

Household Leverage and the Recession*

Virgiliu Midrigan[†]

Thomas Philippon[‡]

May 2016

Abstract

A salient feature of the Great Recession is that regions that experienced larger declines in household debt also experienced larger declines in employment. We study a model in which liquidity constraints amplify the response of employment to changes in debt. We estimate the model using panel data on consumption, employment, wages and debt for U.S. states. Though successful in matching the cross-sectional evidence, the model predicts that deleveraging cannot, by itself, account for the large drop in aggregate employment in the U.S. The 25% decline in household debt observed in the data leads to a modest 1.5% drop in the natural rate of interest, and is easily offset by monetary policy. Household deleveraging is more potent, however, in the presence of other shocks that trigger the zero lower bound on interest rates. In the presence of such shocks household deleveraging accounts for about half of the decline in U.S. employment.

Keywords: Great Recession, Household Debt, Regional Evidence, Zero Lower Bound.

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

*We thank Sonia Gilbukh and Callum Jones for excellent research assistance. We thank Robert Shimer and several anonymous referees for valuable advice on a much older version of this paper. We are also indebted to Fernando Alvarez, Andy Atkeson, David Backus, Patrick Kehoe, Andrea Ferrero, Mark Gertler, Veronica Guerrieri, Erik Hurst, Ricardo Lagos, Guido Lorenzoni, Robert Lucas, Atif Mian, Tom Sargent, Nancy Stokey, Amir Sufi, Ivan Werning, Mike Woodford and numerous seminar participants for comments.

[†]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. regions that experienced the largest declines in household (HH) debt 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 against changes in household debt from 2007 to 2010.

One interpretation of this evidence 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 to reduce consumption and led, due to price rigidities, to a reduction in employment.¹

While existing work has argued that HH debt played a major role in accounting for consumption and employment dynamics *across* regions, two important questions remain unanswered. First, can a quantitative model replicate the *regional* evidence? Second, what are the model's *aggregate* implications? Our goal in this paper is to answer these questions. We present a model that reproduces well the panel data on household spending, wages, debt, employment and house prices across U.S. states from 2001 to 2012. We then use the model to ask: what are the aggregate implications of exogenous fluctuations of HH 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 HH debt, the magnitude observed during the Great Recession?

The model we study is one of liquidity-constrained households. Our choice is motivated by the work of Kaplan and Violante (2014) who show that a large fraction of U.S. households is liquidity constrained.² Unlike Kaplan and Violante (2014), we do not assume transactions costs of converting housing wealth into a liquid asset.³ Considerations of computational tractability led us to follow the approach of Lucas (1990) in assuming that agents must allocate their wealth between housing and a liquid asset *before* an *idiosyncratic* shock to preferences is realized. We assume that such shocks are i.i.d. and use a family construct to eliminate the distributional consequences of asset market incompleteness.⁴ These assumptions simplify our analysis and allow us to study an economy with aggregate shocks and a zero lower bound (ZLB) on nominal interest rates. As we show, the latter is critical in shaping the model's aggregate implications.

Despite risk sharing at the family level, the quantity of HH debt in our economy has im-

¹Mian and Sufi (2011, 2014) and Mian, Rao and Sufi (2013) provide empirical evidence in support of this view. Our paper focuses on state-level data, but the cross-sectional patterns hold at a finer level as well.

²See also Lucas and Stokey (2011) who emphasize the role of liquidity frictions in the recent financial crisis.

³See also recent work by Kaplan, Mittman, Violante (2015) and Gorea and Midrigan (2015).

⁴Challe et. al. (2015) also employ a family construct in order to characterize the aggregate properties of an economy with uninsurable unemployment risk.

portant aggregate consequences. The presence of idiosyncratic uncertainty leads agents in our environment to save for precautionary reasons. In a flexible-price variant of the model, the equilibrium interest rate is below the rate of time preference and is pinned down by the amount agents can borrow. A tightening of credit leads 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 and wages are sticky, the response of real variables to credit shocks depends on the extent to which the nominal interest rate tracks the dynamics of the natural rate. Absent the ZLB, 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. At the ZLB, however, the economy's dynamics are highly non-linear: the marginal effect of a shock to HH debt depends greatly on all other aggregate shocks. This non-linearity gives rise to an identification problem, which we solve by following the approach of Jones (2015) who uses data on forecasts of how long the ZLB will last. We follow his approach and show that, conditional on a given expected duration of the ZLB, the marginal effect of a HH debt shock is uniquely pinned down. We therefore use both state-level data on the comovement of HH debt, wages, consumption and employment, as well as information on the expected path of future interest rates, to study the macroeconomic implications of household deleveraging.

We pin down the model's key parameters using an indirect inference approach. 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 fluctuations in household debt. We then choose the key parameters, including the persistence of credit shocks, the maturity of securities agents trade, the degree of wage rigidity and degree of openness of individual states, by requiring that the coefficients in the auxiliary regressions from the model match those in the data. We show, by bootstrapping our estimates, that the model's parameters are well-identified by the state-level data, with small standard errors around the estimates.

A key parameter in our model is the degree of idiosyncratic uncertainty faced by households. This parameter is pinned down by the comovement between consumption and debt in the state-level data. To understand why this is the case, consider an individual state in which agents experience a sudden tightening of their credit limits. If the amount of idiosyncratic uncertainty agents face is high, they anticipate a high likelihood of being liquidity constrained and find it optimal to continue saving, and therefore reduce consumption sharply. If, in contrast, the amount of idiosyncratic uncertainty is low, the precautionary savings motive is weak and

households can simply reduce the asset side of their balance sheet to respond to the tightening of credit, resulting in a mild drop in consumption.

The model thus captures, in a parsimonious way, the notion that the sensitivity of the economy to changes in HH credit depends on the amount of idiosyncratic uncertainty. By choosing this parameter appropriately, the model can replicate the state-level comovement between consumption and debt in the data. In contrast, models that assume permanent differences in the households' discount factors cannot match the data. In those models impatient agents have no assets and are thus forced to cut their consumption by the full amount of the drop in credit, implying counterfactual 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 aggregate credit shocks. When idiosyncratic uncertainty is high, agents' savings are not very sensitive to changes in interest rates and large reductions in the real interest rate are necessary to ensure that the asset market clears following a tightening of credit. If this is the case, monetary policy would not be able to offset credit shocks because of the ZLB. If, in contrast, the degree of demand uncertainty is low, agents' savings are sensitive to changes in interest rates. In this case the natural rate falls little and monetary policy can offset the credit shock.

Our parameter estimates imply that changes in household debt of the magnitude observed in the Great Recession generate a decline in the natural rate of interest of about 1.5%. Since the Fed Funds rate was equal to 5% at the onset of the recession, the drop in the natural rate caused by the HH credit tightening could have been easily offset without triggering the ZLB. Because of this, household credit shocks *alone* generate a modest, 1.4% drop in employment.

Our estimates also imply, however, that the drop in the natural rate of interest was persistent. Thus, if HH credit shocks are accompanied by other shocks that trigger the ZLB, the marginal impact of HH credit shocks is much larger. We capture this possibility by considering a shock to the spread between the Fed Funds rate and the interest rate, in the spirit of Gilchrist and Zakrajšek (2012). We show that in the presence of this *spread shock*, the economy's responses to changes in HH debt are highly non-linear. We thus use private sector forecast on the expected number of quarters until a Fed Funds rate increase to discipline the model. We find, when we do so, that shocks to HH credit account for about half of the drop in employment in the U.S. data. In particular, while the model fails to account for the rapid decline in employment in 2008, the fact that HH debt shocks are persistent allows the model to account for its slow recovery.

We conclude, therefore, that HH credit is an important source of employment fluctuations

during the Great Recession, but not the only one. Shocks to HH credit can have sizable effects on real activity in the presence of additional shocks in the economy, but not on their own.

Related Work In addition to the work of Mian and Sufi, our paper is related to Guerrieri and Lorenzoni (2015) and Eggertsson and Krugman (2012) who also study the responses of an economy to a HH credit crunch. These researchers find, as we do, that a credit crunch has a minor effect on employment away from the ZLB. While they study a closed-economy setting, our model is that of a monetary union composed of a large number of states. Moreover, our focus is on estimating the model using state-level 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 that wages in individual states comove quite strongly with employment. Matching this evidence implies a steep slope of the Phillips curve in the aggregate, further reinforcing our message that household credit shocks alone cannot account for the bulk of the recession. 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 regional evidence in identification.

Our paper is also related to the literature on housing wealth and consumption.⁵ An important reference is Iacoviello (2005), who 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 reduce the interest rate below the rate of time preference. Much of the recent literature has focused on households' balance sheets but Giroud and Mueller (2015) show that firms balance sheets played an important role in the transmission of consumer demand shocks during the Great Recession.

Finally, our work is related to the literature on financial intermediation, originating with Bernanke and Gertler (1989), Kiyotaki and Moore (1997), Bernanke, Gertler and Gilchrist (1999) and more recently Mendoza (2010), Gertler and Karadi (2011), Gertler and Kiyotaki (2010), and Gilchrist and Zakrajšek (2012). This literature focuses on understanding the role of shocks that disrupt financial intermediation, which we argue must accompany household credit shocks for the model to be able to replicate the large decline in U.S. employment.

⁵See also the work Lustig and van Nieuwerburgh (2005), Garriga, Manuelli and Peralta-Alva (2014), Favilukis, Ludvigson and van Nieuwerburgh (2015), Burnside, Eichenbaum and Rebelo (2015), Landvoigt, Piazzesi and Schneider (2015) who study the determinants of house prices.

2 A Baseline Closed-Economy Real Model

We first describe our model of liquidity constraints in a closed economy without price rigidities. We explain how the precautionary savings motive and household's ability to borrow against the value of their homes interact to determine the equilibrium interest rate.

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.

Technology and Preferences Competitive firms produce output y_t with labor n_t subject to

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

Competition pins down the real wage so $w_t = 1$. 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.

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, 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 drawn from a Pareto distribution

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

Here $\alpha > 1$ determines the amount of uncertainty about v . A lower α implies more uncertainty.

