrowing Constraints

We let x_t be the amount of funds the household transfers to the goods market. Since individual members are ex-ante identical and of measure 1, x_t is also the amount of funds any individual member has available for consumption when entering the goods market. We assume that each member's consumption is limited by the amount of funds it has available:

$$c_{it} \leq x_t \tag{5}$$

We refer to the constraint in (5) as the *liquidity constraint*. The household's flow budget constraint states that

$$x_t + e_t(h_{t+1} - h_t) = w_t n_t + q_t l_t - b_t + (1 + \gamma q_t) a_t \tag{6}$$

This says that the amount of resources the household has available for consumption x_t and housing purchases, $e_t(h_{t+1} - h_t)$, is limited by the amount of labor income it earns in that period, $w_t n_t$; the amount it receives from selling l_t units of the long-term security at price q_t , net of the required coupon payments b_t ; as well as the market value of the a_t securities it owns. Each unit of the security the household owns pays off one unit in coupon payments and can be sold at a price γq_t reflecting the geometric decay of the payments.

We assume a *borrowing constraint* that limits the household's ability to issue *new* loans. In particular, the face value of the new loans the value issues is limited to be below a multiple m_t of the value of the household's end of period housing stock:

$$q_t l_t \leq m_t e_t h_{t+1}. \tag{7}$$

We assume that the parameter governing the credit limit, m_t , follows an AR(1) process and is the only source of aggregate uncertainty in this baseline version of the model:

$$\log m_t = (1 - \rho) \log \bar{m} + \rho \log m_{t-1} + \varepsilon_t, \tag{8}$$

where ε_t is a normal random variable. Shocks to m_t generate variation in the amount individual households are able to borrow over time.

Notice that our specification the borrowing limit restricts a household's ability to take on new loans, not its choice of total debt b_{t+1} . We make this assumption in order to capture the idea that a tightening of the credit limit precludes agents from taking on new loans, but does not force prepayment of old debt. Had we assumed a limit on the stock of debt, a tightening of credit limits would force agents to deleverage immediately rather than gradually, which would be counterfactual.

2.1.4 Savings

Individual households in this economy simultaneously borrow and save using the long-term security. A household's savings are simply the unspent funds of its shoppers in the goods market. The total amount of securities a household purchases at the end of the shopping period is simply

$$a_{t+1} = \frac{1}{q_t} \left(x_t - \int_0^1 c_{it} di \right), \quad (9)$$

where recall that q_t is the price of one unit of the long-term security.

2.1.5 Timing

We conclude the description of the model by illustrating, in Figure 2, the timing assumptions we make. The household enters the period with a_t units of savings, h_t units of housing and b_t units of debt. The uncertainty about the collateral limit m_t is realized at the beginning of the period. The household then chooses how much to work n_t , how much housing to purchase h_{t+1} , how much to borrow b_{t+1} , and how much to transfer to each individual member in the goods market x_t . After the transfer to the goods market is complete, each individual members' preference for consumption v_{it} is realized and the individual members purchase c_{it} units of the consumption good. At the end of the shopping period all unspent funds are pooled by the household to purchase a_{t+1} units of the long-term security.

2.2 Decision Rules

The household's problem is to choose c_{it} , x_t , h_{t+1} , b_{t+1} and n_t in order to maximize the household's life-time utility in (2) subject to the liquidity constraint in (5), the flow budget constraint in (6), the borrowing constraint in (7) and the law of motion for the household's savings in (9). We capture the timing assumption that transfers x_t are chosen prior to the realization of the idiosyncratic preference shock v_{it} with the measurability restriction that x_t is the same for all household members i .

Let μ_t denote the shadow value of wealth, that is, the multiplier on the flow budget constraint (6), ξ_{it} denote the multiplier on the liquidity constraint (5) and λ_t denote the multiplier on the borrowing constraint (7). Finally, let R_{t+1} denote the ex-post realized return of the long-term security:

$$R_{t+1} = \frac{1 + \gamma q_{t+1}}{q_t}. \quad (10)$$

The first-order condition that determines x_t is then

$$\mu_t = \beta \mathbb{E}_t \mu_{t+1} R_{t+1} + \int_0^1 \xi_{it} di, \quad (11)$$

where \mathbb{E}_t is the conditional mathematical expectation operator. Since the loan-to-value limit m_t is the only source of aggregate uncertainty, \mathbb{E}_t is simply the expectation operator over the realization of the credit shock ε_t .

This expression is quite intuitive. One additional unit of a transfer x_t is valued at μ_t , the shadow value of wealth in period t . Since unspent funds can be used to purchase long-term assets, the transfer provides a return R_{t+1} in the following period and is valued at $\beta\mu_{t+1}R_{t+1}$. In addition, transfers to the goods market provide a liquidity service by relaxing the liquidity constraint of individual members. Since transfers are chosen prior to the realization of the taste shock, the liquidity services provided by transfers are given by the expected value of the multiplier of the liquidity constraint of individual members. The second term on the right hand side of (11) is thus the expectation operator over the realization of the idiosyncratic taste shifter v .

Consider next the household's choice of how much to borrow. The first-order condition for b_{t+1} is

$$\mu_t = \beta\mathbb{E}_t\mu_{t+1}R_{t+1} + \lambda_t - \beta\gamma\mathbb{E}_t\lambda_{t+1}\frac{q_{t+1}}{q_t}, \quad (12)$$

where recall that λ_t is the multiplier on the borrowing constraint. This expression is intuitive as well. The benefit to borrowing an additional unit is equal to the shadow value of wealth μ_t and the cost of doing so is next period's repayment, valued at $\beta\mu_{t+1}R_{t+1}$. Borrowing an extra unit has an additional cost, by tightening the borrowing constraint (λ_t), but also a benefit, as it reduces the multiplier on the borrowing constraint in period $t + 1$. The last term in this expression simply reflects the long-term nature of securities and our assumption that the credit limit applies to new, rather than old debt.

Consider next the choice of housing. The first-order condition is given by

$$e_t\mu_t - \beta\mathbb{E}_t\mu_{t+1}e_{t+1} = \beta\mathbb{E}_t\frac{\bar{\eta}}{h_{t+1}} + \lambda_t m_t e_t. \quad (13)$$

The left hand side of this expression is the user cost of housing: the difference between the purchase price and next period's selling price, appropriately discounted. The right hand side is the marginal utility of an additional unit of housing services $\frac{\bar{\eta}}{h_{t+1}}$ as well as the collateral value of housing $\lambda_t m_t e_t$.

We finally discuss the optimal choice of consumption of individual members. Given log preferences and the fact that unspent funds are valued at $\beta\mathbb{E}_t\mu_{t+1}R_{t+1}$, the choice of consumption is simply

$$c_{it} = \min \left[\frac{v_{it}}{\beta\mathbb{E}_t\mu_{t+1}R_{t+1}}, x_t \right], \quad (14)$$

reflecting the possibility of a binding liquidity constraint.

2.3 Equilibrium

The equilibrium is characterized by a sequence of prices e_t , w_t , q_t and allocations such that firms and households optimize and the housing, labor and asset markets clear. The asset market clearing condition is simply

$$a_{t+1} = b_{t+1}, \quad (15)$$

which also implies that the goods market clears, $c_t = y_t = n_t$. The supply of labor is given by the household's first-order condition for employment,

$$n_t^\nu = \mu_t w_t. \quad (16)$$

Recall that firm optimization implies $w_t = 1$ and that the housing stock is in fixed supply, normalize to one.

2.4 The Workings of the Model

We next provide some discussion about the workings of the model. Let

$$\underline{c}_t = \frac{1}{\beta \mathbb{E}_t \mu_{t+1} R_{t+1}} \quad (17)$$

denote the consumption of a member with the lowest realization of the demand shock, $v_{it} = 1$. (We assume restrictions on the parameter space that ensure that the liquidity constraint does not bind for such members.) The multiplier on an individual member's liquidity constraint, ξ_{it} , is simply the difference between that member's marginal utility of consumption and the valuation of unspent wealth:

$$\xi_{it} = \max \left[\frac{v_{it}}{c_{it}} - \beta \mathbb{E}_t \mu_{t+1} R_{t+1}, \quad 0 \right] \quad (18)$$

Integrating across individual members and using the assumption that v is Pareto-distributed gives

$$\underline{c}_t \int_0^1 \xi_{it} di = \frac{1}{\alpha - 1} \left(\frac{x_t}{\underline{c}_t} \right)^{-\alpha}, \quad (19)$$

which says that the expected multiplier on the liquidity constraint is proportional to the fraction of liquidity constrained members (with $v_{it} > x_t / \underline{c}_t$), that is $\left(\frac{x_t}{\underline{c}_t} \right)^{-\alpha}$.

Finally, combining (19) with (11) gives:

$$\frac{1}{\alpha - 1} \left(\frac{x_t}{\underline{c}_t} \right)^{-\alpha} = \left(\beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} R_{t+1} \right)^{-1} - 1 \approx \rho_t - r_t, \quad (20)$$

where $\rho_t = -\log \beta \mathbb{E}_t \left(\frac{\mu_{t+1}}{\mu_t} \right)$ is the subjective discount rate and $r_t = \log \mathbb{E}_t(R_{t+1})$ is the interest rate. Intuitively, the right-hand side of (20) is equal (up to a first-order approximation)

to the difference between the discount rate and the interest rate, while the left-hand size is proportional to the fraction of constrained household members. As the gap between the discount rate and the interest rate increases, it becomes costlier for households to save, transfers fall relative to consumption, so more members end up constrained. Conversely, as the gap between the discount rate and the interest rate falls to 0, so does the fraction of constrained household members.

Consider next the household's total consumption expenditures, $c_t = \int_0^1 c_{it} di$. Using the fact that unconstrained members consume $c_{it} = v_{it} \underline{c}_t$, while constrained members consume $c_{it} = x_t$, we have

$$\frac{c_t}{\underline{c}_t} = \frac{\alpha}{\alpha - 1} \left(1 - \frac{1}{\alpha} \left(\frac{x_t}{\underline{c}_t} \right)^{1-\alpha} \right) \quad (21)$$

This expression is intuitive as well. On one extreme, as the gap between the discount rate and interest rate falls to 0, the ratio x_t/\underline{c}_t goes to infinity, and so the mean/min consumption ratio reaches $\alpha/(\alpha - 1)$, the mean of the taste shocks v under the Pareto distribution. On the other extreme, as x_t/\underline{c}_t goes to 1, all members have identical consumption so that the mean/min ratio is equal to 1.

Finally, letting $\Delta_t = \left(\beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} R_{t+1} \right)^{-1} - 1$ denote the gap between the discount rate and the interest rate, the savings to consumption ratio can be written as:

$$\frac{q_t a_{t+1}}{c_t} = \frac{x_t - c_t}{c_t} = \left(\frac{\alpha}{\alpha - 1} [(\alpha - 1) \Delta_t]^{\frac{1}{\alpha}} - \Delta_t \right)^{-1} - 1, \quad (22)$$

which can be shown is increasing as $\Delta_t \approx \rho_t - r_t$ decreases and is steeper the higher Δ_t is.

Consider now the household's decision of how much to borrow. A comparison of (19) and (20) makes it clear that the gap Δ_t between the discount rate and the interest rate is positive as long as the ratio x_t/\underline{c}_t is finite, as is the expected multiplier on the liquidity constraint. Further, a comparison of (19) and (12) reveals that the multiplier on the borrowing constraint, λ_t , is then positive as well. Since $\underline{c}_t > 0$ and x_t is bounded by an agent's initial financial and housing wealth, labor income and ability to borrow, the borrowing constraint binds in equilibrium whenever the loan-to-value ratio m_t is bounded, as is the case in our application.

2.5 Steady State Equilibrium Interest Rate

Consider next how the equilibrium interest rate is determined in the steady state of our model with a constant credit limit $m_t = \bar{m}$. As discussed above, the household's asset position, $q_t a_{t+1} = x_t - c_t$ is increasing in the equilibrium interest rate and asymptotes as the interest rate approaches the rate of time preference.

Consider next the relationship between debt and interest rates. Since the amount of debt is constrained by the borrowing limit $qb = \frac{1}{1-\gamma}\bar{m}eh$ and the housing stock is constant, the amount of debt in the economy is proportional to the price of houses. The price of houses reflects both the marginal valuation of housing, as captured by the preference parameter $\bar{\eta}$, as well as the collateral value of housing. The latter declines as interest rates increase – higher interest rates make borrowing less attractive. To see this, notice that in the steady state the Euler equation for housing (13) reduces to

$$e = \frac{1}{\mu} \frac{\bar{\eta}}{\bar{h}} \frac{1}{\rho - \frac{\bar{m}}{1-\beta\gamma}(\rho - r)}, \quad (23)$$

where \bar{h} is the fixed stock of housing. Clearly, when \bar{m} is positive a higher interest rate reduces the price of housing and thus the amount the household can borrow.

Figures 3a and 3b illustrate these points. Figure 3a assumes a relatively large degree of idiosyncratic uncertainty (low α) about the preference shocks. Notice how the intersection of the upward-sloping asset supply curve and the downward-sloping debt curve determines the equilibrium interest rate. A tightening of the debt limit reduces the demand for debt, thus reducing the interest rate, potentially below 0. Figure 3b assumes a relatively low degree of idiosyncratic demand uncertainty. In this case agents save less since the precautionary savings motive is reduced and the equilibrium interest rate is therefore higher. Moreover, the intersection of the asset and debt curves occurs now at a point at which the asset supply curve is relatively flat, implying that a given decline in the debt limit is associated with a smaller reduction in the equilibrium interest rate.