Securities The only security traded in this economy is a long-term perpetuity with coupon payments that decay geometrically at a rate determined by a parameter γ . A seller of such a security issues one unit at a price q_t in period t and repays 1 unit of the 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.⁶ The household trades this security with perfectly competitive financial intermediaries.

⁶See Hatchondo and Martinez (2009) and Arellano and Ramanarayanan (2012) who describe the properties of these securities in more detail. Assuming that savings have a shorter duration than debt, or allowing for a spread between the interest rate on debt and savings would be straightforward in this setup but would clutter the algebra without offering much additional insights.

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

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

Similarly, we let a_t denote the amount of coupon payments the household is entitled to receive in period t . Thus a_t represents the household's assets, while b_t denotes its debt.

Budget and Borrowing Constraints 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:

$$c_{it} \leq x_t. \quad (5)$$

We refer to the constraint in (5) as the *liquidity constraint*.

The budget constraint states that

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

In words, 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 also assume a *borrowing constraint* that limits the household's ability to issue *new* loans. The face value of new loans issued is limited by a multiple m_t of the value of one's home:

$$q_t l_t \leq m_t e_t h_{t+1}. \quad (7)$$

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:

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

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 of the borrowing limit restricts a household's ability to take on new loans, not its total debt $q_t b_{t+1}$. We make this assumption 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, which would be counterfactual.

Savings Individual households both borrow and save using the long-term security. A household's savings are 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 then

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

Timing We summarize, 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 x_t . Each individual members' preference for consumption v_{it} is realized and individual members purchase c_{it} units of consumption. At the end of the period all unspent funds are pooled to purchase a_{t+1} units of the 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 to maximize its 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 assumption that transfers x_t are chosen prior to the realization of the idiosyncratic preference shock v_{it} by imposing 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). Let R_{t+1} denote the 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 expectation operator. Since the loan-to-value limit m_t is the only source of aggregate uncertainty, \mathbb{E}_t is taken over the realization of the credit shock ε_t .

This expression is quite intuitive. The 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 $\beta\mathbb{E}_t\mu_{t+1}R_{t+1}$. In addition, transfers provide a liquidity service by relaxing the liquidity constraint of individual members. Since transfers are chosen prior to the realization of the taste shock, these liquidity services are equal to 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 v .

Consider next the household's choice of debt. 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. 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 $\beta\mathbb{E}_t\mu_{t+1}R_{t+1}$. Borrowing an extra unit tightens today's borrowing constraint (λ_t), but relaxes next period's constraint (λ_{t+1}) because of 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: the difference between the purchase price and next period's selling price, appropriately discounted. The right hand side is the marginal utility of housing services $\frac{\bar{\eta}}{h_{t+1}}$ as well as the collateral value of housing $\lambda_t m_t e_t$.

Consider finally the choice of consumption of individual members. With logarithmic preferences⁷ the choice of consumption reduces to

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

2.3 Equilibrium

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

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

⁷Our robustness section below studies an economy with more general CRRA preferences.

The supply of labor is given by

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

Recall finally that firm optimization implies $w_t = 1$ and that the housing stock is in fixed supply.

2.4 The Workings of the Model

We next briefly discuss the workings of the model. Let

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

denote the consumption of a member with the lowest taste shock, $v_{it} = 1$. Using the first-order conditions above, one can show that the transfer to individual members satisfies

$$\frac{1}{\alpha - 1} \left(\frac{x_t}{\underline{c}_t} \right)^{-\alpha} = \left(\beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} R_{t+1} \right)^{-1} - 1 \approx \rho_t - r_t, \quad (18)$$

where $\rho_t = -\log \beta \mathbb{E}_t \left(\frac{\mu_{t+1}}{\mu_t} \right)$ is the discount rate and $r_t = \log \mathbb{E}_t(R_{t+1})$ is the interest rate.

Intuitively, the right-hand side of (18) is equal (up to a first-order approximation) to the difference between the discount rate and the interest rate, while the left-hand side is proportional to the fraction of constrained household members, that is, those with $v_{it} > x_t/\underline{c}_t$. 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.

Consider next the household's total consumption expenditures, $c_t = \int_0^1 c_{it} di$. We have

$$\frac{c_t}{\underline{c}_t} = \frac{\alpha}{\alpha - 1} \left(1 - \frac{1}{\alpha} \left(\frac{x_t}{\underline{c}_t} \right)^{1-\alpha} \right). \quad (19)$$

The lower the gap between the discount rate and interest rate is, the lower the fraction of constrained members, and therefore the larger the mean/min consumption ratio.

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:

$$\frac{q_t a_{t+1}}{c_t} = \left(\frac{\alpha}{\alpha - 1} [(\alpha - 1)\Delta_t]^{\frac{1}{\alpha}} - \Delta_t \right)^{-1} - 1, \quad (20)$$

which increases as Δ_t decreases and is steeper the higher Δ_t is.

Consider now the household's decision of how much to borrow. Because the taste shocks are unbounded, the expected multiplier on the liquidity constraint, the LHS of (18), is always positive. Thus, a comparison of (11) and (12) reveals that the multiplier on the borrowing constraint, λ_t , is always positive as well. Intuitively, agents would like to borrow as much as possible in this economy as long as the interest rate is below the rate of time preference, which is always the case unless we eliminate the uncertainty about preference shocks.

2.5 Steady State Equilibrium Interest Rate

Consider next how the equilibrium interest rate is determined in the steady state with a constant credit limit $m_t = \bar{m}$. We have already discussed the supply of assets in the previous section. Consider next the demand for assets. Because the borrowing limit binds, $qb = \frac{1}{1-\gamma}\bar{m}eh$, the amount of debt in the economy is proportional to the value of houses. The value of houses reflects both their service flow, as well as their collateral value. The latter declines as the interest rate increases since a higher interest rate make borrowing less attractive. To see this, notice that in the steady state the Euler equation for housing (13) reduces to

$$eh = \frac{\bar{\eta}}{\mu} \frac{1}{\rho - \frac{\bar{m}}{1-\beta\gamma}(\rho - r)}. \quad (21)$$

Clearly, a higher interest rate reduces housing values and the amount the household can borrow.

Figures 3a and 3b illustrate how the interest rate is determined in the steady state of the model. Figure 3a assumes a relatively large degree of idiosyncratic uncertainty about the taste shocks. Notice how the intersection of the upward-sloping savings 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. Figure 3b assumes a relatively low degree of idiosyncratic uncertainty. In this case agents save less and the equilibrium interest rate is higher. Moreover, the intersection of the asset and debt curves now occurs 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 , which, due to the long-term nature of securities, leads to a gradual reduction in household debt. We contrast the responses in economies with a high and low degree of idiosyncratic uncertainty.

Notice that the equilibrium interest rate falls in both economies. The interest rate falls more in the economy with greater demand uncertainty, reflecting the steeper savings curve. In contrast, output barely falls. Although a tightening of credit magnifies the consumption-leisure distortions, these are small here, as in cash-in-advance models.⁸

⁸Cooley and Hansen (1991).

3 An Island Monetary Economy with Price Rigidities

We next embed the frictions described above into a monetary economy with price and wage rigidities. The economy consists 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. We assume that intermediate goods producers are monopolistically competitive. Individual households on each island belong to unions that sell differentiated varieties of labor. Prices and wages are subject to Calvo adjustment frictions. Labor is immobile across islands and the housing stock on each island is in fixed supply.

3.1 Setup

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. 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 last reset its wage. The problem of a household that belongs to labor union z is to

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int_0^1 v_{it}(s) \log(c_{it}(s)) di + \bar{\eta} \log(h_t(s)) - \frac{1}{1+\nu} n_t(z, s)^{1+\nu} \right] \quad (22)$$

subject to the budget constraint

$$p_t(s)x_t(s) + e_t(s)(h_{t+1}(s) - h_t(s)) = w_t(z, s)n_t(z, s) + q_t l_t(s) - b_t(s) + (1 + \gamma q_t)a_t(s) + T_t(z, s), \quad (23)$$

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 firms on that particular island.

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:

$$c_{it}(s) \leq x_t(s), \quad (24)$$

a borrowing constraint

$$q_t l_t(s) \leq m_t(s) e_t(s) h_{t+1}(s), \quad (25)$$

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

$$q_t a_{t+1}(s) = p_t(s) \left(x_t(s) - \int_0^1 c_{it}(s) di \right). \quad (26)$$

There are no barriers to capital flows, so all islands trade securities at a common price q_t .

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 non-tradable goods produced locally and $y_t^T(s, j)$ units of tradable goods produced on island j :

$$y_t(s) = \left(\omega^{\frac{1}{\sigma}} y_t^N(s)^{\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} \left(\int_0^1 y_t^T(s, j)^{\frac{\kappa-1}{\kappa}} dj \right)^{\frac{\kappa-1}{\kappa-1} \frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (27)$$

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. Letting $p_t^N(s)$ and $p_t^T(s)$ denote the prices of these goods on island s , the final goods price on an island is

$$p_t(s) = \left(\omega p_t^N(s)^{1-\sigma} + (1-\omega) \left(\int_0^1 p_t^T(j)^{1-\kappa} dj \right)^{\frac{1-\sigma}{1-\kappa}} \right)^{\frac{1}{1-\sigma}}. \quad (28)$$

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

$$y_t^N(s) = \omega \left(\frac{p_t^N(s)}{p_t(s)} \right)^{-\sigma} y_t(s), \quad (29)$$

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

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

Intermediate Goods Producers Traded and non-traded goods on each island are themselves CES composites of varieties of differentiated intermediate inputs with an elasticity of substitution ϑ . The demand for an individual variety is $y_t^T(z, s) = (p_t^T(z, s)/p_t^T(s))^{-\vartheta} y_t^T(s)$.