2.6 Impulse Response to a Credit Shock

Figure 4 reports the baseline's economy impulse responses to a one-time negative shock to credit, ε_t . Since loan-to-value ratio m_t follows a persistent autoregressive process, a one-time innovation to this process leads to a persistent decline in the household's ability to borrow. Recall, however, that we assume that the collateral constraint limits a consumer's ability to take on new loans. A decline in the loan-to-value ratio thus leads to a gradual reduction in total household debt. The latter evolves according to

$$q_t b_{t+1} = \gamma q_t b_t + m_t e_t,$$

and thus contracts gradually in response to the credit shock. Figure 4 contrasts the response of debt, interest rates, output and house prices in economies with a relatively high and low degree of idiosyncratic uncertainty, as captured by the tail parameter α . To make things comparable, we choose values for the discount factor in each of the two economies so as to ensure that the steady interest rate is equal to 2% in both.

Notice that the equilibrium interest rate falls in both economies in response to a tightening of credit. Since the shock reduces households' ability to borrow, the interest rate must decline so as to ensure that total amount of household assets falls as well. As discussed above, the interest rate falls more in the economy with greater demand uncertainty, reflecting the stronger precautionary savings motive and thus steeper savings curve. In contrast, output (and thus consumption and employment) barely fall in response to the shock, by only 0.05%. Although a tightening of credit increases the consumption-leisure distortions, these distortions are quantitatively small here, as is the case in the flexible-price cash-in-advance model (see Cooley and Hansen (1991)). Finally, notice that the response of house prices to a tightening of credit is ambiguous. Although a reduction of the loan-to-value ratio reduces the collateral value of housing, the reduction in interest rates increases it, thus resulting in an ambiguous effect of credit shocks on the price of houses.

3 An Island Monetary Economy with Price Rigidities

We next embed the setup described above into a richer monetary economy with price and wage rigidities amenable to quantitative analysis. The economy is composed of a continuum of ex-ante identical islands of measure 1 that belong to a monetary union and trade among themselves. Consumers on each island derive utility from the consumption of a final good, leisure and housing. The final good is assembled using inputs of traded and non-traded goods which themselves are composites of varieties of differentiated intermediate products. We assume that intermediate goods producers are monopolistically competitive. In addition, we assume that individual households on each island are organized in unions that sell differentiated varieties of labor to intermediate goods producers on each island. Both intermediate goods prices and union wages are subject to a Calvo-type adjustment friction. Finally, we assume that labor is immobile across islands and that the housing stock on each island is in fixed supply.

3.1 Setup

We start by describing the assumptions we make on preferences and technology, and then on how prices are determined.

3.1.1 Household Problem

The representative household on each island has preferences identical to those described in the previous section. We let s index an individual island and $p_t(s)$ denote the price of the final consumption good on the island. We assume perfect risk-sharing across households

belonging to different labor unions on a given island. Because of separability in preferences, risk-sharing implies that all households on an island make identical consumption, housing and savings choices, even though their labor supply differs depending on when the union that represents them was last able to reset its wage. The problem of an individual household that belongs to labor union z is to

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int_0^1 v_{it}(s) \log(c_{it}(s)) di + \bar{\eta} \log(h_t(s)) - \frac{1}{1+\nu} n_t(z, s)^{1+\nu} \right] \quad (24)$$

subject to the budget constraint

$$p_t(s)x_t(s) + e_t(s)(h_{t+1}(s) - h_t(s)) = w_t(z, s)n_t(z, s) + q_t l_t(s) - b_t(s) + (1 + \gamma q_t)a_t(s) + T_t(z, s), \quad (25)$$

where $T_t(z, s)$ collects the profits households earn from their ownership of intermediate goods firms, transfers from the government aimed at correcting the steady state markup distortion, as well as the transfers stemming from the risk-sharing arrangement. We assume that households on island s exclusively own all the firms on that particular island and have implicitly imposed the fact that consumption, housing and savings choices of households belonging to different unions are identical.

As earlier, the household also faces a liquidity constraint limiting the consumption of an individual member to be below the amount of real balances the member holds:

$$c_{it}(s) \leq x_t(s), \quad (26)$$

a borrowing constraint

$$q_t l_t(s) \leq m_t(s)e_t(s)h_{t+1}(s), \quad (27)$$

and the law of motion for a household's assets is given by

$$q_t a_{t+1}(s) = p_t(s) \left(x_t(s) - \int_0^1 c_{it}(s) di \right). \quad (28)$$

Implicit in this formulation is our assumption that there are no barriers to capital flows across islands, implying that the price of the long-term security q_t is common to all islands.

3.1.2 Final Goods Producers

Final goods producers on island s produce $y_t(s)$ units of the final good using $y_t^N(s)$ units of the intermediate non-tradable good produced on the island and $y_t^T(s, j)$ units of tradable goods produced in island j :

$$y_t(s) = \left(\omega^{\frac{1}{\sigma}} y_t^N(s)^{\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} \left(\int_0^1 y_t^T(s, j)^{\frac{\kappa-1}{\kappa}} dj \right)^{\frac{\kappa}{\kappa-1} \frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (29)$$

where ω determines the share of non-traded goods, σ is the elasticity of substitution between traded and non-traded goods and κ is the elasticity of substitution between varieties of the traded goods produced on different islands. We assume that the law of one price holds across locations. That is, final goods producers on all islands face a price $p_t^T(j)$ for the traded intermediate input sourced from island j . Letting $p_t^N(s)$ denote the price of the traded goods on island s , the final goods price on an island is

$$p_t(s) = \left(\omega p_t^N(s)^{1-\sigma} + (1-\omega) \left(\int_0^1 p_t^T(j)^{1-\kappa} dj \right)^{\frac{1-\sigma}{1-\kappa}} \right)^{\frac{1}{1-\sigma}}. \quad (30)$$

The demand for non-tradable intermediate goods produced on an island is

$$y_t^N(s) = \omega \left(\frac{p_t^N(s)}{p_t(s)} \right)^{-\sigma} y_t(s), \quad (31)$$

while demand for an island's tradable goods is an aggregate of what all other islands purchase:

$$y_t^T(s) = (1-\omega) p_t^T(s)^{-\kappa} \left(\int_0^1 p_t^T(j)^{1-\kappa} dj \right)^{\frac{\kappa-\sigma}{1-\kappa}} \left(\int_0^1 p_t(j)^\sigma y_t(j) dj \right). \quad (32)$$

3.1.3 Intermediate Goods Producers

Traded and non-traded goods on each island are themselves composites of varieties of differentiated intermediate inputs. The elasticity of substitution between such varieties is equal to ϑ in both the tradable and nontradable sectors. We thus have, for example,

$$y_t^T(s) = \left(\int_0^1 y_t^T(z, s)^{\frac{\vartheta-1}{\vartheta}} dz \right)^{\frac{\vartheta}{\vartheta-1}}, \quad (33)$$

and similarly for non-traded goods. Demand for an individual producer's variety is therefore $y_t^T(z, s) = (p_t^T(z, s)/p_t^T(s))^{-\vartheta} y_t^T(s)$.

Individual producers of intermediate goods are subject to Calvo-type price adjustment frictions with a constant adjustment hazard. Let λ_p denote the probability that a firm does not reset its price in a given period. Alternatively, the firm is allowed to reset its price with probability $1 - \lambda_p$. The problem of a firm that resets its price is to maximize the present discounted flow of profits weighted by the probability that the price it chooses at t , $p_t^{T*}(s)$ will still be in effect at any particular date in the future. As earlier, the production function is linear in labor so that the unit cost of production is simply the island's wage $w_t(s)$.

For example, a traded intermediate goods firm that resets its price solves

$$\max_{p_t^{T*}(s)} \sum_{k=0}^{\infty} (\lambda_p \beta)^k \mu_{t+k}(s) (p_t^{T*}(s) - \tau_p w_t(s)) \left(\frac{p_t^{T*}(s)}{p_t^T(s)} \right)^{-\vartheta} y_t^T(s), \quad (34)$$

where $\mu_{t+k}(s)$ is the shadow value of wealth of the representative household on island s , that is, the multiplier on the flow budget constraint (25) and $\tau_p = \frac{\vartheta-1}{\vartheta}$ is a tax the government levies to eliminate the steady state markup distortion. This tax is rebated lump sum to households on island s .

The composite price of traded or non-traded goods is then a weighted average of the prices of individual differentiated intermediates. For example, the composite price of traded goods evolves according to

$$p_t^T(s) = \left((1 - \lambda_p) p_t^{T*}(s)^{1-\vartheta} + \lambda_p p_{t-1}^T(s)^{1-\vartheta} \right)^{\frac{1}{1-\vartheta}}, \quad (35)$$

where we have used the fact that the adjustment hazard is constant.

3.1.4 Wage Setting

We assume that individual households are organized in unions that supply differentiated varieties of labor. The total amount of labor services available in production is

$$n_t(s) = \left(\int_0^1 n_t(z, s)^{\frac{\psi-1}{\psi}} dz \right)^{\frac{\psi}{\psi-1}}, \quad (36)$$

where ψ is the elasticity of substitution. Demand for an individual union's labor services given its wage $w_t(z, s)$ is therefore $n_t(z, s) = (w_t(z, s)/w_t(s))^{-\psi} n_t(s)$. The problem of a union that resets its wage is therefore to

$$\max_{w_t^*(s)} \sum_{k=0}^{\infty} (\lambda_w \beta)^k \left[\tau_w \mu_{t+s} w_t^*(s) \left(\frac{w_t^*(s)}{w_t(s)} \right)^{-\psi} n_t(s) - \frac{1}{1+\nu} \left(\left(\frac{w_t^*(s)}{w_t(s)} \right)^{-\psi} n_t(s) \right)^{1+\nu} \right], \quad (37)$$

where λ_w is the probability that a given union leaves its wage unchanged and $\tau_w = (\psi-1)/\psi$ is a labor income subsidy aimed at correcting the steady state markup distortion. The composite wage at which labor services are sold to producers is

$$w_t(s) = \left((1 - \lambda_w) w_t^*(s)^{1-\psi} + \lambda_w w_{t-1}(s)^{1-\psi} \right)^{\frac{1}{1-\psi}}. \quad (38)$$

The elasticity of substitution ψ across varieties of labor services is a key parameter that determines the extent to which individual wages respond to credit shocks in this model. To see this, log-linearize the optimal choice of reset wages that solves (37) around the steady-state without aggregate or island-level shocks:

$$\hat{w}_t^*(s) = \beta \lambda_w \mathbb{E}_t \hat{w}_{t+1}^*(s) + \frac{1 - \beta \lambda_w}{1 + \psi \nu} (-\hat{\mu}_t(s) + \psi \nu \hat{w}_t(s) + \nu \hat{n}_t(s)), \quad (39)$$

where hats denote log-deviations from the steady state. This shows that the term $\psi \nu$ dampens the elasticity of reset wages to changes in, say, the shadow value of wealth, $\mu_t(s)$. Although

workers would like to respond to an increase in μ_t by reducing wages and thus supplying more hours, they are less inclined to do so when the elasticity of substitution between labor varieties ψ is high or when the inverse of the Frisch elasticity of labor supply ν is high. Intuitively, if the elasticity of substitution ψ is high, cutting wages would lead to a large increase in the amount of labor supplied by an individual union and thus its members' disutility from work. Thus, for a given frequency of wage changes, as determined by λ_w , a higher ψ dampens the response of wages to shocks.

3.1.5 Island Equilibrium

As earlier, the supply of housing is fixed and normalized to 1:

$$h_{t+1}(s) = 1. \quad (40)$$

The total amount of the composite labor service is used by producers of both tradable and non-tradable goods:

$$n_t(s) = \int_0^1 \left(\frac{p_t^N(z, s)}{p_t^N(s)} \right)^{-\vartheta} y_t^N(s) dz + \int_0^1 \left(\frac{p_t^T(z, s)}{p_t^T(s)} \right)^{-\vartheta} y_t^T(s) dz, \quad (41)$$

where $y_t^N(s)$ and $y_t^T(s)$ are given by (31) and (32).

The agents' consumption savings choices are identical to those described earlier. For example, the minimum consumption level is equal to

$$\underline{c}_t(s) = \frac{1}{\beta \mathbb{E}_t \mu_{t+1}(s) R_{t+1}} \frac{1}{p_t(s)}, \quad (42)$$

where recall that $p_t(s)$ is the price of the final good on the island. The choice of transfers $x_t(s)$ is identical to that in (11) above, while total household consumption is given by (21) as earlier. Finally, total consumption on a given island must equal to the total amount of the final good produced:

$$c_t(s) = y_t(s), \quad (43)$$

since the latter, unlike the intermediate traded inputs, is non-tradable.

As discussed above, we assume free flows of capital across islands so that the amount agents on an individual island save, $q_t a_{t+1}(s)$, is not necessarily equal to the amount they borrow, $q_t b_{t+1}(s)$. The island's net asset position evolves according to:

$$q_t(a_{t+1}(s) - b_{t+1}(s)) = (1 + \gamma q_t)(a_t(s) - b_t(s)) + w_t(s)n_t(s) + T_t(s) - p_t(s)c_t(s), \quad (44)$$

where we have aggregated across all members of the various unions and used the fact that the composite wage index satisfies $w_t(s)n_t(s) = \int_0^1 w_t(z, s)n_t(z, s)dz$. In words, an island's net asset position increases if wage income and profits received by individual agents on the island exceeds the amount they consume.