Individual producers of intermediate goods are subject to Calvo price adjustment frictions. Let λ_p denote the probability that a firm does not reset its price in a given period. A firm that resets its price maximizes the present discounted flow of profits weighted by the probability that the price it chooses at t will still be in effect at any particular date. 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

$$\max_{p_t^{T^*}(s)} \sum_{k=0}^{\infty} (\lambda_p \beta)^k \mu_{t+k}(s) (p_t^{T^*}(s) - \tau_p w_t(s)) \left(\frac{p_t^{T^*}(s)}{p_t^T(s)} \right)^{-\vartheta} y_t^T(s), \quad (31)$$

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 (23), 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 price of traded goods is

$$p_t^T(s) = \left((1 - \lambda_p) p_t^{T^*}(s)^{1-\vartheta} + \lambda_p p_{t-1}^T(s)^{1-\vartheta} \right)^{\frac{1}{1-\vartheta}} \quad (32)$$

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

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

where ψ is the elasticity of substitution. Demand for an individual union's labor 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 to choose a new wage $w_t^*(s)$ to

$$\max_{w_t^*(s)} \sum_{k=0}^{\infty} (\lambda_w \beta)^k \left[\tau_w \mu_{t+s} w_t^*(s) \left(\frac{w_t^*(s)}{w_t(s)} \right)^{-\psi} n_t(s) - \frac{1}{1+\nu} \left(\left(\frac{w_t^*(s)}{w_t(s)} \right)^{-\psi} n_t(s) \right)^{1+\nu} \right], \quad (34)$$

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

$$w_t(s) = \left((1 - \lambda_w) w_t^*(s)^{1-\psi} + \lambda_w w_{t-1}(s)^{1-\psi} \right)^{\frac{1}{1-\psi}}. \quad (35)$$

The elasticity of substitution ψ determines the extent to which wages respond to credit shocks. To see this, log-linearize the optimal choice of reset wages that solves (34) around the steady-state:

$$\hat{w}_t^*(s) = \beta \lambda_w \mathbb{E}_t \hat{w}_{t+1}^*(s) + \frac{1 - \beta \lambda_w}{1 + \psi \nu} (-\hat{\mu}_t(s) + \psi \nu \hat{w}_t(s) + \nu \hat{n}_t(s)), \quad (36)$$

where hats denote log-deviations from the steady state. 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 supplying more hours, they are less inclined to do so when the elasticity of substitution ψ is high. Intuitively, if ψ is high, cutting wages would lead to a large increase in the amount of labor supplied by a union and its members' disutility from work.

Island Equilibrium The composite labor services are used by producers of both tradable and non-tradable goods:

$$n_t(s) = \int_0^1 \left(\frac{p_t^N(z, s)}{p_t^N(s)} \right)^{-\vartheta} y_t^N(s) dz + \int_0^1 \left(\frac{p_t^T(z, s)}{p_t^T(s)} \right)^{-\vartheta} y_t^T(s) dz. \quad (37)$$

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

$$\underline{c}_t(s) = \frac{1}{\beta \mathbb{E}_t \mu_{t+1}(s) R_{t+1}} \frac{1}{p_t(s)}, \quad (38)$$

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 (19) as earlier.

The island's net asset position evolves according to:

$$q_t(a_{t+1}(s) - b_{t+1}(s)) = (1 + \gamma q_t)(a_t(s) - b_t(s)) + w_t(s)n_t(s) + T_t(s) - p_t(s)c_t(s). \quad (39)$$

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

$$1 + i_t = \mathbb{E}_t R_{t+1} \quad (40)$$

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:

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

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 $\bar{\pi}$.⁹ Since an individual island is of measure zero, monetary

⁹We 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 all prices and wages are automatically indexed to $\bar{\pi}$. See Coibion and Gorodnichenko (2014) and Blanco (2015) who study the size of these costs in the absence of indexation.

policy does not react to island-specific disturbances. The monetary union is closed so aggregate savings must equal to aggregate debt:

$$\int_0^1 a_{t+1}(s)ds = \int_0^1 b_{t+1}(s)ds. \quad (41)$$

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 for this class of models, credit shocks alone cannot generate movements in house prices as large as those in the data. We thus assume shocks to both the loan-to-value ratio as well as the consumer's preference for housing.¹⁰ In particular, we modify the utility function in (22) to introduce time-varying weights on housing in preferences, $\eta_t(s)$.

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

$$\log m_t(s) = (1 - \rho) \log \bar{m} + \rho \log m_{t-1}(s) + \varepsilon_t(s), \quad (42)$$

as does the preference weight on housing:

$$\log \eta_t(s) = (1 - \rho) \log \bar{\eta} + \rho \log \eta_{t-1}(s) + \sigma_\eta \varepsilon_t(s). \quad (43)$$

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_η 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. This follows because housing is separable in the utility function and the housing stock is fixed.

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 tightening of household credit. We start by discussing the responses in an economy with flexible prices and then those in an economy with price stickiness.

Figure 5a shows the responses in a flexible price economy in which the degree of demand uncertainty is relatively high. In contrast to the closed economy, the asset holdings of agents on

¹⁰See Kiyotaki, Michaelides and Nikolov (2011) for an illustration of the problem and Garriga, Manuelli and Peralta-Alva (2014) and Favalukis, Ludvigson and Van Nieuwerburgh (2015) for approaches to resolve it.

the island do not fall nearly as much as their debt does. The island's net foreign asset position thus increases. Financing this increase requires that agents on the island reduce consumption and increase employment. Wages fall as well to ensure that the rest of the economy buys more of the island's traded goods.

Figure 5b shows that the responses of all real 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 savings or by cutting consumption and leisure. When demand uncertainty is low, it is relatively costless to reduce savings and so an island's assets fall nearly as much as its debt. Both sides of the island's balance sheet thus contract, with little impact on other variables. In contrast, when demand uncertainty is high, reducing savings is costly since individual members are more likely to be liquidity constrained. The household thus finds it optimal to respond to the credit tightening by cutting consumption by a lot more.

To summarize, in our model credit constraints prevent households from smoothing the marginal utility of consumption both across members as well as across time. Households can respond to a tightening of credit constraints by either reducing overall consumption, thus worsening the intertemporal allocation, or by reducing savings, thus worsening the allocation of consumption across agents. The more dispersed the idiosyncratic shocks are, the more the household chooses to reduce overall consumption to avoid the high costs of variation of the marginal utility of consumption across its members.

Figure 5c illustrates the role of price rigidities. The upper row of the figure shows that wages and prices react much more gradually to the credit shock when prices are sticky, while consumption falls 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 occur with flexible prices is now muted. Intuitively, 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.¹¹ 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 associated with a decline in non-tradable employment which is no longer offset by an increase in tradable employment. Consequently, employment on the island falls, more so when non-traded goods account for a larger fraction of spending.

¹¹See 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.

3.5 Impulse Response to an Aggregate Credit Tightening

The Appendix provides a more detailed discussion about the mechanics of the responses to an aggregate credit crunch. Briefly, consumption and thus output and employment fall in our model for similar reasons as in simpler New Keynesian models in which exogenous shocks to the discount factor drive movements in the natural rate.¹² An important difference, however, is that in our model the natural rate is determined endogenously, by the availability of credit as well as the strength of the precautionary-savings motive. To see this, notice that the shadow value of wealth μ_t satisfies an Euler equation similar to that in the simpler New Keynesian model:

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

where Δ_t is implicitly a function of the ratio of household debt to consumption:

$$\left(\frac{\alpha}{\alpha - 1} [(\alpha - 1)\Delta_t]^{\frac{1}{\alpha}} - \Delta_t \right)^{-1} - 1 = \frac{q_t b_{t+1}}{c_t}, \quad (44)$$

where we have used (20) and the asset market clearing condition, $a_{t+1} = b_{t+1}$. Since consumption is inversely related to the shadow value of wealth, an increase in μ_t reduces overall consumption, and, due to price and wage rigidities, overall output and employment.

As in simpler New Keynesian models, μ_t increases after a credit tightening only if monetary policy does not offset the credit shock. In particular, 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, that is, the *natural interest rate*. To the extent to which monetary policy does not follow such a policy, either due to the fact that it uses a less aggressive Taylor rule or due to a zero lower bound, credit shocks generate fluctuations in real variables.

Figure 5d compares the responses of aggregate variables after an aggregate credit shock to those of island variables to an equally-sized island-level shock. Notice in Figure 5d that the initial drop in both consumption and employment is greater at the island level than it is in the aggregate. This owes to the fact that the Taylor rule partly offsets the shock in the aggregate but, by assumption, does not react to an island-specific shock. Notice also that the recession in the aggregate is much more persistent than in an individual island, even though the credit shock is equally persistent. Intuitively, an individual island's consumption is pinned down by the *change* in household credit, 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.

¹²See, for example, the work of Christiano, Eichenbaum and Rebelo (2011) and Werning (2012).

The broader implication of these impulse responses is that we cannot simply extrapolate the state-level correlations to inform about the role of credit in generating fluctuations in the aggregate. While the state-level evidence is useful in identifying the key parameters of the model, we need to use the structure of the model to study the role of aggregate credit shocks.

4 Quantification

We next describe how we have chosen parameters for our model. 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 values to based on existing evidence or steady-state considerations, and then the ones we choose using indirect inference and the state-level 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. All our parameter choices are reported in Table 3.

We pin down three additional parameters using steady state considerations. The discount factor β is chosen so that the steady state real 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 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 model, these two last two targets imply an aggregate debt to income ratio of $2.5 \times 0.29 = 0.725$.

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 preference shocks, σ_η ; 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 conduct indirect inference by estimating auxiliary panel regressions in both our model and the data, and choosing parameters in the model to ensure that the coefficient estimates in these auxiliary regressions are as close as possible in the model and in the data.¹³ We consider a set of panel regressions using state-level data on household debt from the FRB New York Consumer Credit Panel,¹⁴ house price data from the FHFA, as well as data on employment, wages and consumption expenditures from the BEA.¹⁵ Let s denote a particular state and let t denote time. In actual and model-simulated data, we estimate panel regressions of the form

$$\ln Y_{s,t} = d_s + f_t + \chi_1 \text{Debt}_{s,t} + \chi_2 \text{Debt}_{s,t-1} + u_{s,t} \quad (45)$$

where Y is: 1) the employment/population ratio, 2) wages, 3) consumption spending or 4) house prices. We include state fixed effects d_s and time fixed effects f_t . In all these regressions, $\text{Debt}_{s,t}$ 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 2001 population.