3.2 Monetary Policy

Let $y_t = \int_0^1 \frac{p_t(s)}{p_t} y_t(s) ds$ be total real output in this economy, where $p_t = \int_0^1 p_t(s) ds$ is the aggregate price index. Let $\pi_t = p_t/p_{t-1}$ denote the rate of inflation and

$$1 + i_t = \mathbb{E}_t R_{t+1} \tag{45}$$

be the expected nominal return on the long-term security, which we refer to as the nominal interest rate. We assume that monetary policy is characterized by a Taylor-type interest rate rule subject to a zero lower bound:

$$1 + i_t = \max \left[(1 + i_{t-1})^{\alpha_r} \left[(1 + \bar{i}) \pi_t^{\alpha_\pi} \left(\frac{y_t}{\bar{y}} \right)^{\alpha_y} \left(\frac{y_t}{y_{t-1}} \right)^{\alpha_x} \right]^{1-\alpha_r}, 1 \right],$$

where α_r determines the persistence of the interest rate rule, while α_π , α_y and α_x determine the extent to which monetary policy responds to inflation, deviations of output from its steady state level, and output growth, respectively. We assume that \bar{i} is set to a level that ensures a steady state level of inflation of $\bar{\pi}$.³ Notice that because an individual island is of measure zero, monetary policy does not react to island-specific disturbances. Also recall that the monetary union is closed so aggregate savings must equal to the aggregate debt:

$$\int_0^1 a_{t+1}(s) ds = \int_0^1 b_{t+1}(s) ds \tag{46}$$

3.3 Source of Shocks

For our quantification, we introduce shocks to housing preferences in addition to credit shocks. We do so because, as is well known, credit shocks alone cannot generate movements in house prices in this class of models nearly as large as those observed in the data. We thus assume shocks to both the loan-to-value ratio as well as the consumer's preference for housing.⁴ In particular, we modify the utility function in (24) to introduce time-varying weights on housing in preferences, $\eta_t(s)$.

Specifically, we now assume that the loan-to-value ratio on each island, $m_t(s)$, follows an autoregressive process:

$$\log m_t(s) = (1 - \rho) \log \bar{m} + \rho \log m_{t-1}(s) + \varepsilon_t(s), \tag{47}$$

³We assume in our quantitative analysis that $\bar{\pi}$ is equal to 2% per year. We eliminate the steady-state costs of positive inflation by assuming that prices and wages are automatically indexed to $\bar{\pi}$. See Coibion and Gorodnichenko (2014) and Blanco (2015) for a discussion of the size of these costs.

⁴See Kiyotaki, Michaelides and Nikolov (2011) for an illustration of the problem and Garriga, Manuelli and Peralta-Alva (2014) and Favilukis, Ludvigson and Van Nieuwerburgh (2015) for approaches to resolve it.

as does the preference weight on housing:

$$\log \eta_t(s) = (1 - \rho) \log \bar{\eta} + \rho \log \eta_{t-1}(s) + \sigma_\eta \varepsilon_t(s). \quad (48)$$

For simplicity, we assume that these two processes have the same persistence ρ and are driven by a single disturbance $\varepsilon_t(s)$. Thus, periods in which the loan-to-value ratio is lower are also periods in which the demand and thus the price of houses falls, further restricting agents' ability to borrow. The parameter σ_h governs the relative variability of the housing preference shocks. We continue to refer to the shocks $\varepsilon_t(s)$ as *credit shocks*, since changes in both housing preferences (and thus house prices) as well as changes in the loan-to-value ratio only affect island and economy-wide variables through their effect on the amount of debt households can take on.

3.4 Impulse Response to an Island-Level Credit Tightening

We next illustrate the workings of this richer model by reporting how an individual island responds to an island-specific credit shock that reduces the households' ability to borrow. We start by discussing the responses in an economy with flexible prices and wages and then those in an economy with price and wage stickiness.

Figure 5a shows the responses in a flexible price economy in which the degree of demand uncertainty is relatively high. The upper-left panel shows that debt contracts gradually, but in contrast to a closed economy, the asset holdings of agents on the island do not fall nearly as much. The island's net foreign asset position thus increases. Financing this increase in the net foreign asset position requires that agents on the island spend less than they earn, which, given that leisure and consumption are normal goods, leads to a decline in consumption, an increase in employment, as well as a drop in wages. Intuitively, wages must fall on the island so as to induce the rest of the economy to buy more of that particular island's traded goods.

Figure 5b shows that the responses of all these variables are muted in the economy with lower demand uncertainty. The intuition is as follows. An island can respond to a tightening of credit in two ways: either by reducing its asset position or by cutting consumption and leisure. When demand uncertainty is low, it is relatively costless to reduce transfers to individual members of the household and so an island's assets fall nearly as much as its debt does. Both sides of the island's balance sheet thus contract, with little impact on other variables. In contrast, when demand uncertainty is high, reducing assets is costly since individual shoppers are more likely to end up liquidity constrained. The representative household on the island thus finds it optimal to respond to the credit tightening by simultaneously cutting consumption, leisure as well as the amount transferred to each individual member. The

model thus captures, in a very parsimonious way, the idea that responding to credit shocks is more costly when the precautionary savings motive is stronger.

We next explain the role of price and wage rigidities in shaping the response of an individual island to a credit shock. Figure 5c reports the economy’s response to a credit shock in an economy with price rigidities and compares it to those in an economy with flexible prices. The upper row of the figure shows that both wages and prices react much more gradually to the credit shock when prices are sticky, while consumption responds much more. Moreover, employment falls now, both because non-tradable employment experiences a bigger decline, as well as because the increase in tradable employment that would otherwise occur with flexible prices is now muted. The reason for these results is quite intuitive. Wage rigidities in this environment act like a tax on labor supply while price rigidities lead to an increase in firm markups and thus reduce real wages. Both of these forces prevent employment from increasing following a credit tightening.⁵ In fact, since a large fraction of an island’s consumption is on locally produced non-tradable goods, the large reduction in consumption associated with the credit tightening is now associated with a decline in non-tradable employment which, due to wage and price rigidities, is no longer offset by an increase in tradable employment. Consequently, employment on the island falls, more so when a larger fraction of spending is on non-tradable goods.

3.5 Impulse Response to an Aggregate Credit Tightening

We conclude the description of the setup of the model by describing the responses to an economy-wide credit shock under the assumption that monetary policy follows a Taylor rule. As we have shown earlier, when prices are flexible a credit tightening in the aggregate leads to a reduction in real interest rates and essentially no change in other real variables, such as employment. As is well known, a policy of strict inflation targeting would mimic the flexible price responses even in the presence of price and wage rigidities. Such a policy would ensure that the real interest rate in the sticky price economy mimics that in the flexible price economy, the latter often referred to as the *natural interest rate*, which we denote by r_t^n .

When prices are sticky and monetary policy does not fully offset the credit shock by a policy of strict inflation targeting, the real interest rate is too high compared to that under the flexible price allocations. If this is the case, consumption falls, as Figure 5d shows, as does output and employment — all these variables respond in an identical way to an aggregate credit shock up to a first-order approximation. The drop in consumption is larger, the greater

⁵See Kehoe, Midrigan and Pastorino (2016) for cross-sectional evidence from the U.S. Great Recession that both of these margins account for the drop in employment in states that have experience the largest declines in household credit.

the gap between the real interest rate and its natural level. Interestingly, even though we have assumed a unitary elasticity of intertemporal substitution, the decline in consumption in our model is not equal to the expected value of the sum of future interest rate gaps, $\mathbb{E}_t \sum_j (r_{t+j}^n - r_{t+j})$, as it would be the case in simpler versions of the New Keynesian models in which shocks to r_t^n are assumed exogenous. This point is illustrated by the upper-right panel of Figure 5d which shows that this sum falls by about twice more than consumption does.

The lower panels of Figure 5d provide some additional description of the mechanics behind the drop in consumption in the economy with price rigidities. Recall from (21) that aggregate consumption is the product of two terms: the minimum amount \underline{c}_t consumed by the $v = 1$ members of the household, times the mean-min consumption ratio. The minimum consumption level is given by

$$\underline{c}_t = \frac{1}{\mu_t} \frac{1}{\beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} \frac{R_{t+1}}{\pi_{t+1}}}, \quad (49)$$

where recall that μ_t is the shadow value of wealth (the multiplier on the budget constrained) and π_{t+1} is the inflation rate. Similarly, the mean-min consumption ratio is given by

$$\text{mean-min consumption ratio} = \frac{\alpha}{\alpha - 1} \left(1 - \frac{1}{\alpha} \left(\frac{x_t}{\underline{c}_t} \right)^{1-\alpha} \right).$$

The lower panel of Figure 5d shows that with sticky prices μ_t increases in response to the credit shock, thus leading agents to reduce the minimum level of consumption. This effect is amplified by the tightening of liquidity constraints associated with the lower real interest rates which make it optimal for households to reduce the amount transferred to individual members, x_t/\underline{c}_t , thus reducing the mean-min consumption ratio. This latter effect is more modest, however, as shown by the lower-right panel of Figure 5d, and accounts for only about a quarter of the overall drop in consumption. Absent price rigidities, the shadow value of wealth would actually fall slightly, owing to the greater liquidity frictions, and the sharper decline in real interest rates would lead households to increase the amount consumed by the $v = 1$ members of the household: the lower-center panel of Figure 5d shows that in this case the minimum consumption actually increases, by about 0.5%. This effect is essentially offset by the tightening of liquidity constraints, which reduce the mean-min consumption ratio, thus leading to a very small drop of consumption in the aggregate.

To conclude, consumption (and thus output and employment) fall in our model mostly because of a sharp increase in the shadow value of wealth μ_t , rather than due to a tightening of liquidity constraints per se. In turn, the shadow value of wealth μ_t satisfies

$$(1 + \Delta_t) \beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} \frac{R_{t+1}}{\pi_{t+1}} = 1,$$

an Δ_t is implicitly a function of the amount of credit available, as shown by (22). The workings of the model are thus fairly similar to those of the simple New Keynesian model, in which $\mu_t = 1/c_t$. The key difference is that in our setup the mapping from changes in credit to disturbances in this Euler equation and thus consumption and output responses is a more complicated one and depends on the details of the parameterization. In particular, the response of output to a given reduction in household debt depends critically on the strength of the precautionary-savings motive, as determined by the uncertainty parameter α . We thus proceed next to estimating this, and other parameters of the model, using state-level data.

4 Quantification and Steady State Implications

We next describe how we have chosen parameters for our model and discuss the model's steady state implications. We assume that all islands are identical in the initial steady state, which we associate with 2001 in the data. We first discuss the parameters we assign based on existing evidence or steady-state considerations, and then the parameters we choose using an indirect inference approach and the state-level panel data from 2001 to 2012.

4.1 Assigned Parameters

The period is one quarter. We assume a Frisch elasticity of labor supply of 1/2 and thus set ν equal to 2. We assume that prices and wages are reset on average once a year, so we set λ_p and λ_w , the hazards of not adjusting, equal to 0.75. We follow the trade literature in setting the elasticity of substitution between traded and non-traded goods, σ , equal to 0.5, and that between varieties of traded goods produced in different islands, κ , equal to 1.5. We use the Justiniano and Primiceri (2008) estimates of the parameters characterizing the Taylor rule. These choices are reported in the left column of Panel C of Table 1.

We have three additional parameters that are pinned down by steady state considerations. The discount factor β is chosen so that the steady state interest rate is equal to 2% per year. The weight of housing in preferences $\bar{\eta}$ is chosen so that the aggregate housing to income ratio is equal to 2.5, a number that we compute using the 2001 Survey of Consumer Finances (SCF). Finally, the steady state loan-to-value ratio is chosen so that in the aggregate debt to housing ratio is equal to 0.29, a number once again computed from the SCF. Since the debt constraint binds in the equilibrium of our model, these two last two targets imply an aggregate debt to income ratio of $2.5 \times 0.29 = 0.725$. Panel C of Table 1 reports these parameters as well.

4.2 Parameters Chosen Using Indirect Inference

We have six additional parameters that determine the dynamic responses to a credit shock: the Pareto tail of the distribution of idiosyncratic preference shocks, α ; the persistence of coupon payments, γ ; the persistence of the shocks ρ ; the relative volatility of housing preferences, σ_h ; the elasticity of substitution between labor varieties, ψ , as well as the weight on non-tradable goods in the utility function ω . We choose these parameters using panel information on the comovement of household debt, consumption, employment and wages in the cross-section of U.S. states from 2001 to 2012. We discuss our indirect inference approach next.

Indirect inference is based on estimating an auxiliary model in both the theory and the data, and choosing parameters in the theory so as to ensure that the estimates of the auxiliary model are as close as possible in the theory and in the data.⁶ Our auxiliary model is a set of panel regressions using state-level data on household debt from the FRB New York Consumer Credit Panel,⁷ house price data from the FHFA, as well as data on employment, wages and consumption expenditures from the BEA.⁸ Specifically, we estimate panel regressions of the form

$$\ln Y_t^i(s) = d^i(s) + f_t^i + \chi_1^i \text{Debt}_t(s) + \chi_2^i \text{Debt}_{t-1}(s) + u_t^i(s) \quad (50)$$

in both our theory and in the data, where Y^i is one of four series: 1) the employment-population ratio, 2) wages (employee compensation divided by total employment), 3) per-capita consumption spending and 4) house prices; d_s is state-specific fixed effect f_t is a time effect and Debt_{st} is the amount of household debt in an individual state scaled by that state's 2001 income. We weigh individual states in these regressions by the state's population.

Our choice of this particular auxiliary model is motivated by the observation of Mian and Sufi that there was a strong positive correlation between household debt and state-level aggregates during the boom and bust in house prices around the Great Recession. To be clear, the regressions in (50) are not meant to capture any particular causal relationship, but rather the dynamic pattern of the correlations between these variables. As the second-to-last column of Panel B of Table 1 illustrates, variation in debt is indeed strongly correlated with employment, consumption, wages and house prices across states and over time: the R^2 in these regressions ranges from 0.56 for wages to about 0.70 for consumption and employment and 0.87 for house prices. Moreover, most of this variation is not driven by an aggregate component common to all states. As the last column of Panel B of Table 1 shows, not

⁶See, for example, Guvenen and Smith (2014).

⁷We include credit card debt, auto loans and student loans, in addition to mortgage debt in our measure of household credit.