Our choice of these auxiliary regressions is motivated by the strong correlation between HH debt and state-level variables documented by Mian and Sufi. To be clear, the regressions in (45) are not meant to capture any particular causal relationship, but rather the dynamic pattern of the correlations between these variables. As Table 2 illustrates, the R^2 in these regressions is high, ranging 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: not including the time effects f_t in these regressions reduces the R^2 in these regressions by only 0.15 to 0.20. Thus, about half of the variation in state-level wages, consumption and employment is associated with variation in debt across states. Figure 6a illustrates this pattern for a subset of states in the data by plotting the actual employment-population ratio in each state against the fitted values from the estimates of (45). The two track each other closely.

We estimate identical panel regressions in our model by choosing the path for credit shocks $\varepsilon_{s,t}$ for each period and state to ensure that the model matches the path for debt in the data perfectly.¹⁶ Thus, the independent variables on the right hand side of (45) are the same in both

¹³See, for example, Guvenen and Smith (2014) for a more detailed description of the approach.

¹⁴We include credit card debt, auto loans and student loans, in addition to mortgage debt in our measure of household credit.

¹⁵See the Appendix for a more detailed description of the data we use.

¹⁶Our model is at a quarterly frequency while the data is annual, so we linearly interpolate the data to estimate a quarterly series of shocks, and then estimate the auxiliary model at the original annual frequency.

our model and in the data. We then choose parameters in the model to ensure that the fitted values produced by the auxiliary regressions using model-simulated data are as close as possible to those in the U.S. data. More precisely, let $\hat{y}_{s,t}^{\text{model}} = \chi_1^{\text{model}} \text{Debt}_{s,t} + \chi_2^{\text{model}} \text{Debt}_{s,t-1}$ be the fitted values of each of the four series in the model and $\hat{y}_t^{\text{data}}(s)$ be the fitted values in the data. We choose the six parameters of the model to minimize

$$\sum_{k=1}^4 \sum_{s=1}^{50} \sum_{t=2002}^{2012} \left(\hat{y}_{s,t}^{k,\text{model}} - \hat{y}_{s,t}^{k,\text{data}} \right)^2. \quad (46)$$

where k denotes the four different state-level variables.

Tables 1 and 2 illustrate how well we do in matching the patterns in the data. Recall that our model is over-identified – we have six parameters to match eight coefficients in the auxiliary regressions, yet the model does a reasonably good job of matching the coefficients and thus fitted values in the data. The column labeled ‘ R_{Model}^2 ’ in Table 2 shows that our measure of the model’s goodness of fit (1 - the sum of squared deviations of $\hat{y}_{s,t}^{k,\text{model}}$ from $\hat{y}_{s,t}^{k,\text{data}}$, scaled by the sum of squares of $\hat{y}_{s,t}^{k,\text{data}}$) is equal to 0.96 for employment and consumption, 0.98 for wages and 0.87 for house prices. Consequently, the correlation between the fitted values in the model and the data is nearly one, as is their relative standard deviation. Figure 6b illustrates this point, by contrasting the fitted values in the model and in the data for a subset of the states. Overall, 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.

Table 3 reports the estimated parameter values. 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 about 20 quarters.¹⁷ 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 or refinance their mortgages, as well as borrow using shorter maturity home equity lines of credit.

We also find 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_η is equal to 7.49. Intuitively, although debt and house prices have fluctuated a lot in the cross-section and in the time series, the debt to housing ratio was relatively stable, as pointed out by Justiniano, Primiceri and Tambalotti (2015). The model thus requires small changes in the loan-to-value ratio to account for the patterns in the data.

¹⁷We follow Hatchondo and Martinez (2009) in defining duration as the weighted average maturity of cash flows. This is given by $\frac{1}{q} \sum_{t=1}^{\infty} t \left(\frac{1}{1+r} \right)^t \gamma^{t-1} = \frac{1+r}{1+r-\gamma}$.

Notice also that we estimate a relatively strong degree of *real wage rigidity*, in that the elasticity of substitution between labor varieties is high, $\psi = 5.4$. This is necessary to account for the fact that wages in the data are about one-half as volatile as employment (see 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 consistent with the evidence that non-tradable distribution services are an important component of the cost of retailing even for highly tradable goods.¹⁸

The right column of Table 3 also reports the standard errors around these estimates, which we have computed by bootstrapping and resampling individual states in the panel with replacement. These are low, suggesting that the parameters are well-identified.

The Appendix provides a detailed discussion about what features of the data identify each parameter. Here we briefly summarize that discussion. The degree of demand uncertainty, α , is identified by the relative variability of consumption and debt. The degree of real rigidity, ψ , is identified by the relative volatility of wages. An island's degree of openness, ω , is identified by the relative volatility of employment. Finally, the parameters γ and ρ , determining the duration of securities and the persistence of shocks, are identified by the relative volatility and correlation of household debt and housing prices.

4.3 Steady State Implications

We conclude this section by briefly discussing the steady-state properties of the model. Given the parameters we estimate, the gap between the discount rate and the interest rate is about 1.5% per year. The ratio of transfers to minimum consumption, x/\underline{c} is equal to 2.1, implying that $(x/\underline{c})^{-\alpha} = 1.7\%$ of household members are liquidity constrained. Finally, the implicit employment tax levied by the liquidity constraints is small: employment would increase by only 0.1% in their absence. Overall, the estimates of the model imply that the steady state distortions arising from the liquidity constraints are not too large, and mainly manifest themselves in a 1.5% gap between the discount rate and the interest rate.

5 Aggregate Implications

We next study the aggregate implications of the model. Our focus is on understanding the role of HH debt in shaping the dynamics of employment during the Great Recession. We first describe our solution method and explain why the marginal effect of HH debt shocks is solely a function of the expected ZLB duration. We then show that the model predicts that the observed

¹⁸See the evidence in Burstein, Eichenbaum and Rebelo (2005).

decline in household credit after 2007 leads to a modest decline in the natural rate of interest, of about 1.5%, and cannot, on its own, trigger the ZLB. We finally study the interaction between household credit shocks and an additional shock, to interest rate spreads, that triggers the ZLB.

5.1 Solution Algorithm

We use an algorithm developed by Jones (2015), a variant of methods developed by Eggertsson and Woodford (2003) and Guerrieri and Iacoviello (2015).¹⁹ We first log-linearize all equations of the model in the absence of the ZLB. We then define two regimes, one in which the ZLB binds and another one in which it does not. The algorithm is based on a piece-wise linear solution of the equilibrium conditions in these two regimes, under the assumption that agents observe the credit shock in each period, but believe that no other shocks are possible in the future.

Given a conjectured date T at which the ZLB will stop binding, we iterate backwards by using either the equilibrium conditions under the ZLB regime or the non-ZLB regime, depending on which one is conjectured to be in effect (note that we allow for the possibility that a shock at t triggers the ZLB at some *future* period). This backward recursion gives a path for all variables of the model, including the interest rate, which we use to update the initial conjecture until convergence. Once the algorithm has converged, the solution of the model is

$$\mathbf{x}_t = \mathbf{J}_t + \mathbf{Q}_t \mathbf{x}_{t-1} + \mathbf{G}_t \epsilon_t,$$

where \mathbf{x}_t collects the endogenous aggregate variables and ϵ_t collects the aggregate shocks. Since we compute this solution by iterating backward starting from the date T at which the ZLB is conjectured to stop binding, the marginal impulse response to a given shock ϵ_t , encoded in \mathbf{Q}_t and \mathbf{G}_t , is only a function of the date T at which the economy exits the ZLB, in addition to the primitive parameters of the model. We use this result in our identification scheme below.

5.2 Role of Household Debt in the Great Recession

As we show in the Appendix, the U.S. HH debt to income ratio exhibits a trend. Since we do not allow for trends in our model, we detrend the data by subtracting a linear trend and then smooth the resulting series to eliminate high-frequency noise. We finally use a Kalman filter in order to back out the path for aggregate credit shocks that the model requires to match the resulting debt-to-income series and study the model's implications for various macro variables.

Figure 7 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 HH debt is accompanied by a fairly

¹⁹We are extremely grateful to Callum Jones for sharing his codes and helping us with the computations.

modest rise and then drop in the natural interest rate, from 1.5% in 2001 to 2.8% at the peak in 2008, to 1.4% by the end of 2013. Note also the almost perfect correlation between the natural rate and debt-to-income implied by the model: the two rise and fall in lockstep. The nominal interest rate mimics a similar pattern, owing to the Taylor rule: it increases from about 3.6% in 2001 to 5.3% at the peak in 2007 to about 3% by 2013.

Consider next the response of employment, which we measure as the total number of employees on non-farm payrolls, scaled by the U.S. population.²⁰ As the lower-left panel of Figure 7 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 only one-fifth of the actual decline. Moreover, this drop in employment is gradual, mimicking the gradual reduction in debt, as opposed to sudden, as observed in the data. Notice also that inflation is fairly volatile in our model, owing to the relatively steep slope of the Phillips curve implied by our estimates. This result echoes the findings of Beraja, Hurst, Ospina (2015). The fact that the New Keynesian model predicts a large drop in inflation is well-documented, and several explanations have been proposed recently (see for instance Coibion and Gorodnichenko (2015), Del Negro et al. (2015), or Gilchrist et al. (2015)).

The bottomline of this exercise is that changes in HH debt in our model cannot, on their own, to generate the large reduction in employment and interest rates observed in the data. This result is a direct implication of the relatively low uncertainty of idiosyncratic shocks we have estimated using state-level data. As we show in the Appendix, had we assumed a greater volatility of idiosyncratic shocks, the model would have produced much larger employment responses. Such a parameterization would imply, however, a much greater sensitivity of real variables to credit shocks at the state-level and would thus be counterfactual.

5.3 Robustness Checks

We considered three robustness checks. First, we studied an economy with a lower elasticity of intertemporal substitution. Second, we considered an alternative estimation, using state-level wage data adjusted for worker composition from Beraja, Hurst and Ospina (2015). Third, we studied a version of our model with a lower steady-state equilibrium interest rate. The Appendix discusses these robustness checks in some detail. Here we briefly summarize that discussion.