⁸See the Appendix for a more detailed description of the data we use.

including the time effects f_t in these regressions reduces the R^2 in these regressions by only 0.15 to 0.20. Thus, a large fraction of the variation in state-level aggregates is associated with variation in debt across states. Figure 6a illustrates this pattern for subset of states in the data by plotting the actual employment-population ratio in each state against the fitted values from the estimates of (50). The two track each other closely.

To estimate the auxiliary model in the theory, we choose the path for credit shocks $\varepsilon_t(s)$ for each period and state in our theory so that the model matches the path for debt in the data perfectly. Recall that a period in our theory is one quarter, while the state-level data is sampled once a year. For this reason we first linearly interpolate the debt series in the data in order to back out a path for shocks for each quarter, but then estimate the auxiliary model at the original annual frequency. Thus, the independent variables on the right hand side of (50) are the same in both our theory and in the data. We then choose parameters in the theory to ensure that the fitted values produced by the auxiliary regressions in the theory are as close as possible to those in the data. In particular, let $\hat{y}_t^{i,\text{theory}}(s) = \chi_1^{i,\text{theory}} \text{Debt}_t(s) + \chi_2^{i,\text{theory}} \text{Debt}_{t-1}(s)$ be the fitted values of each of the four series in the theory and similarly let $\hat{y}_t^{i,\text{data}}(s)$ be the fitted values in the data. Our objective is then to choose the six parameters of the theory to minimize

$$\sum_{i=1}^4 \sum_{s=1}^{50} \sum_{t=2002}^{2012} \left(\hat{y}_t^{i,\text{theory}}(s) - \hat{y}_t^{i,\text{data}}(s) \right)^2. \quad (51)$$

Panels A and B of Table 1 illustrates how well we do in matching the patterns in the data. Notice that the model is over-identified (recall that we have six parameters in our theory to match eight coefficients in the auxiliary regression), yet the theory does a reasonably good job of matching the coefficients and thus fitted values in the data. Panel B shows that the sum of squared errors between the fitted values in the theory and in the data is equal to 0.04 of the total variance of the employment and consumption series, 0.02 of the wage series and 0.13 of the house price series. The correlation between the fitted values in the model and the data is nearly one, as is the relative standard deviation of variables in the model relative to the data – the model understates the volatility of consumption and house prices by about 10%, and overstates that of wages by about 5%. Figure 6b illustrates the strong fit of the model by contrasting the fitted values in the model and in the data for a subset of the states. Overall, we conclude that our model provides a successful account of the correlation between household debt, consumption, employment and wages across U.S. states in the period surrounding the Great Recession.

Panel C of Table 1 reports the parameter values needed to achieve this fit. The value of the Pareto tail parameter is equal to 5.5, implying a standard deviation of the logarithm of v equal to $1/\alpha = 0.18$. The parameter governing the duration of long-term securities, γ is

equal to 0.953, implying a duration of the security of about 20 quarters.⁹ This is shorter than the duration of a 30-year mortgage (about 13 years), but we prefer to directly estimate this parameter, rather impose a particular duration, since households in the U.S. have the option to prepay their mortgages, as well as borrow using much shorter maturity home equity lines of credit and credit card loans etc.

Panel C of Table 1 also shows that shocks to credit are fairly persistent, with an AR(1) coefficient of 0.76, and that shocks to housing preferences are much more volatile than changes in the loan-to-value ratio itself – the value of σ_h is equal to 7.49. Intuitively, although debt and house prices have fluctuated dramatically in the cross-section and in the time series, the debt to housing ratio is relatively stable, as pointed out by Justiniano, Primiceri and Tambalotti (2015), and the model thus requires fairly small changes in the loan to value ratio to account for the patterns in the data.

Notice also that the model estimates a fairly strong degree of *real rigidity*, in that the elasticity of substitution between labor varieties is fairly high, $\psi = 5.4$, implying that wages move fairly slow in response to credit shocks. Intuitively, this is necessary to account for the fact that wages in the data are about one-half and one-third as volatile as employment and consumption, respectively (see Panel A of Table 1). Finally, the model implies a fairly high degree of home bias in preferences, with a weight on non-tradables equal to $\omega = 0.87$. This is in line with the findings in the international macro literature that non-tradable distribution services are an important component of the cost of retailing even highly tradable goods.¹⁰

The right panel of Panel C of Table 1 also reports the standard errors around these estimates, which we have computed by resampling the individual states in the panel with replacement. Notice that the standard errors are fairly low, suggesting that the parameters are well-identified.

We next provide some intuition for what features of the data pin down individual parameters of our model by perturbing each of the six parameters and reporting how the model’s fit changes. We report these experiments in Table 2.

We first reduce α to 4, thus increasing the volatility of the taste shocks. Table 2 shows that all state-level aggregates are now much more volatile, with employment volatility increasing most (by about 80%), while house price volatility increasing least (about 40%). In contrast, the correlation between the model and data-produced series continues to be fairly high, in excess of 0.9. Thus, increasing the uncertainty about taste shocks makes state-level variables more sensitive to credit shocks. To understand this result, recall our discussion of the impulse

⁹We follow Hatchondo and Martinez (2009) in defining duration as the weighted average maturity of cash flows. This is given by $\frac{1}{q} \sum_{t=1}^{\infty} t \left(\frac{1}{1+r} \right)^t \gamma^{t-1} = \frac{1+r}{1+r-\gamma}$.

¹⁰See the evidence in Burstein, Eichenbaum and Rebelo (2005).

responses in Figures 5a and 5b. When uncertainty about taste shocks is high, the household finds it costly to reduce its level of assets in order to respond to a tightening of credit, and thus its consumption falls by more. Due to price and wage rigidities, more volatile state-level consumption fluctuations translate into more volatile employment and wage fluctuations as well.

We next increase γ to 0.975, thus increasing the duration of securities to 32 quarters. Table 2 shows that this perturbation again increases the volatility of state-level variables, more so for house prices (which are now 60% more volatile) than real variables which are now 20-30% more volatile. Intuitively, a greater γ implies that a current credit tightening will have effects on the household's ability to borrow far into the future, thus leading to a greater reduction in the collateral value of housing on impact. Importantly, increasing the duration of securities counterfactually reduces the correlation between house prices in the model and in the data. When γ is relatively high, house prices fall much more rapidly in response to a credit tightening than debt does, changing the correlation between house prices and debt away from that observed in the data. Increasing the persistence of shocks, ρ , has a qualitatively similar impact as increasing γ and to conserve on space we do not report the results of these experiments here.

We next reduce the elasticity of substitution of labor varieties, ψ , to 3. We now see that wages are about 25% more volatile compared to the Benchmark model or the data, while employment is about 10% less volatile. Clearly, this parameter is identified based on the relative variability of wages and employment in the data.

Finally, we reduce the share of non-tradable goods in preferences to $\omega = 0.75$, thus decreasing the fraction of employment in the non-tradable goods sectors. We now find that employment, consumption and wages are all much less volatile (about 1/3 less volatile than in the data or the Benchmark) model. Intuitively, the larger the fraction of the workforce that produces tradable goods, the more insulated a state is from state-specific credit shocks.

4.3 Steady State Implications

We conclude this section by briefly discussing the key steady-state predictions of the model. Given the degree of demand uncertainty we estimate, the (annualized) gap between the discount rate and the interest rate is about 1.5%. The ratio of transfers to minimum consumption, x/\underline{c} is equal to about 2.1, implying that $(x/\underline{c})^\alpha = 1.7\%$ are liquidity constrained in any given period. Finally, the implicit employment tax levied by the liquidity constraints is not very large: output and employment would increase by only 0.09% in the absence of the liquidity constraints. Overall, we conclude that the estimates of the model imply that the steady state distortions arising from the liquidity constraints are not too large, and these

constraints mainly manifest themselves in a 1.5% reduction in the annualized real interest rate compared to the rate of time preference.

5 Aggregate Implications

We next study the model's aggregate implications. Our focus is, in particular, on understanding the dynamics of real variables in the 2008-2009 Great Recession. We proceed in three steps. First, we describe the impulse responses of various variables to a one-time credit shock using the parameter values we have estimated in the previous section. Second, we report the model's predictions when we confront it with a path for credit shocks that allow the model to exactly reproduce the time-series of the U.S. household debt to income. We show that the model predicts that the observed decline in household credit during the Great Recession led to a fairly modest decline in the natural rate of interest, of about 1.5%, and was thus incapable, on its own, to trigger the zero lower bound on interest rates. We finally study the interaction between household credit shocks and an additional shock that triggers the zero lower bound and show that in the presence of this additional shock the effect of household credit on real variables is considerably amplified, owing to the fact that credit shocks have quite persistent effects on the natural rate of interest in our model.

5.1 Impulse Responses to a One-Time Credit Shock

Figure 7a reports the impulse responses to a one-time decline in ε_t common to all islands. As earlier, the fact that the credit limit m_t follows a persistent autoregressive process as well as the fact that the credit limit applies only to new debt issues, implies that debt falls gradually and then recovers slowly as well. This is shown by the solid line in the upper-left panel of Figure 7a which scales the value of household debt $q_t b_{t+1}$ by the steady state income and reports deviations from the steady state. In this experiment we have chosen the size of the credit shock to ensure a maximal drop in debt of 0.25 of steady state income, roughly the drop in household debt observed in the U.S. during the 2007-2010 period.

The solid line in the lower-left panel of Figure 7a shows the response of output (and, up to a first-order approximation, consumption and employment), expressed as % deviations from the steady state. These real variables fall immediately by about 1.5% in response to the 25% decline in household credit, and gradually mean-revert.

The right panels of Figure 7 show the response of the nominal interest rate and inflation. The nominal interest rate falls gradually, by about 2% (annualized) at its trough and follows a path similar to that of the debt series. Inflation, in contrast, falls immediately, by about 1.5% (annualized), thus by almost as much as output. Our parameter estimates thus imply a

fairly steep slope of the Phillips curve. Moreover, they imply that interest rates fall by about 2%, far too small compared to the drop in interest rates observed in the U.S. data, a point to which we return below.

Figure 7a also contrasts the responses in our Benchmark model, estimated using the U.S. state level data, with the predictions of an economy in which prices and wages are permanently fixed at their steady-state values, an assumption often made in the recent literature. Notice that the response of output is now much larger (about 3.5% at its trough which occurs about 12 quarters after the shock) and much more persistent: output returns to only about 2.8% below its steady state level even 30 quarters after the shock and only gradually recovers. Thus, with permanently fixed prices the model predicts what essentially amounts to a *secular stagnation* – a long-term decline in both output and interest rates below their long-run trends.

Figure 7b provides some intuition for why the economy with greater price stickiness predicts a larger and more persistent response of output to household credit shocks. The figure reports the response of the real interest rate in our Benchmark model, the economy with fixed prices, as well as the natural rate of interest. Recall from Figure 5d that if monetary policy were to mimic the latter, inflation would not react to a credit shock in our model, while output would fall by very little. Figure 7b shows that although the real interest rate falls by more initially in the economy with fixed prices, it recovers much more quickly than the real interest rate in our Benchmark model and is thus higher than the natural rate of interest starting from about four years (16 quarters) after the shock. This is a direct consequence of the particular parameterization of the Taylor rule that we have used, one that puts a relatively high weight on inflation and relatively little on output, thus implying a less responsive Taylor rule in the fixed price economy. Since output in the model depends on the entire expected path of real interest rates, the anticipation of persistently high real interest rates (relative to the natural rate) in the fixed price economy reduces output in a persistent fashion.

Figure 7c compares the responses of aggregate variables after an aggregate credit shock to those of island variables after a credit shock that only affects an individual island. We choose the size of the shocks in the two experiments in order to ensure that debt falls, at its trough, by the same amount, 25% of steady state income. Notice that the initial drop in both consumption and employment is more than twice greater at the island level than it is in the aggregate. This owes to the fact that monetary policy can react to a shock in the aggregate but, by assumption, does not react to an idiosyncratic shock to an individual island. Thus, even though the cross-sectional comovement between consumption, employment and household credit allows us to identify the key parameters of the model,

we cannot simply extrapolate the state-level correlations to inform about the role of credit in generating fluctuations in the aggregate. Notice also that the recession in the aggregate is much more persistent than that of an individual island, even though the credit shock is equally persistent. Intuitively, an individual island's consumption is pinned down by the *change* in the amount households can borrow, while the aggregate consumption responses are pinned down by the *level* of credit since it is the latter that pins down the natural rate of interest.

Overall, we conclude that our estimated model implies fairly modest responses of real variables to a shock that forces households to reduce credit. Even though the magnitude of the debt contraction we have considered here is fairly large, 25% of steady state income, the decline in output is equal to only about 1.5%. Even this decline in output can be avoided: the drop in the natural rate of interest caused by such a credit shock is equal to about 1.3%, too modest to trigger the zero lower bound constraint on nominal interest rates.

We emphasize, however, that the decline in the natural rate of interest predicted by our model is, though small, fairly persistent. Thus, if household credit shocks are accompanied by other shocks that reduce the Fed's ability to cut interest rates, the resulting effects on output can be much greater. We illustrate this point by considering next a shock that drives a wedge between the federal funds rate f_t and the nominal interest rate i_t at which agents borrow:

$$i_t = f_t + \xi_t,$$

where the wedge itself follows an autoregressive process:

$$\xi_t = \rho_\xi \xi_{t-1} + v_t.$$

This wedge can arise due to frictions in the financial markets and was indeed highly elevated during the U.S. Great Recession. We modify the Taylor rule by assuming that the federal funds rate reacts to changes in such spreads and fully offsets them whenever possible. Thus, we assume that

$$1 + f_t = \max \left[(1 + f_{t-1} + \xi_{t-1})^{\alpha_r} \left[(1 + \bar{i}) \pi_t^{\alpha_\pi} \left(\frac{y_t}{\bar{y}} \right)^{\alpha_y} \left(\frac{y_t}{y_{t-1}} \right)^{\alpha_x} \right]^{1-\alpha_r} - \xi_t, \quad 1 \right],$$

Figure 7d illustrates the economy's responses to a 2.5% increase in the credit spread that mean-reverts gradually. Absent additional shocks, monetary policy fully offset the increase in the spread by reducing its target rate f_t (shown in the upper-right panel of the figure), resulting in no effects on output and inflation. In the presence of the household credit shock (chosen, as earlier, to generate a 25% drop in debt relative to steady state income), monetary

policy can no longer fully respond as it is constrained by the zero lower bound for about 10 quarters. Consequently, output falls much more now, by about 5%, much more than the 1.5% drop in the absence of the increase in the spread. Notice, however, that inflation falls quite a bit now as well, to -1%, owing to the steep slope of the Phillips curve implied by the state-level evidence. Our findings thus echo those of Beraja, Hurst and Ospina (2015) who conclude, using an alternative VAR-based methodology, that “the wage stickiness necessary to get demand shocks to be the primary cause of aggregate employment declines during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.”