Lower Elasticity of Intertemporal Substitution We assumed more general CRRA preferences with an elasticity of intertemporal substitution of $1/3$ and re-estimated all the parameters

²⁰We detrend the log of the resulting series using a linear trend.

of the model using state-level data. As reported in the Appendix, this economy is characterized by a stronger precautionary savings motive, as evidenced by a 3% gap between the discount rate and the steady state real interest rate. We find that both the nominal and natural rate of interest are twice more volatile than in the Benchmark economy with logarithmic preferences. Importantly, HH credit shocks are still not sufficient to trigger the ZLB. Even though the interest rate needs to fall by more now during the recession, the level of the nominal interest rates at their peak is now higher (almost 7% in 2007), owing to the large rise in the natural interest rate during the boom. Monetary policy thus has more room to maneuver during the recession.

Beraja, Hurst and Ospina (2015) Wage Data We originally used a measure of wage data from the BEA. In a recent paper Beraja, Hurst and Ospina (2015) construct an alternative wage panel using the Current Population Survey,²¹ which takes into account variation in the composition of the labor force, and show that it exhibits considerable variation. We have re-estimated our model using this new wage data, and found that a lower degree of real wage rigidity ($\psi = 2.3$) is needed to match the greater volatility of these wage series. This version of the model predicts even more modest movements in employment than our Benchmark model.

Lower Steady State Real Interest Rate We have also re-estimated the model by targeting a real interest rate of 0.5% (compared to 2% in our Benchmark model). The model’s implications for real variables are virtually unchanged. Although both the natural and nominal interest rates shift down by 1.5%, the nominal interest rate continues to be above zero and so monetary policy can offset the household credit shocks just as well as in our Benchmark model.

5.4 Household Debt at the Zero Lower Bound

The decline in the natural rate of interest predicted by our model is fairly persistent. Thus, if household credit shocks are accompanied by other shocks that reduce the Fed’s ability to cut interest rates because of the ZLB, the resulting effects on output can be much greater. We illustrate this point by adding a shock that drives a spread between the federal funds rate f_t and the nominal interest rate i_t at which agents trade securities:

$$i_t = f_t + \xi_t,$$

where the spread itself follows an autoregressive process:

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

²¹We are grateful to Erik Hurst for sharing the data with us.

Such a spread can arise due to frictions faced by financial intermediaries, as in the work of Gertler and Karadi (2011) and Gertler and Kiyotaki (2010).²² We modify the Taylor rule by assuming that the Fed Funds rate reacts to changes in spreads and offsets them if possible:

$$1 + f_t = \max \left[(1 + f_{t-1} + \xi_{t-1})^{\alpha_r} \left[(1 + \bar{i}) \pi_t^{\alpha_\pi} \left(\frac{y_t}{\bar{y}} \right)^{\alpha_y} \left(\frac{y_t}{y_{t-1}} \right)^{\alpha_x} \right]^{1-\alpha_r} - \xi_t, 1 \right]. \quad (47)$$

Clearly, HH debt shocks and shocks to spreads reinforce each other in the presence of the ZLB. A sufficiently large increase in spreads may force monetary policy to cut nominal rates by a lot and thus leave it unable to respond further to household debt shocks. Because of this non-linearity, the marginal effect of household debt shocks on output and employment depends on what one assumes about the size and persistence of the shocks to spreads. We illustrate this point in Figure 8, in which we consider a spread shock small enough that it would have no effect on output in the absence of HH debt shocks because it would be perfectly offset by the modified Taylor rule above.

Figure 8 reports the response of the economy for two experiments, in which the persistence ρ_ξ of the spread shock is either low or high. When ρ_ξ is not very large, the ZLB is not expected to last long and the marginal contribution of shocks to HH debt is nearly as small as in the absence of spread shocks. In contrast, when ρ_ξ is larger, the ZLB is expected to last for a long time and HH debt shocks trigger larger declines in employment. Intuitively, the size of the employment drop is proportional to the expected sum of current and future deviations of the real rate from the natural rate. The longer the ZLB is expected to last, the larger this sum is, and the deeper the recession.

This example suggests that quantifying the impact of HH debt on macroeconomic dynamics at the ZLB is challenging, since the marginal impact of HH debt shocks is sensitive to what one assumes about the additional shocks. Our discussion in Section 5.1 made it clear, however, that the marginal contribution of the HH credit shocks is solely a function of the expected duration of the zero lower bound. We use this insight to discipline our analysis below.

We next use the Morgan Stanley Months to First Rate Hike (MSM1KE) series, which measures the real time forecasts of the duration of the ZLB, to calibrate the persistence of the spread shock in the model. As before, we extract the HH debt shocks to match the actual HH debt series, and impose a spread shock that brings down the nominal interest rate to 0 in 2008Q3, but does not cause a change in employment on its own. We choose $\rho_\xi = 0.982$ to ensure that the model's expected ZLB duration matches that in the data.²³

²²Gilchrist and Zakrajek (2012) provide evidence that this spread was indeed high during the Great Recession.

²³See the Appendix for the series predicted by the model and the data.

Figure 9 shows the responses of macro variables in this experiment. Because HH debt declines gradually, the model predicts a decline in employment that reaches its trough only in 2012 compared to the much more immediate decline in the data. The overall peak-to-trough drop is nevertheless almost as large in the model as in the data. The model thus accounts well for the slow employment recovery, but not for its collapse at the onset of the recession. Overall, the average deviation of employment from trend in the data was about 4.8% from 2008 to 2013, while the corresponding number from the model was about 2.7%. The model thus accounts for about 55% of the overall decline in employment during this period.

Notice finally, in the lower-right panel of Figure 9 that inflation now falls much more in the model compared to the data. It turns out that adding an additional sequence of shocks to the firms' desired markups that remedies this problem does not change the model's predictions for employment. Figure 10 illustrates this point. This result reflects the fact that in this experiment we continue to target a given sequence of expected ZLB durations. Conditional on matching these durations, the marginal effect of HH credit shocks on employment is unchanged. For this reason, adding alternative shocks or allowing the monetary authority to engage in forward guidance would not alter the structural matrices \mathbf{Q}_t and \mathbf{G}_t and thus the model's predictions for how employment changes in response to household credit shocks.

6 Conclusions

A popular account of the U.S. Great Recession is the view that declines in housing wealth and households' ability to borrow 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 co-moves strongly with changes in house prices and household debt in the cross-section of U.S. regions. This paper proposes a theory that captures this view by introducing a role for credit in alleviating household liquidity constraints. We estimate the model using an indirect inference approach and data from a panel of U.S. states. We then study the model's implications for aggregate times series.

Our model predicts that changes in household debt of the magnitude observed during the Great Recession do not, by themselves, generate large movements in the natural rate of interest. Household deleveraging lower this rate by only 1.5%, and monetary policy can offset it by reducing the nominal rate by an equivalent amount. We conclude that additional shocks are necessary to explain why nominal rates fell much more in this period. This conclusion is robust to a number of perturbations of the model.

On the other hand, conditional on the ZLB binding, household credit shocks have a much larger impact on real activity. We show that, despite the important non-linearities brought about by the ZLB, we can identify the marginal contribution of household credit by requiring the model to reproduce the expected duration of the ZLB in the data. When we do so, we find that household credit shocks account for about 55% of the overall decline in employment during the Great Recession. While the model fails to account the large immediate drop in employment during the 2009-2010 period, the persistent nature of household deleveraging allows the model to reproduce 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] Challe, Edouard, Julien Matheron, Xavier Ragot and Juan Rubio-Ramirez. 2016. "Precautionary Savings and Aggregate Demand," Working Paper.
- [8] Christiano, Lawrence, Martin Eichenbaum and Sergio Rebelo. 2011. "When is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119(1) 78-121.
- [9] 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.
- [10] Cooley, Thomas and Gary Hansen. 1991. "The Inflation Tax in a Real Business Cycle Model," *American Economic Review*, 79(4), 733-748.
- [11] Del Negro, Marco and Giannoni, Marc and Schorfheide, Frank. 2015. "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7(1), 168-196.

- [12] Eggertsson, Gauti and Michael Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 2003(1), 139-233.
- [13] 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.
- [14] 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*.
- [15] 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.
- [16] 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.
- [17] Gertler, Mark and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 58(1), 17-34.
- [18] Gilchrist, Simon and Zakrajšek, Egon. 2012. "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102(4), 1692-1720.
- [19] Gilchrist, Simon and Schoenle, Raphael and Sim, Jae and Zakrajšek, Egon. 2015. "Inflation Dynamics During the Financial Crisis," *Working Paper*
- [20] Giroud, Xavier and Holger Mueller. 2015. "Firm Leverage, Consumer Demand, and Employment Losses during the Great Recession," *Quarterly Journal of Economics*, forthcoming.
- [21] Gorea, Denis and Virgiliu Midrigan. 2015. "Liquidity Constraints in the U.S. Housing Market," NYU Working Paper.
- [22] Guerrieri, Luca and Matteo Iacoviello. 2015. "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily." *Journal of Monetary Economics*, 70, 22-38.
- [23] Guerrieri, Veronica and Guido Lorenzoni. 2015. "Credit Crises, Precautionary Savings, and the Liquidity Trap," Chicago Booth Working Paper.
- [24] Guvenen, Fatih and Anthony Smith. 2014. "Inferring Labor Income Risk and Partial Insurance from Economic Choices," *Econometrica*, 82(6), 2085-2129.
- [25] Hatchondo, Juan Carlos and Leonardo Martinez. 2009. "Long-Duration Bonds and Sovereign Defaults," *Journal of International Economics*, 79(1), 117-125.
- [26] Iacoviello, Matteo. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95(3): 739-764.

- [27] Jones, Callum. 2015. “Unanticipated Shocks and Forward Guidance At the Zero Lower Bound.” NYU Working Paper.
- [28] Justiniano, Alejandro and Giorgio Primiceri. 2008. “The Time Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98(3), 604-641.
- [29] Justiniano, Alejandro, Giorgio Primiceri and Andrea Tambalotti. 2015. “Household Leveraging and Deleveraging,” *Review of Economic Dynamics*, 18(1), 3-20.
- [30] Kaplan, Greg, and Gianluca Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4) 1119-1239.
- [31] Kaplan, Greg, Kurt Mitman, and Gianluca Violante. 2015. “Consumption and House Prices in the Great Recession: Model Meets Evidence,” NYU Working Paper.
- [32] Kehoe, Patrick, Virgiliu Midrigan and Elena Pastorino. 2016. “Debt Constraints and the Labor Wedge,” forthcoming, *American Economic Review Papers & Proceedings*.
- [33] 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.
- [34] Landvoigt, Tim, Monika Piazzesi and Martin Schneider. 2015. “The Housing Market(s) of San Diego,” *American Economic Review*, 105(4), 1371-1407.
- [35] Lucas, Robert Jr., 1990. “Liquidity and Interest Rates,” *Journal of Economic Theory*, 50(2) 237-264.
- [36] Lucas, Robert Jr. and Nancy Stokey, 2011. “Liquidity Crises.” Federal Reserve Bank of Minneapolis Economic Policy Papers.
- [37] Lustig, Hanno and Van Nieuwerburgh, Stijn. 2005. “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 60(3), 1167-1219.
- [38] Mendoza, Enrique 2010. “Sudden Stops, Financial Crises, and Leverage,” *American Economic Review*, 100(5), 1941-1946.
- [39] 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.
- [40] Mian, Atif, Kamalesh Rao and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128(4), 1687-1726.
- [41] Mian, Atif, and Amir Sufi. 2014. “What Explains the 2007-2009 Drop in Employment?” *Quarterly Journal of Economics*, 82(6), 2197-2223.
- [42] Nakamura, Emi and Jón Steinsson. 2014. “Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions.” *American Economic Review*, 104(3), 753-792.
- [43] Werning, Ivan. 2012 “Managing A Liquidity Trap: Monetary and Fiscal Policy.” MIT Working Paper.

Table 1: Coefficients in Auxiliary 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

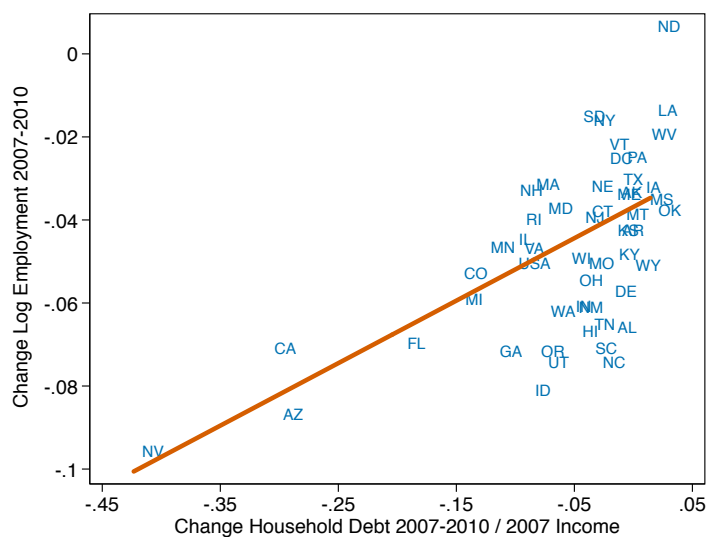
Table 2: Model and Data Fit

	R^2_{Model}	R^2_{Data}	R^2_{Data} (no time effects)
Employment	0.96	0.69	0.55
Consumption	0.96	0.71	0.56
Wages	0.98	0.56	0.35
House Prices	0.87	0.87	0.71

Table 3: Parameter Values

<i>Assigned</i>			<i>Chosen with indirect inference</i>			
					s.e.	
	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_η	7.49	(1.37)	Volatility housing shocks
σ	0.5	Elasticity traded/non-traded	ψ	5.38	(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	$-\ln(\beta)$	3.49		Annual discount rate, %
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

Figure 1: Employment and Household Debt



Notes: The x axis reports the change in the real value of household debt from 2007 to 2010 scaled by 2007 income. The y axis reports the log change of the employment to population ratio in a given state. We exclude construction-sector employment from our measure of employment. See the Appendix for the data sources.