5.2 Matching Household Debt Dynamics in the Great Recession

We next explicitly study the role of changes in household credit in accounting for the drop in output and employment during the Great Recession. We do so by choosing a path for the aggregate credit shocks in our economy, ε_t , that ensure that the model reproduces the time-paths for the U.S. household debt to income ratio.

The left panel of Figure 8 shows that the household debt to income ratio exhibits a trend, starting from about 0.5 in 1975 to about 1 in the last decade. Since we do not allow for trends in our model, we detrend the data by subtracting a linear trend. The resulting series is shown in the right panel of Figure 8. We smooth this series to eliminate high frequency noise, by projecting it on a cubic spline of order 15 — the smoothed series is reported with dotted lines in the figure. We finally use a Kalman filter in order to back out the path for credit shock that the model requires to match the smoothed detrended debt-to-income series and study the model’s implications for various macro variables starting from 2001 to 2014.

Figure 9 presents the model’s responses to the sequence of credit shocks ε_t uncovered with the Kalman filter. The nearly 25% boom and bust in household debt in our model is accompanied by a fairly modest rise and then drop in interest rates, from about 4% in 2001 to 5% at the peak in 2007 to about 3% by 2013, much more modest than that observed in the data in this period.

Consider next the response of employment, which we measure in the data as the total number of employees on non-farm payrolls, scaled by U.S. population. (We detrend the log of the resulting series using a linear trend.) As the lower-left panel of Figure 9 shows, the Great Recession was associated with an extremely persistent nearly 7.5% drop in the employment-population ratio. The model, in contrast, predicts a much more modest 1.4% drop, thus accounting for less than 20% of the decline in employment observed in the data. Moreover, this drop in employment is gradual, mimicking the gradual reduction in debt, as opposed to sudden as observed from 2007 to 2009 in the data. Finally, notice in the lower-right panel of

Figure 9 that inflation in our model is fairly volatile, almost as much as employment.

Overall, these results echo those based on the impulse response analysis. We find, as earlier, that the changes in debt to income in our model are insufficient, on their own, to generate the large reduction in employment and interest rates observed in the data.

As discussed earlier, if household credit shocks are also accompanied by other shocks that reduce the Fed's ability to cut interest rates, the resulting effects on employment can be much greater. We illustrate this point in Figure 10 in which we consider an experiment in which a persistent shock to the credit spread ξ_t reduces the federal funds rate to 0 in the second quarter of 2008. As earlier, such a shock would, on its own, have no effect on real variables since it can be fully offset by monetary policy. In contrast, when such a shock is also accompanied by a shock to household credit, the resulting decline in employment can be quite large. As Figure 10 shows, the drop in employment is now much greater, about 3.7%, thus half of the drop in the data. As earlier, the inflation puzzle persists, with inflation now falling below 0.5%, a drop much greater than that observed in the data.

The dynamics of the economy at the zero lower bound is quite sensitive to what exactly one assumes about the size and persistence of the additional shocks that trigger the zero lower bound, as well as the extent to which the Fed uses forward guidance to increase expectations of future output and inflation. For example, raising the persistence of the shock to spreads can dramatically increase the decline in employment, while assuming that the Fed commits to keeping interest rates at zero past the date justified by the Taylor rule would offset the decline. We thus find it useful, and more informative, to summarize our results in terms of the decline in the natural rate of interest implied by the path of household debt observed in the data. We find that the natural rate of interest has declined by about 1.5% and did so in a very persistent fashion. Although on its own such a drop implies fairly mild declines in employment, it can be quite large in the present of other shocks given the non-linear dynamics of the economy at the zero lower bound.

We conclude by illustrating, in Figure 11, the usefulness of using state-level data to inform about the parameters of our model. In this figure we report the results of an experiment in which we increase the degree of demand uncertainty by reducing α to 4.25 (compared to 5.5) in the Benchmark model. We also assume a greater degree of prices stickiness, by increasing the elasticity of substitution ψ to 100, thus dampening the extent to which reset wages respond to aggregate shocks. The figure shows that in such an environment both employment and interest rates are much more sensitive to household credit shocks than they are in our Benchmark model. The peak to trough drop in employment is equal to about 5%, thus about three-quarters than that observed in the data, interest rates drop by about 4%, while the drop in inflation is, not surprisingly, much more modest. As we have argued earlier,

however, such a parameterization is inconsistent with the state-level data as it implies a much greater sensitivity of consumption and employment to changes in debt in the cross-section, and much weaker volatility of state-level wages. Moreover, even though this version of the model generates sufficiently volatile fluctuations in employment, it has trouble capturing the timing of the employment declines, owing to the fact that household debt in the U.S. data has fallen much more gradually than employment.

6 Conclusions

A popular account of the U.S. Great Recession is the view that declines in housing wealth and the households' ability to borrow have led to a reduction in consumption and employment due to price rigidities and constraints on monetary policy. This view is motivated, in part, by the observation that employment comoves strongly with changes in house prices and household debt in the cross-section of U.S. states. This paper proposes a theory that captures this view by stressing the role of liquidity constraints in generating a precautionary-savings motive, giving rise to borrowing and lending in equilibrium. We estimate the key parameters of our model using an indirect inference approach that requires the model to match the salient cross-sectional features of the data. We then study the model's implications for the aggregates. Our model predicts that changes in household debt of the magnitude observed in the Great Recession generate fairly small movements in the natural rate of interest, of about 1.5%, and can, on their own, be easily offset by monetary policy even in the present of a zero lower bound constraint. As a caveat, the decline in the natural rate of interest is quite persistent, however, so shocks to household credit can have quite powerful effects in the presence of additional shocks, such as shocks to spreads, that trigger the zero lower bound. We conclude, therefore, that shocks to household credit, on their own, were not potent enough to have accounted for the bulk of the contraction in employment during the Great Recession. Yet such shocks can have sizable effects on real activity in the presence of additional shocks in the economy, and, due to their persistent nature, can partly account for the slow recovery of U.S. employment in the aftermath of the financial crisis.

References

- [1] Beraja, Martin, Erik Hurst and Juan Ospina. 2015. "The Aggregate Implications of Regional Business Cycles," Chicago Booth Working Paper.
- [2] Bernanke, Ben and Mark Gertler. 1989 "Agency Costs, Net Worth and Business Fluctuations," *American Economic Review*, 79, 14-31.

- [3] Bernanke, Ben, Mark Gertler and Simon Gilchrist. 1999. "The financial accelerator in a quantitative business cycle framework," in Taylor, John B. and Michael Woodford, ed., *Handbook of Macroeconomics* vol. 1C.
- [4] Blanco, Andres. 2015. "Optimal Inflation Target in an Economy with Menu Costs and an Occasionally Binding Zero Lower Bound," University of Michigan Working Paper.
- [5] Burnside, Craig, Martin Eichenbaum and Sergio Rebelo. 2015. "Understanding Booms and Busts in Housing Prices." Northwestern University Working Paper.
- [6] Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4), 742-784.
- [7] Christiano, Lawrence, Martin Eichenbaum and Sergio Rebelo. 2011. "When is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119(1) 78-121.
- [8] Coibion, Olivier, Yuriy Gorodnichenko and Johannes Wieland. 2012. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise their Inflation Targets in Light of the ZLB?" *Review of Economic Studies*, 79, 1371-1406.
- [9] Cooley, Thomas and Gary Hansen. 1991. "The Inflation Tax in a Real Business Cycle Model," *American Economic Review*, 79 (4), 733-748.
- [10] Eggertsson, Gauti, and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics*, 127(3) 1469-1513.
- [11] Favilukis, Jack, Sydney Ludvigson and Stijn Van Nieuwerburgh. 2015. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," forthcoming, *Journal of Political Economy*.
- [12] Garriga, Carlos, Rodolfo Manuelli and Adrian Peralta-Alva. 2014. "A Macroeconomic Model of Price Swings in the Housing Market," Federal Reserve Bank of Saint Louis Working Paper.
- [13] Gertler, Mark and Nobuhiro Kiyotaki. 2010. "Financial Intermediation and Credit Policy in Business Cycle Analysis," in Friedman, Benjamin and Michael Woodford, ed., *Handbook of Monetary Economics*, vol 3, 547-599.
- [14] Gertler, Mark and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 58(1), 17-34.

- [15] Gorea, Denis and Virgiliu Midrigan. 2015. “Liquidity Constraints in the U.S. Housing Market,” NYU Working Paper.
- [16] Guerrieri, Veronica, and Guido Lorenzoni. 2015. “Credit Crises, Precautionary Savings, and the Liquidity Trap,” Chicago Booth Working Paper.
- [17] Hatchondo, Juan Carlos and Leonardo Martinez. 2009. “Long-Duration Bonds and Sovereign Defaults,” *Journal of International Economics*, 79(1), 117-125.
- [18] Guvenen, Fatih and Anthony Smith. 2014. “Inferring Labor Income Risk and Partial Insurance from Economic Choices,” *Econometrica*, 82(6), 2085-2129.
- [19] Iacoviello, Matteo. 2005. “House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle,” *American Economic Review*, 95(3): 739-764.
- [20] Justiniano, Alejandro and Giorgio Primiceri. 2008. “The Time Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98(3), 604-641.
- [21] Justiniano, Alejandro, Giorgio Primiceri and Andrea Tambalotti. 2015. “Household Leveraging and Deleveraging,” *Review of Economic Dynamics*, 18(1), 3-20.
- [22] Kaplan, Greg, and Gianluca Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4) 1119-1239.
- [23] Kaplan, Greg, Kurt Mitman, and Gianluca Violante. 2015. “Consumption and House Prices in the Great Recession: Model Meets Evidence,” NYU Working Paper.
- [24] Kehoe, Patrick, Virgiliu Midrigan and Elena Pastorino. 2016. “Debt Constraints and the Labor Wedge,” forthcoming, *American Economic Review Papers & Proceedings*.
- [25] Kiyotaki, Nobuhiro, Alexander Michaelides and Kalin Nikolov. 2011. “Winners and Losers in Housing Markets,” *Journal of Money, Credit and Banking*, 43(2-3), 255-296.
- [26] Landvoigt, Tim, Monika Piazzesi and Martin Schneider. 2015. “The Housing Market(s) of San Diego,” *American Economic Review*, 105(4), 1371-1407.
- [27] Lucas, Robert Jr., 1990. “Liquidity and Interest Rates,” *Journal of Economic Theory*, 50(2) 237-264.
- [28] Lucas, Robert Jr. and Nancy Stokey, 2011. “Liquidity Crises.” Federal Reserve Bank of Minneapolis Economic Policy Papers.

- [29] Lustig, Hanno and Van Nieuwerburgh, Stijn. 2005. "Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective," *Journal of Finance*, 60(3), 1167-1219.
- [30] Mendoza, Enrique 2010. "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100(5), 1941-1946.
- [31] Mian, Atif, and Amir Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis," *American Economic Review*, 101(5), 2132-2156.
- [32] Mian, Atif, Kamalesh Rao and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128(4), 1687-1726.
- [33] Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Quarterly Journal of Economics*, 82(6), 2197-2223.
- [34] Nakamura, Emi and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions." *American Economic Review*, 104(3), 753-792.