Figure 2: Timing

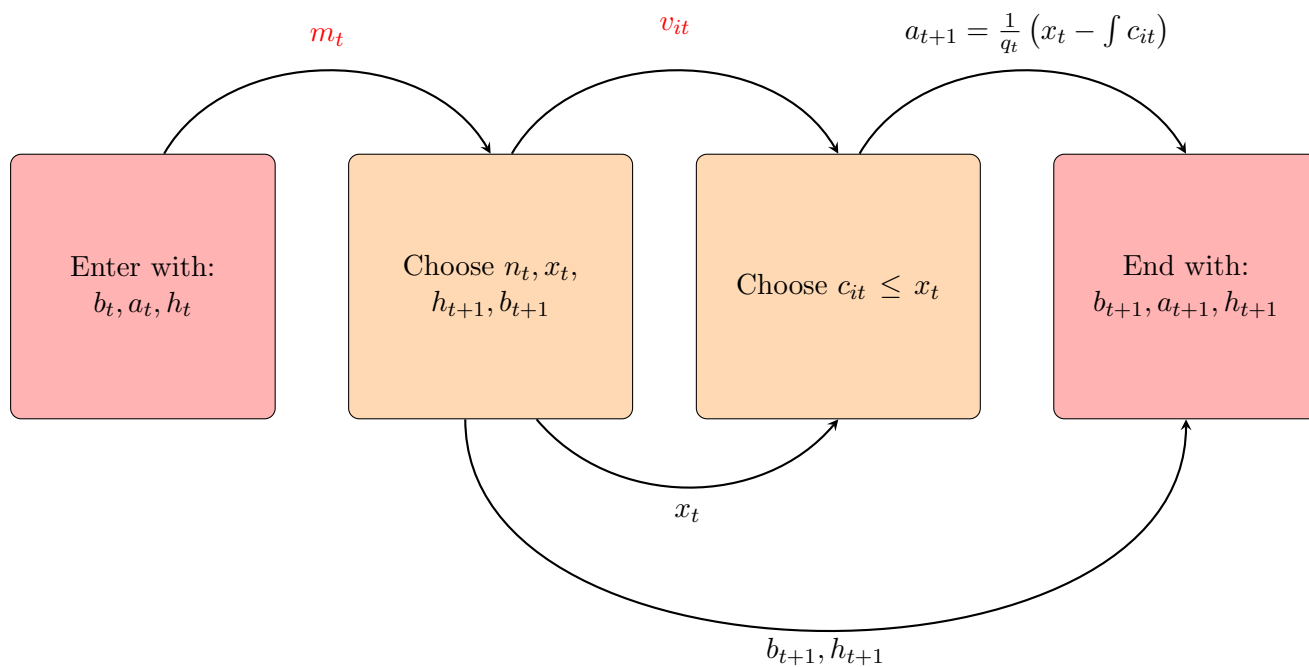


Figure 3a: Steady State Interest Rate, High v Uncertainty

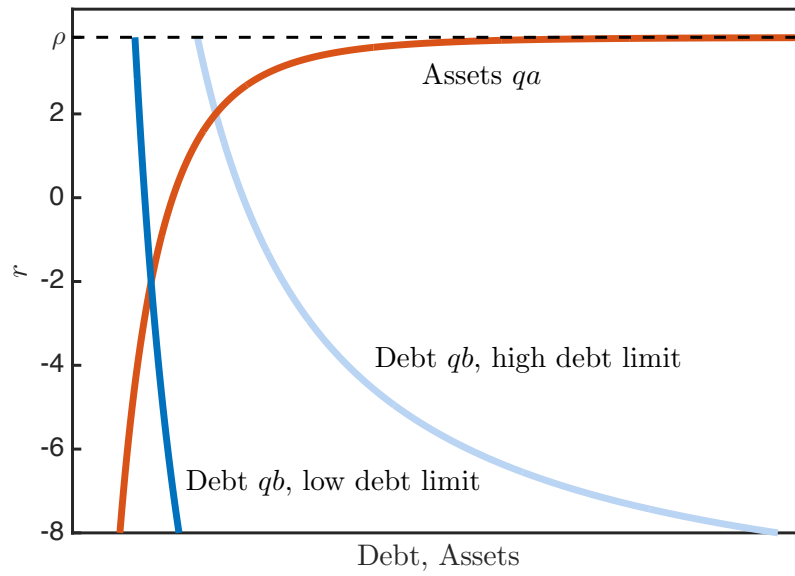


Figure 3b: Steady State Interest Rate, Low v Uncertainty

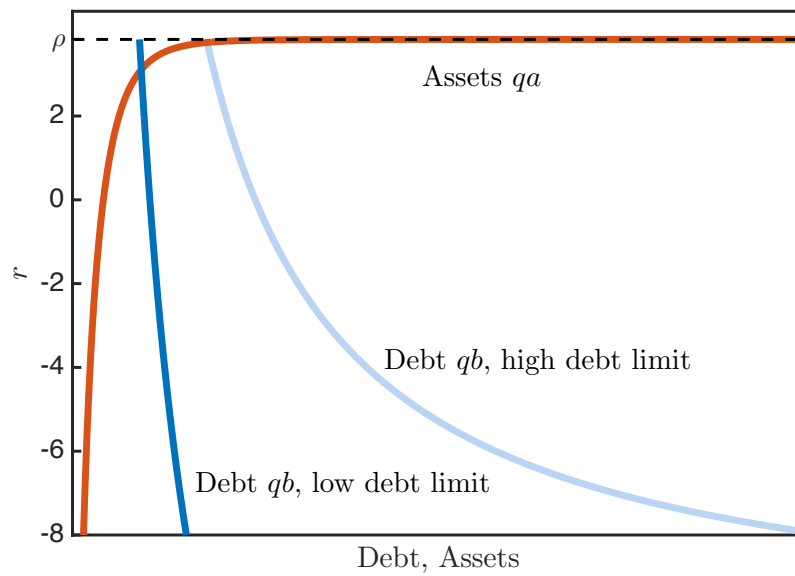


Figure 4: Impulse Response to a Credit Shock, Baseline Closed Economy Model

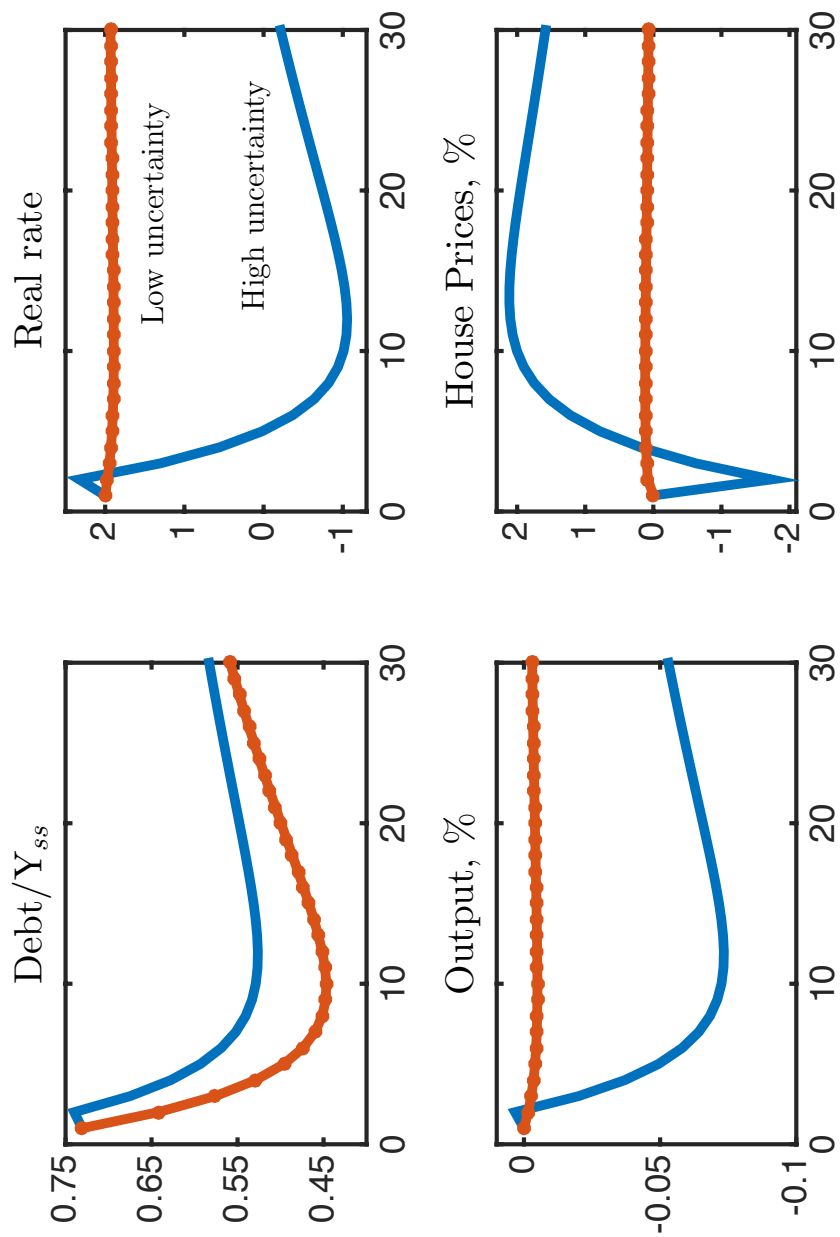


Figure 5a: Island Credit Tightening, Flexible Prices, High v Uncertainty

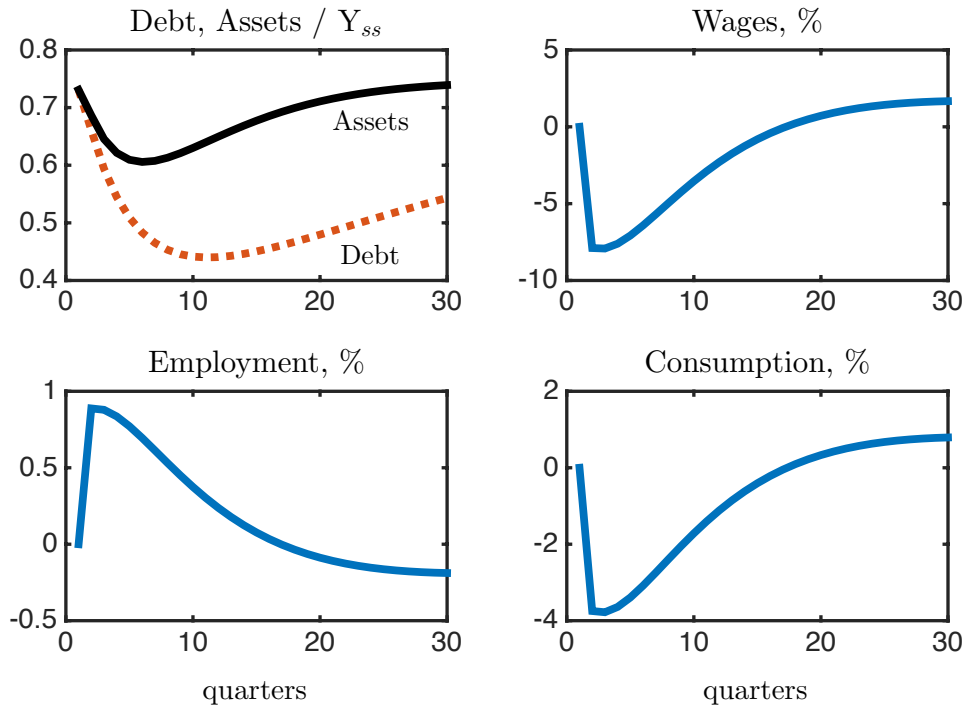


Figure 5b: Island Credit Tightening, Flexible Prices, Low v Uncertainty

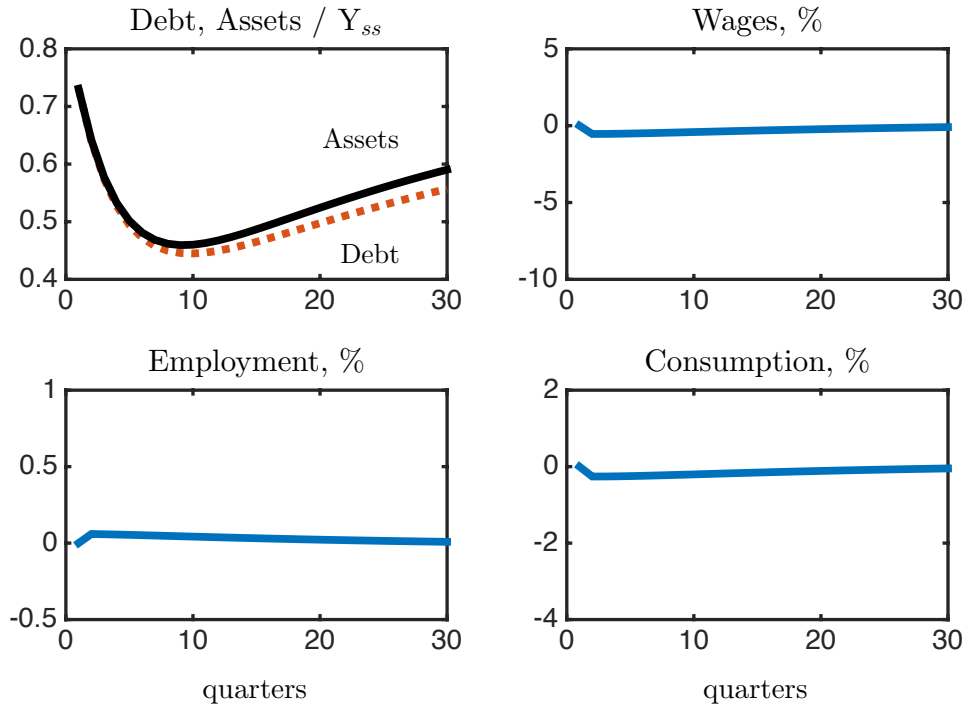


Figure 5c: Island Credit Tightening, Sticky Prices