Table 1: Parameterization

Panel A: Coefficients in Panel Regressions

	Data	Model
log employment on current debt	0.18	0.18
log employment on lagged debt	-0.15	0.15
log consumption on current debt	0.30	0.30
log consumption on lagged debt	-0.20	-0.21
log wages on current debt	0.09	0.11
log wages on lagged debt	-0.04	-0.05
log house prices on current debt	1.94	1.84
log house prices on lagged debt	-1.40	-1.51

Panel B: Distance between Theory and Data

	Distance	Correlation	Relative std. dev.	R^2 Data	R^2 Data no f_t
Employment	0.04	0.98	1.00	0.69	0.55
Consumption	0.04	0.98	0.91	0.71	0.56
Wages	0.02	0.99	1.05	0.56	0.35
House Prices	0.13	0.93	0.88	0.87	0.71

Panel C: Parameter Values

<i>Assigned</i>			<i>Chosen with indirect inference</i>			
	1 quarter	Period length	α	5.50	(0.41)	Tail Pareto taste shocks
ν	2	Inverse labor supply elasticity	γ	0.95	(0.01)	Persistence coupon payments
λ_p	0.75	Calvo price stickiness	ρ	0.76	(0.05)	AR(1) shocks
λ_w	0.75	Calvo wage stickiness	σ_h	7.49	(1.37)	Volatility housing shocks
σ	0.5	Elasticity traded/non-traded	ψ	5.39	(1.91)	Elasticity labor varieties
κ	1.5	Elasticity traded goods	ω	0.87	(0.07)	Weight on non-traded goods
α_r	0.86	Taylor rule persistence				
α_π	1.71	Taylor coefficient inflation	β	0.991		Discount factor
α_y	0.05	Taylor coefficient output	$\bar{\eta}$	0.024		Weight on housing in preferences
α_x	0.21	Taylor coefficient output growth	\bar{m}	0.014		Steady state LTV

Table 2: Identification of Key Parameters

	Data	Benchmark	$\alpha = 4$	$\gamma = 0.975$	$\psi = 3$	$\omega = 0.75$
log employment on current debt	0.18	0.18	0.32	0.22	0.16	0.12
log employment on lagged debt	-0.15	-0.15	-0.28	-0.20	-0.14	-0.10
log consumption on current debt	0.30	0.30	0.55	0.37	0.30	0.25
log consumption on lagged debt	-0.20	-0.21	-0.42	-0.27	-0.21	-0.19
log wages on current debt	0.09	0.11	0.20	0.14	0.14	0.08
log wages on lagged debt	-0.04	-0.05	-0.13	-0.06	-0.08	-0.05
log house prices on current debt	1.94	1.84	2.97	3.17	1.84	1.81
log house prices on lagged debt	-1.40	-1.51	-2.46	-2.83	-1.51	-1.50
relative std dev. employment (model/data)	-	1.00	1.80	1.27	0.89	0.68
relative std dev. consumption (model/data)	-	0.91	1.56	1.12	0.91	0.70
relative std dev. wages (model/data)	-	1.05	1.59	1.49	1.26	0.63
relative std dev. house prices (model/data)	-	0.88	1.43	1.59	0.88	0.87
correlation employment (model vs. data)	-	0.98	0.99	1.00	1.00	0.99
correlation consumption (model vs. data)	-	0.99	0.93	0.98	0.98	0.93
correlation wages (model vs. data)	-	0.99	0.94	1.00	0.98	0.94
correlation house prices (model vs. data)	-	0.93	0.91	0.79	0.93	0.91

Figure 2: Timing

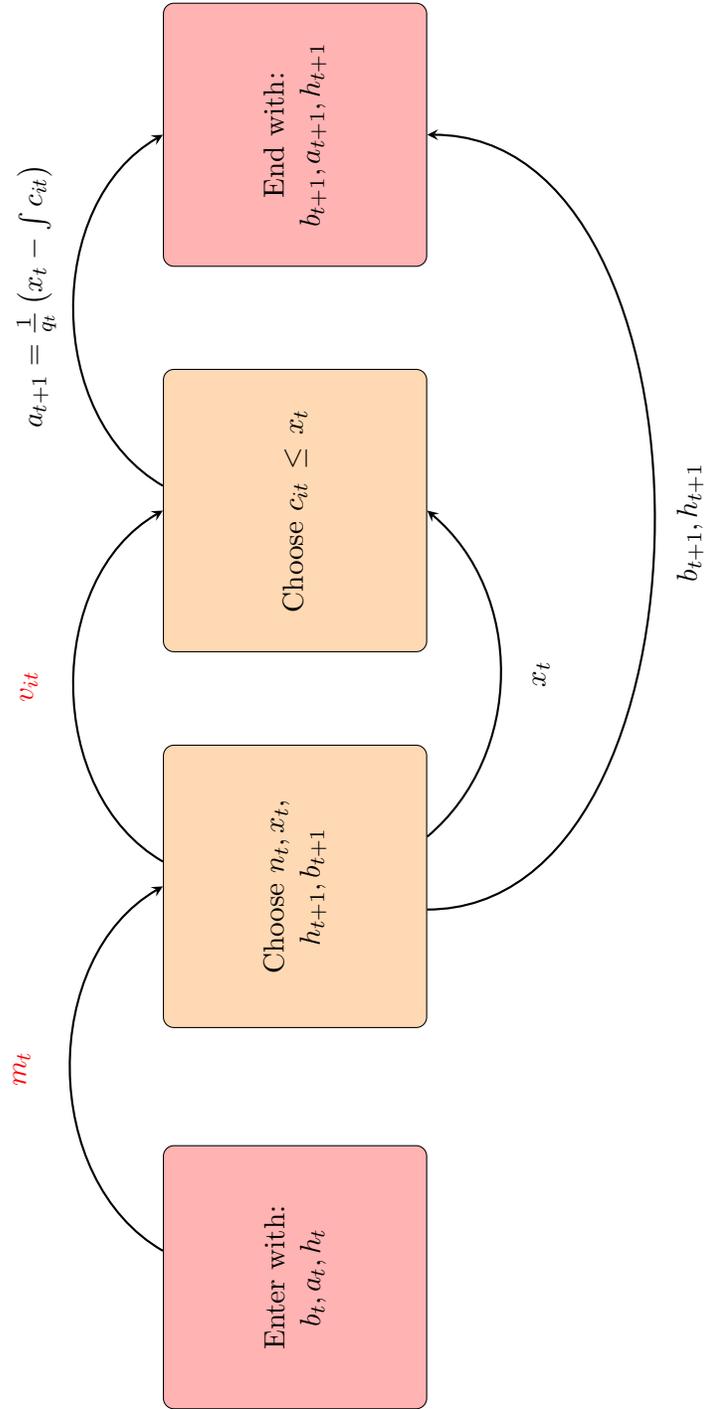


Figure 3a: Steady State Interest Rate, High v Uncertainty

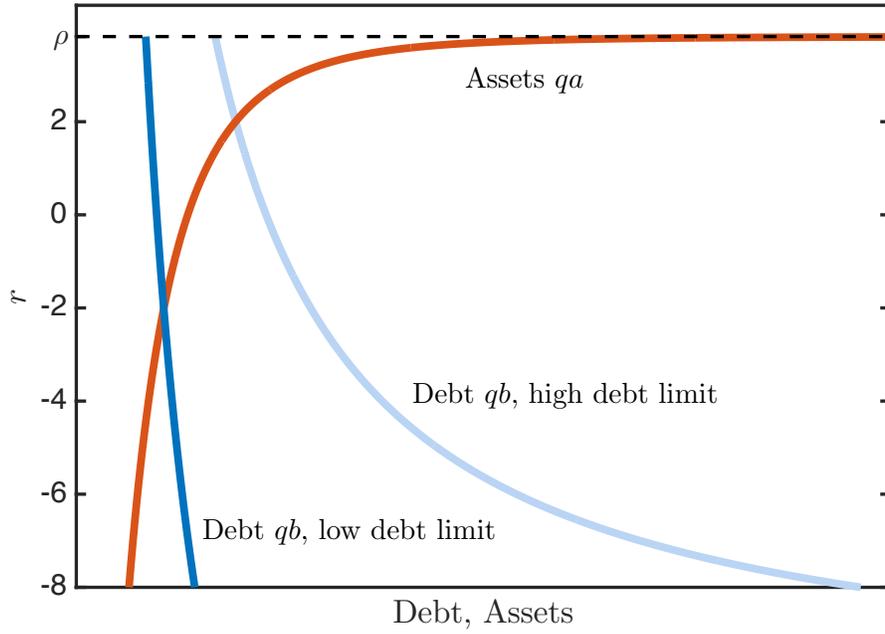


Figure 3b: Steady State Interest Rate, Low v Uncertainty

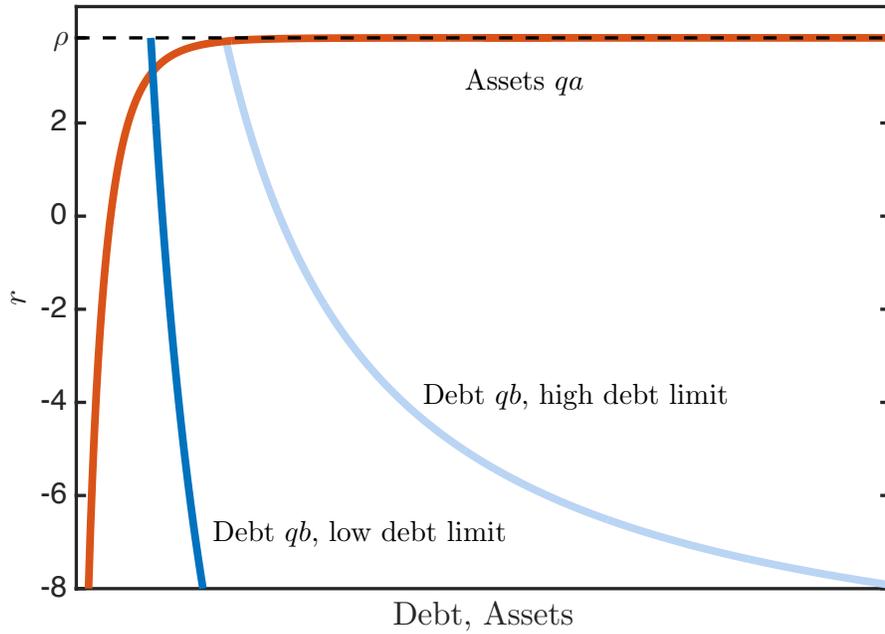


Figure 4: Impulse Response to a Credit Shock, Baseline Closed Economy Model

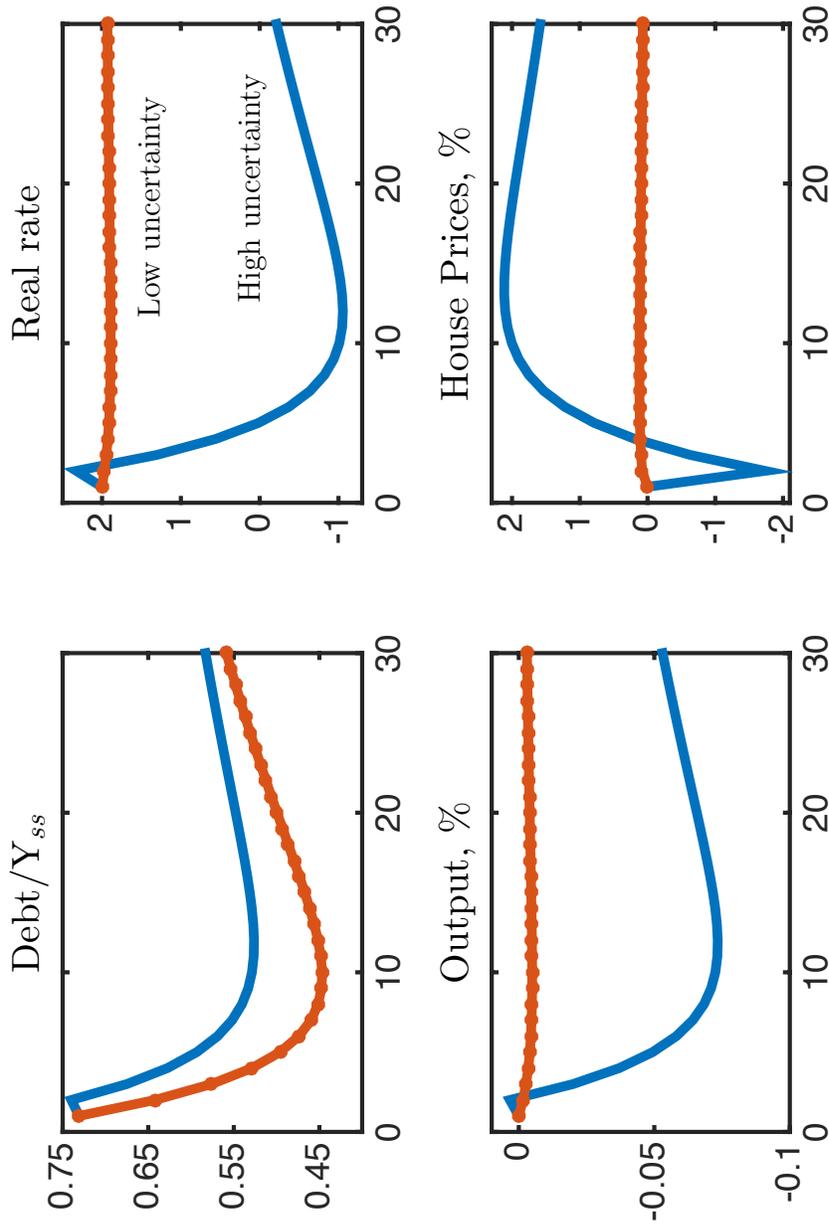


Figure 5a: Island Credit Tightening, Flexible Prices, High v Uncertainty

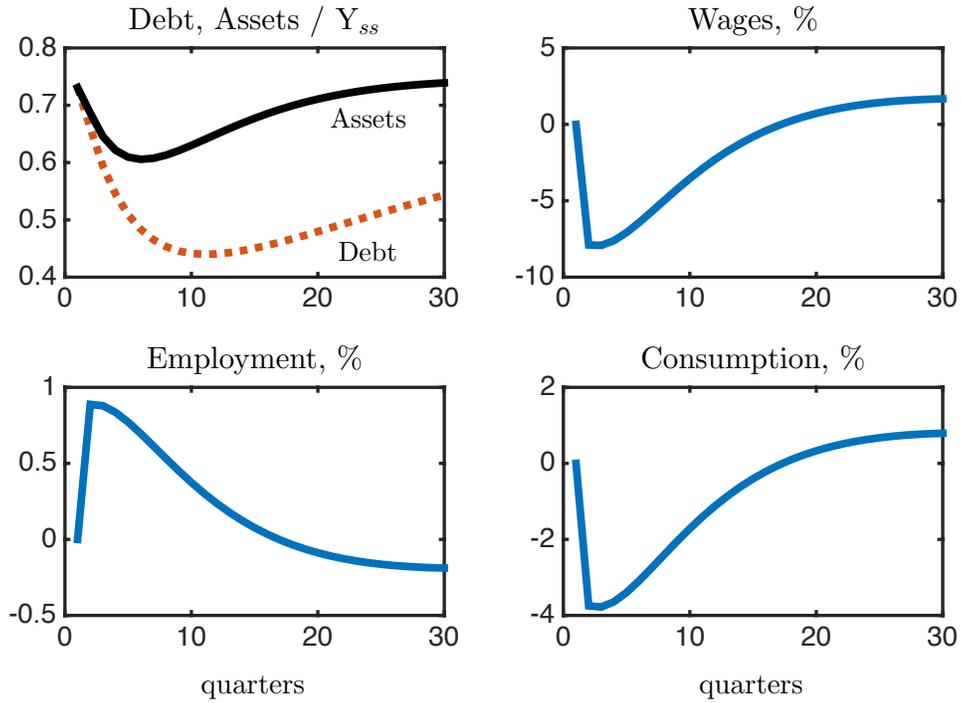


Figure 5b: Island Credit Tightening, Flexible Prices, Low v Uncertainty

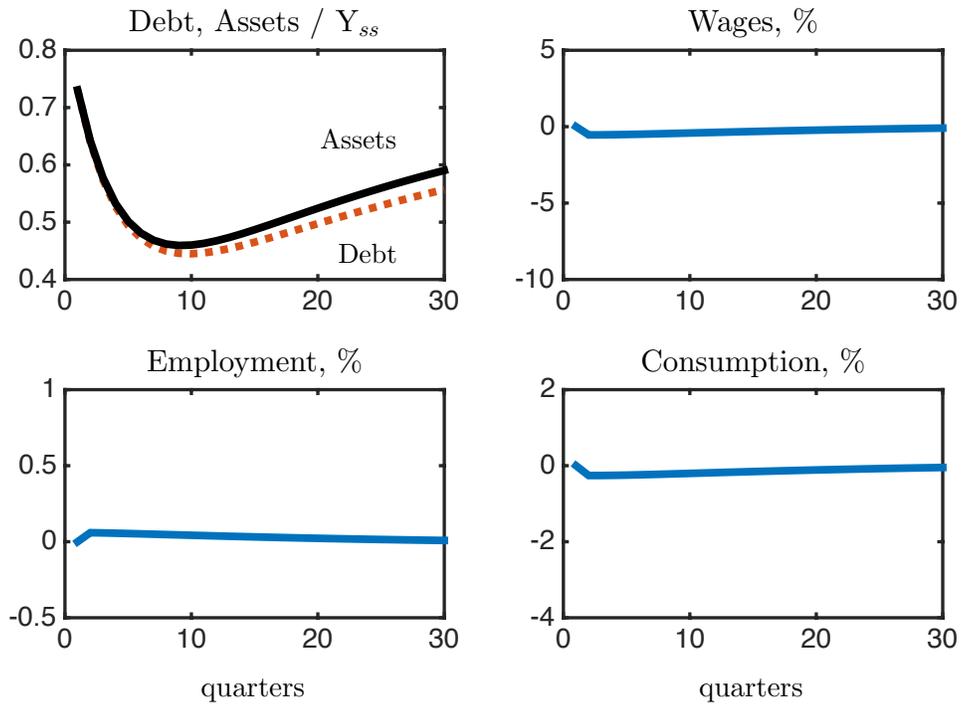


Figure 5c: Island Credit Tightening, Sticky Prices

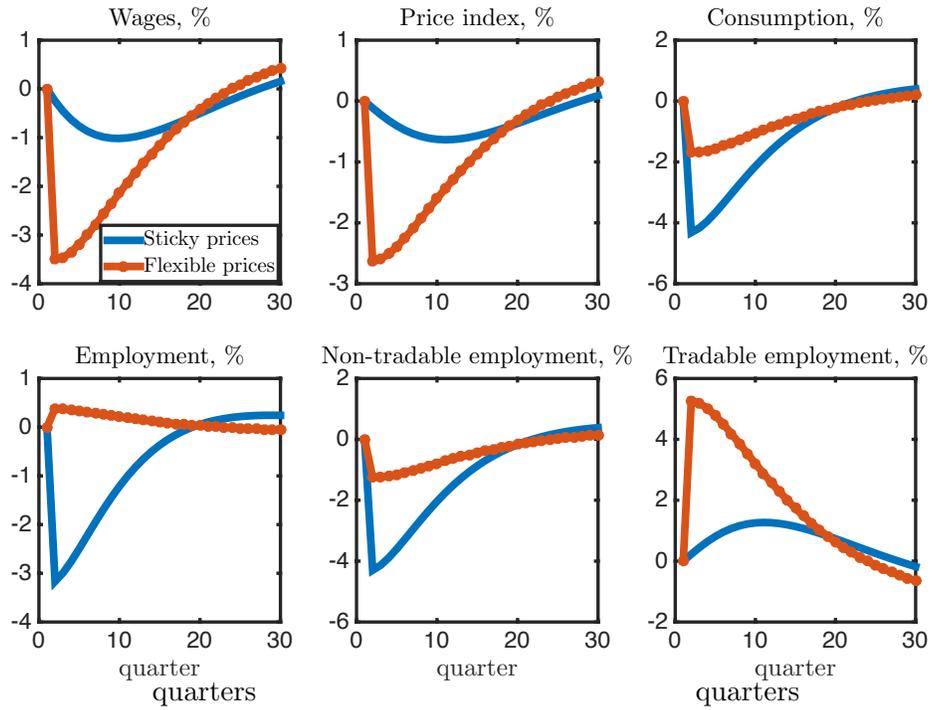


Figure 5d: Aggregate Credit Tightening, Sticky Prices

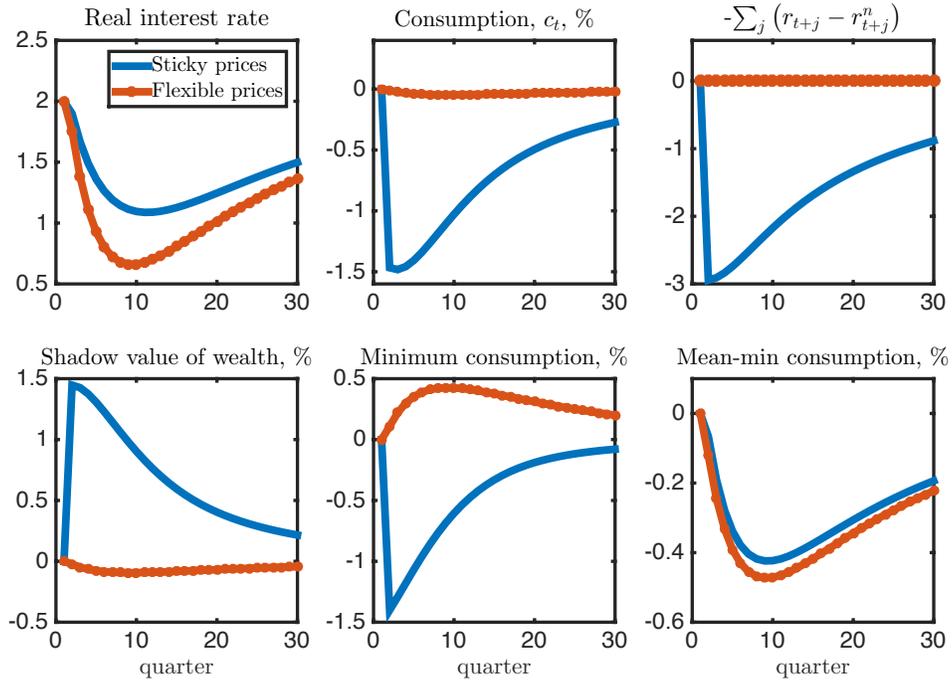


Figure 6a: Employment: Data vs. Fitted Values from Auxiliary Model