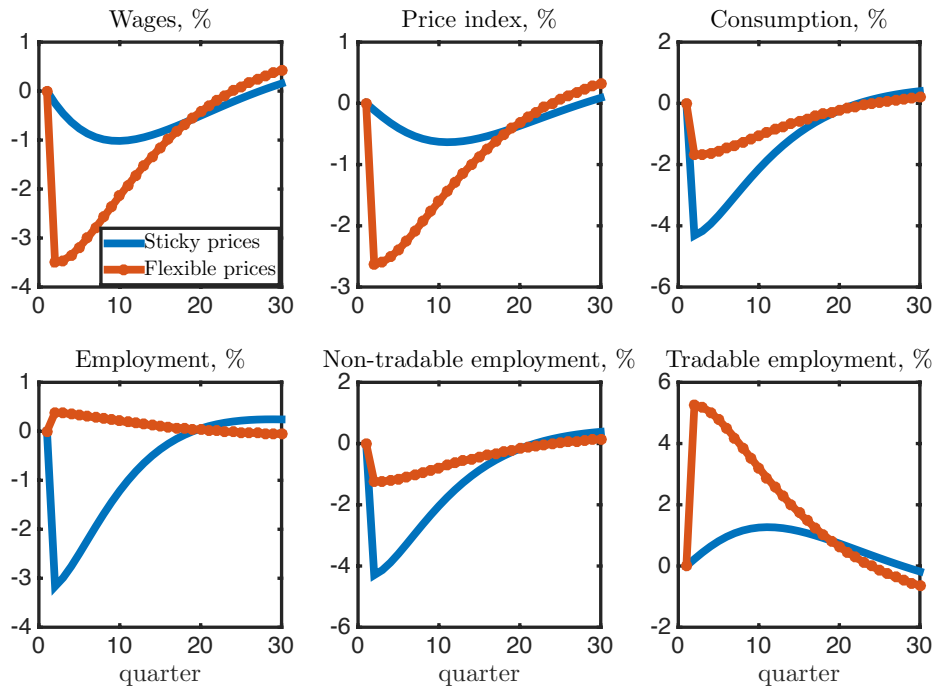


Figure 5d: Comparison of Island and Aggregate Responses

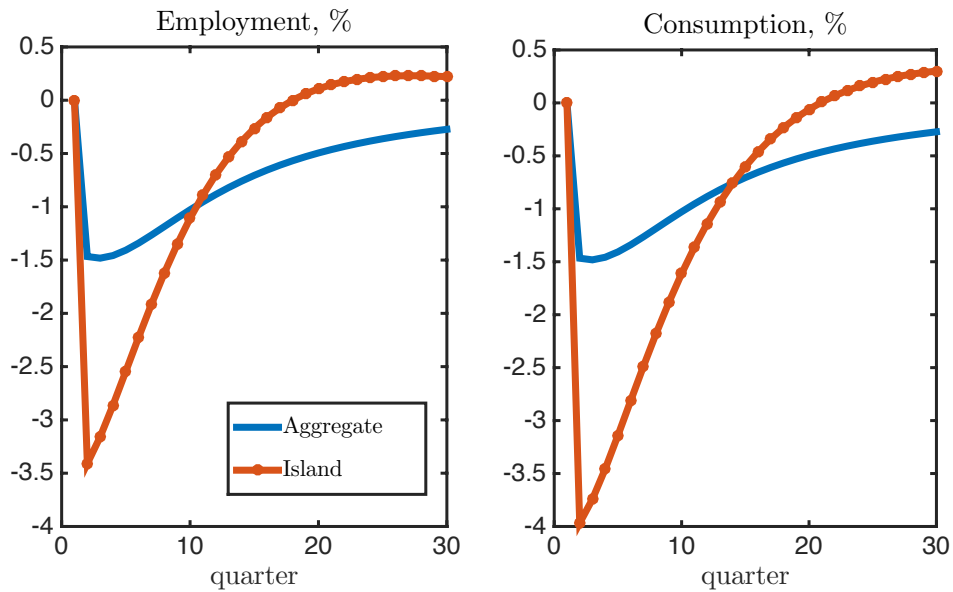


Figure 6a: Employment: Data vs. Fitted Values

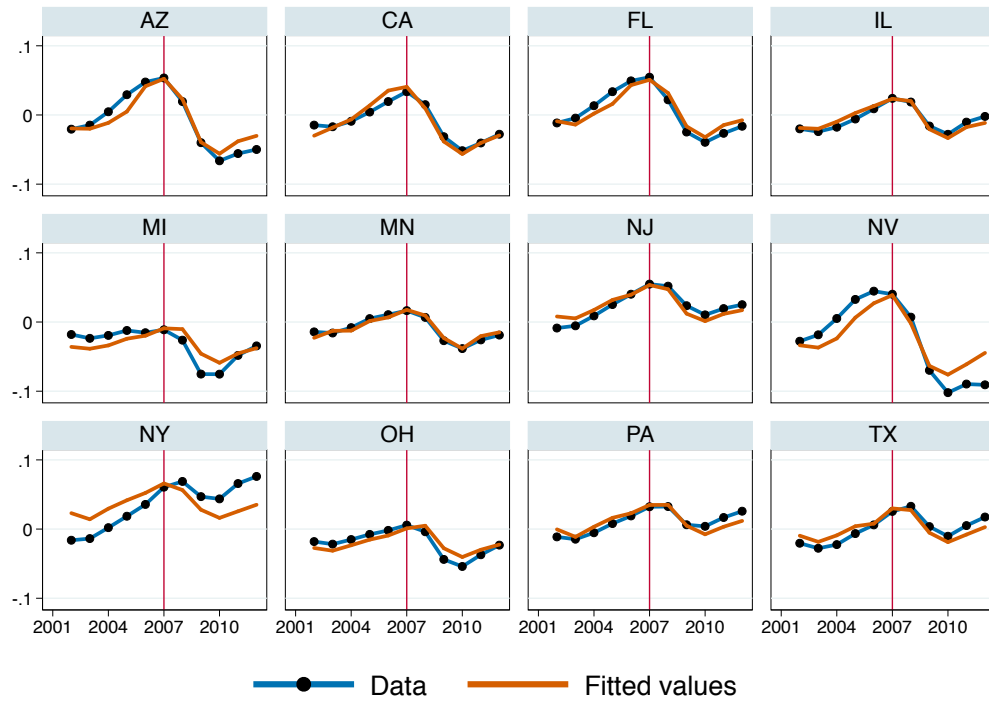


Figure 6b: Employment: Fitted Values Model vs. Data



Figure 7: Dynamics of Model that Matches U.S. Household Debt

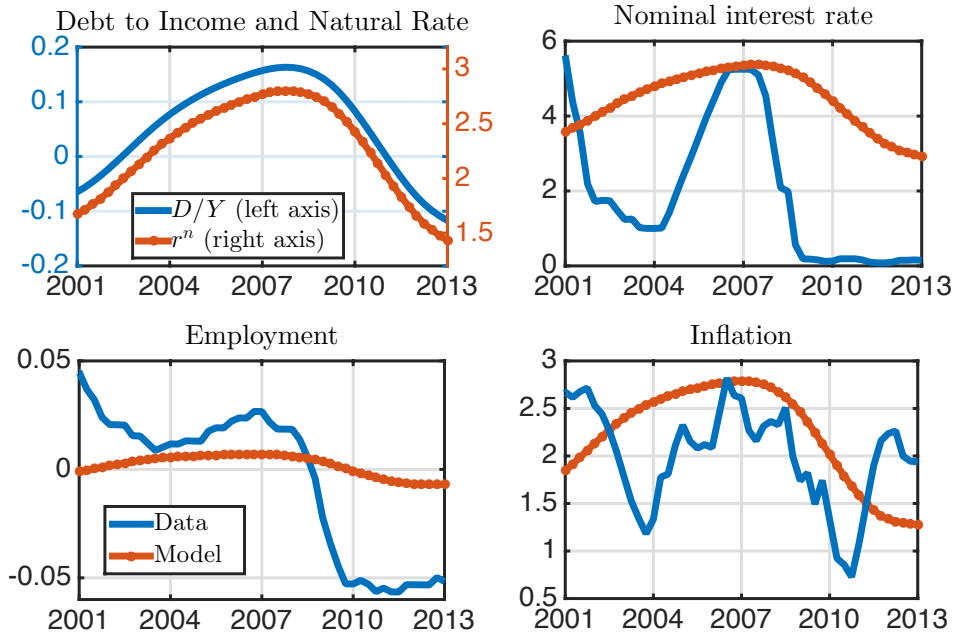


Figure 8: Role of Persistence of Spread Shocks

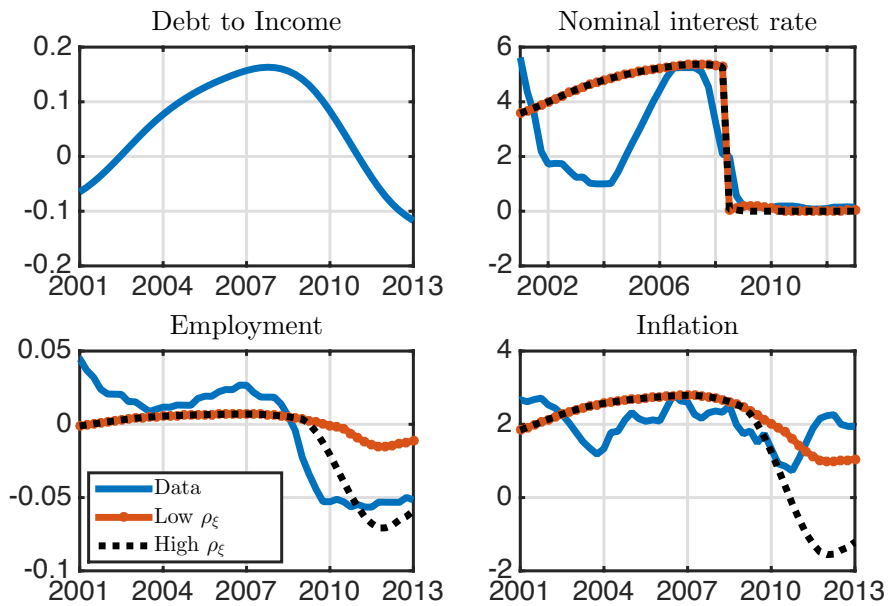


Figure 9: Dynamics of Model that Matches Expected ZLB Duration

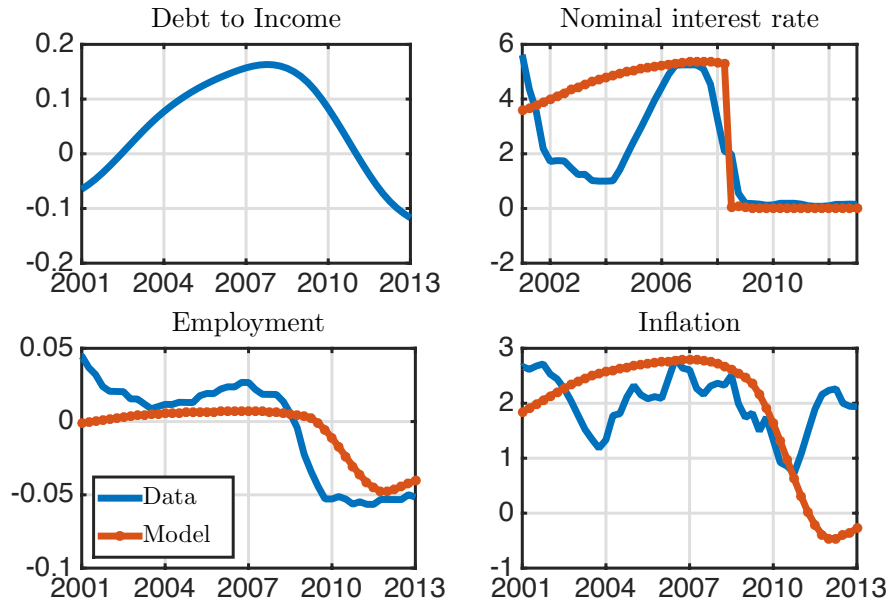


Figure 10: Matching Inflation Dynamics