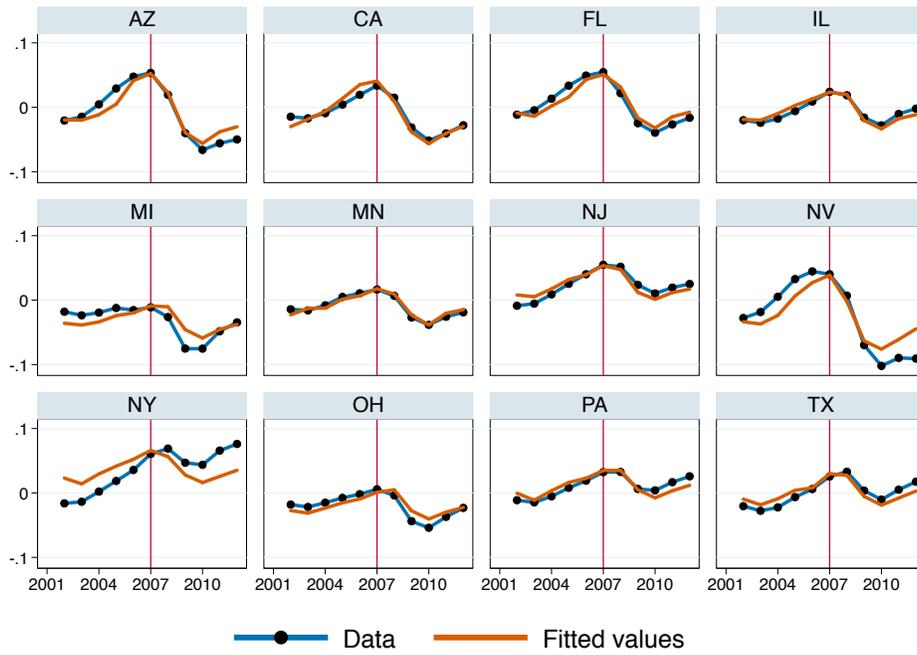


Figure 6b: Employment: Fitted Values Model vs. Data

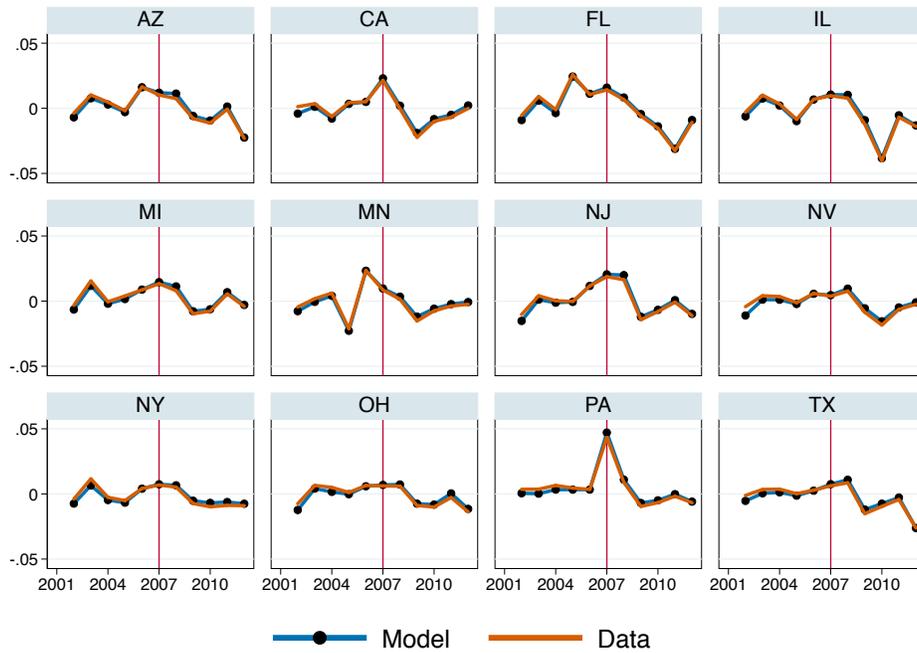


Figure 7a: Impulse Response to an Aggregate Credit Shock

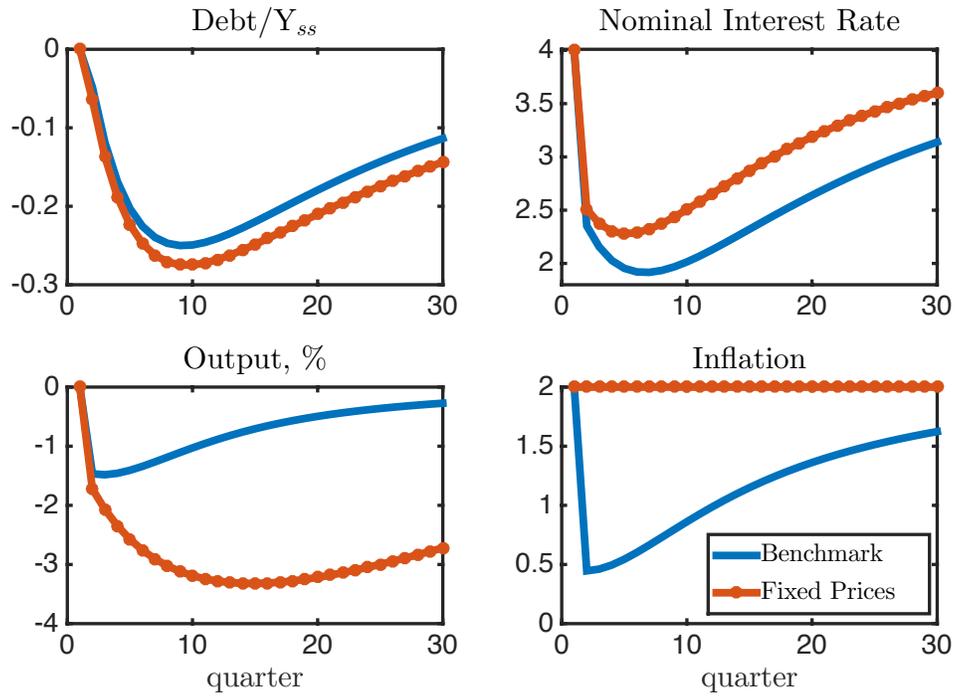


Figure 7b: Impulse Response to an Aggregate Credit Shock

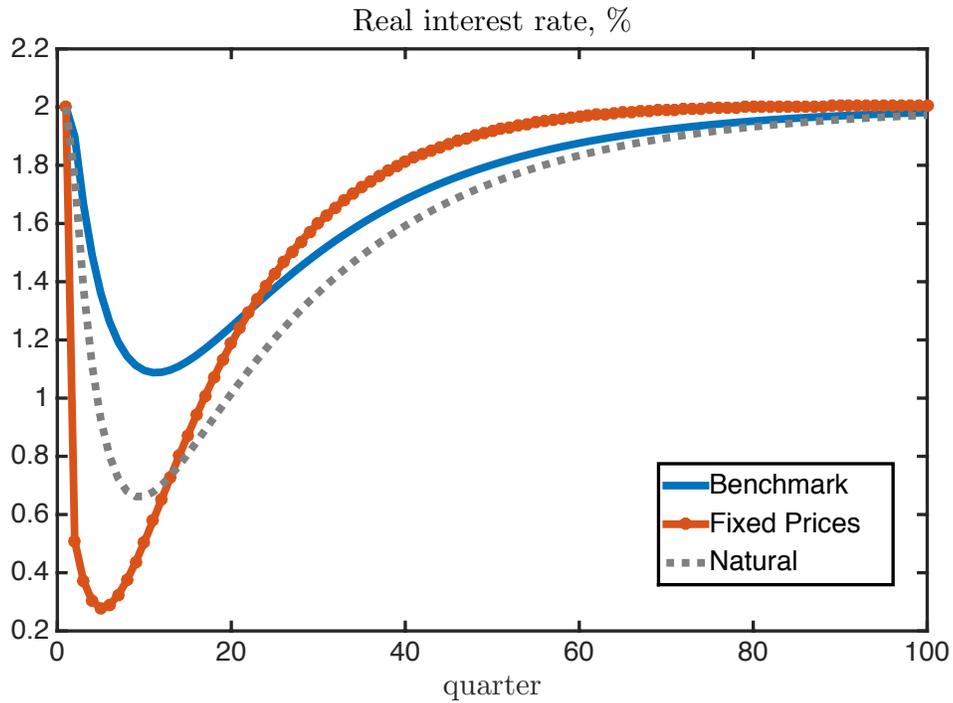


Figure 7c: Comparison of Island and Aggregate Responses

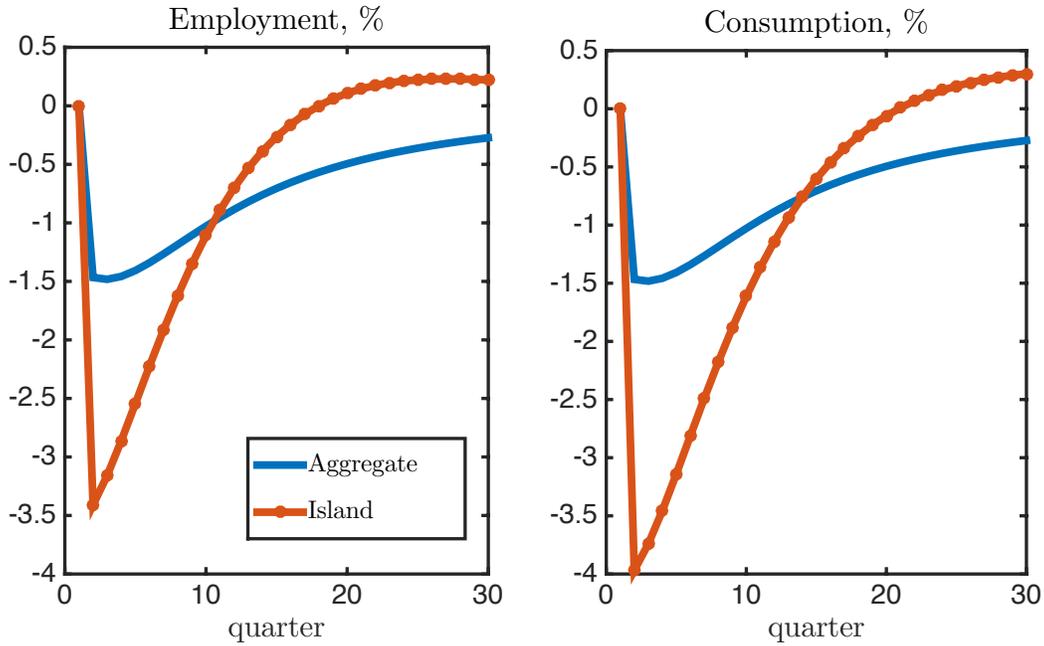


Figure 7d: Impulse Response to a Shock to Spreads

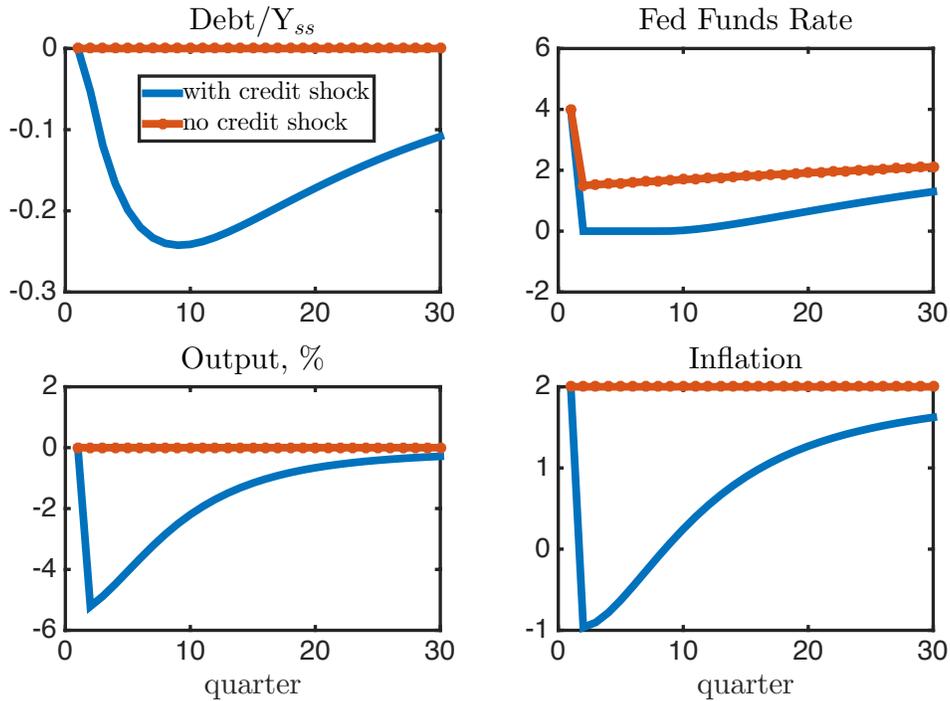


Figure 8: Household Debt in the U.S.

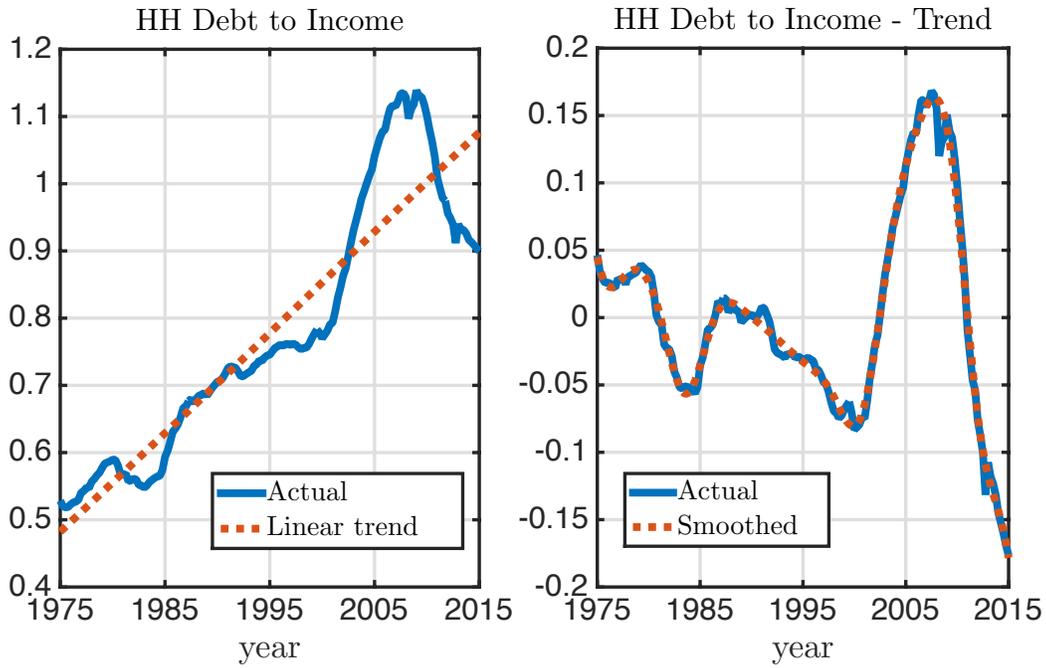


Figure 9: Dynamics of Model that Matches U.S. Household Debt

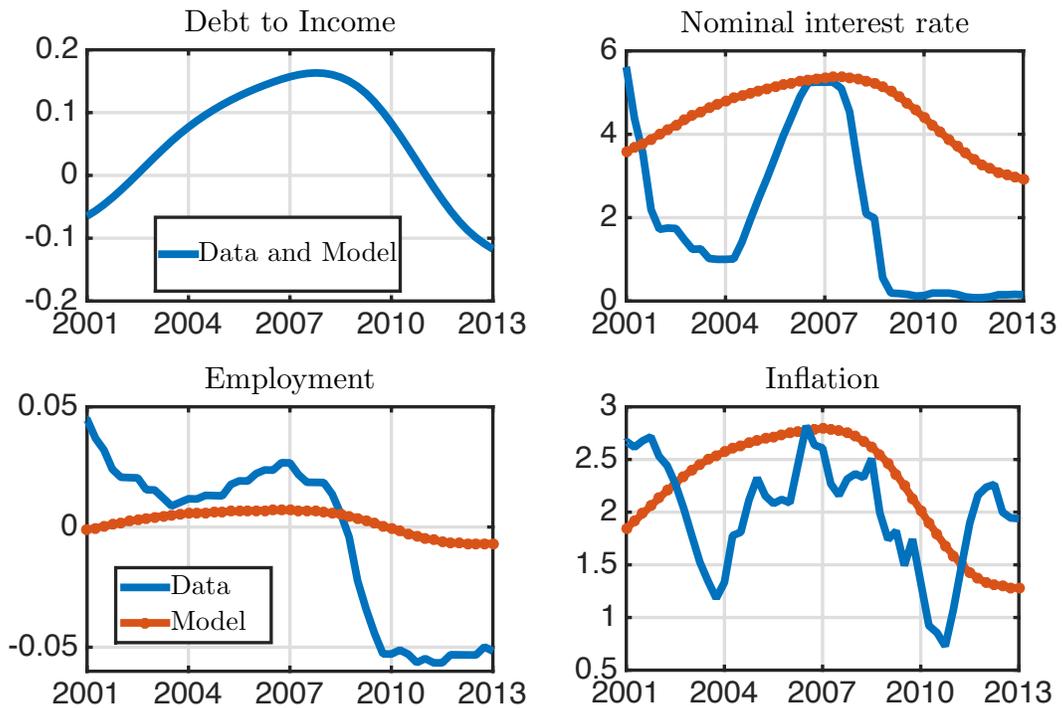


Figure 10: Dynamics of Model in Response to a Spread Shock

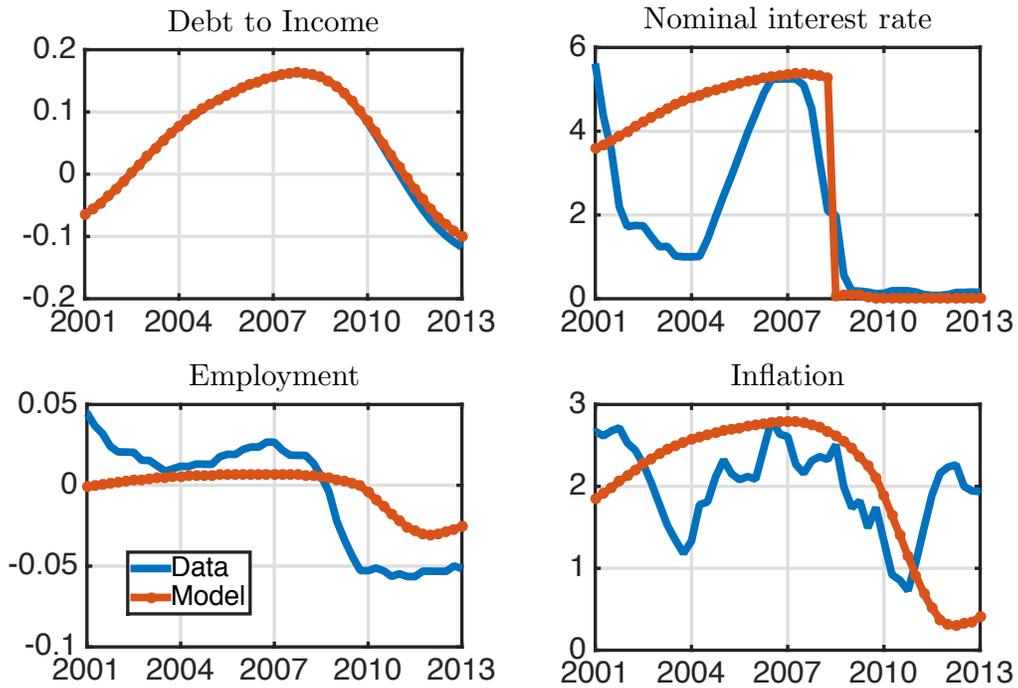


Figure 11: High Uncertainty and Rigidity Experiment

