WP/18/194

IMF Working Paper

Household Leverage and the Recession

by Callum Jones, Virgiliu Midrigan and Thomas Philippon

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

WP/18/194

IMF Working Paper

Household Leverage and the Recession

by Callum Jones, Virgiliu Midrigan and Thomas Philippon

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Research Department

Household Leverage and the Recession

Prepared by Callum Jones, Virgiliu Midrigan, Thomas Philippon

Authorized for distribution by Luis Cubeddu

August 2018

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

We evaluate and partially challenge the ‘household leverage’ view of the Great Recession. In the data, employment and consumption declined more in states where household debt declined more. We study a model where liquidity constraints amplify the response of consumption and employment to changes in debt. We estimate the model with Bayesian methods combining state and aggregate data. Changes in household credit limits explain 40 percent of the differential rise and fall of employment across states, but a small fraction of the aggregate employment decline in 2008-2010. Nevertheless, since household deleveraging was gradual, credit shocks greatly slowed the recovery.

JEL Classification Numbers: E2, E4, E5, G0, G01.

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

Authors’ E-Mail Addresses: cjones3@imf.org, virgiliu.midrigan@nyu.edu,
tphilipp@stern.nyu.edu

1 Introduction

A salient feature of the Great Recession is that U.S. regions that experienced the largest declines in household debt also experienced the largest declines in consumption and employment. Figure 1 illustrates this pattern, originally documented in a series of papers by Mian and Sufi, by plotting the change in consumption and employment against changes in household debt during the 2002 to 2007 boom, as well as during the 2007 to 2010 recession.¹

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 and trade frictions, to a reduction in employment.

The goal of our paper is to quantitatively evaluate the household leverage view of the recession. We ask: how important were shocks to household credit limits in generating the macroeconomic dynamics across regions and in the aggregate? The evidence in Figure 1 does not directly answer this question for two reasons. First, individual states and the U.S. aggregate are subject to a variety of shocks in addition to changes in credit standards. These shocks may themselves generate the correlation between state-level employment and household credit documented in Figure 1. For example, changes in a state's productivity may lead to changes in that state's employment, house prices and the amount households borrow. Second, the response of employment or consumption in an individual state to a given credit shock may be very different than the response to the same shock in the aggregate. Moving from the cross-sectional to the aggregate elasticity requires a careful consideration of trade patterns and monetary policy.

We therefore use a model to isolate the effects of credit shocks and answer our question. We study a parsimonious and tractable model of a monetary union composed of a large number of regions that trade with each other. Households in each region face liquidity constraints that restrict their ability to use the equity in their homes to finance consumption. These constraints manifest themselves in a time-varying endogenous spread between the equilibrium interest rate and the households' rate of time preference. This wedge makes it optimal for households to borrow against the value of their homes, but their ability to do so is limited by a loan-to-value constraint which can change over time. Households in each region are subject to a number of additional exogenous shocks to productivity and preferences. All shocks have both a region-specific and an aggregate component.

We estimate the model and the shocks using Bayesian likelihood methods and data on

¹See Mian and Sufi (2011, 2014) and Mian, Rao and Sufi (2013) for additional evidence.

employment, consumption, wages, house prices and household debt at the state level, together with the corresponding series in the aggregate. An important challenge that arises in our estimation is that the policy interest rate was at the *zero lower bound* (ZLB) during part of the period that we study. We account for the non-linearities arising from this bound by using a piece-wise linear solution method developed by Eggertsson and Woodford (2003), Guerrieri and Iacoviello (2015), and Jones (2017). Even with this piece-wise linear approximation, a direct application of Bayesian methods is infeasible in our setting given the large number of regional and aggregate state variables in the model. We therefore propose a new methodology that allows us to combine both state and aggregate data to evaluate the likelihood function.

A popular approach to introducing a motive for household borrowing is to assume heterogeneity in agents' rate of time preference, as in the work of Iacoviello (2005) and a large literature that builds on it. Given our focus on estimating the model, we pursue a more flexible approach. We assume instead that agents are subject to idiosyncratic shocks to their liquidity needs. These shocks imply that the spread between the discount rate and the interest rate is a smooth function of the amount agents can borrow. This flexibility is useful in estimation because we do not want to impose a priori an important role for credit shocks in determining household-level consumption variation. Depending on how volatile individual agents' liquidity needs are, the model either collapses to a frictionless benchmark in which consumption does not respond to a tightening of credit constraints, or to an environment in which consumption is much more sensitive to how much an agent can borrow.

The liquidity constraints we model prevent households from converting their housing wealth into funds available for consumption. We do not think that this is a controversial assumption: if a household could pay, for example, an unexpected medical bill with the equity in their home, they would not need to borrow against their homes and a loan-to-value constraint would be of no consequence. A number of researchers have indeed documented that frictions that prevent households from extracting home equity to finance consumption are sizable: see the work of Kaplan and Violante (2014) and Gorea and Midrigan (2018). In contrast to this literature, which models liquidity constraints as arising from a fixed cost agents must pay to refinance their mortgage debt, considerations of computational tractability lead us to follow the approach of Lucas (1990) in assuming that agents can tap home equity at no cost, but must do so prior to the realization of an idiosyncratic taste shock. We assume that such shocks are i.i.d. and use a family construct to eliminate the distributional consequences of asset market incompleteness.²

²See also Lucas and Stokey (2011) who emphasize the role of liquidity frictions and Challe et. al. (2015) who use a family construct to characterize an economy with uninsurable unemployment risk.

We view this set of assumptions as a convenient modeling device that allows us to parsimoniously capture the role of household credit in the macroeconomy. Our model is flexible enough to capture many of the insights of richer non-linear Bewley models, yet sufficiently tractable to be used in a Bayesian estimation with data from a large number of regions.

The quantity of household debt in our economy has important 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, but 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, the economy's dynamics are highly non-linear: the marginal effect of a credit shock depends on how long the ZLB is expected to last. We thus use survey data on the expected duration of the ZLB together with key macroeconomic time-series to estimate the model and extract the aggregate shocks that generate the data. With these shocks in hand we then study the marginal contribution of credit shocks on aggregate variables.

We estimate three structural parameters that are key for determining the economy's response to a credit shock – the degree of idiosyncratic uncertainty, and the degrees of wage and price stickiness. We also estimate the persistence and volatility of the state- and aggregate shock processes. At the state-level, we use annual data on household debt, household spending, employment, wages, and house prices from 1999 to 2015. At the aggregate level, we use the aggregate counterparts of the state-level series, together with the Fed Funds rate and inflation from 1984 to 2015.

Directly combining aggregate series with their regional counterparts is computationally challenging in the presence of non-linearities arising from the ZLB, because the state-level observables are a time-varying function of aggregate shocks. To make our Bayesian estimation feasible we exploit the structure of the model and express state-level variables as deviations from the corresponding aggregates. Up to a first-order approximation, these deviations evolve according to a law of motion that is independent of whether the ZLB binds or not. This structure allows

us to separate the model's likelihood function into independent state-level and aggregate components. Nevertheless, both state and aggregate information shape the likelihood function and the posterior distributions of the model parameters. Intuitively our estimation exploits the differential rise and fall of individual states' spending, employment, debt, wages and house prices, in addition to the aggregate comovement of these series, to identify the structural parameters.

Using the model, we generate counterfactual series for employment and consumption using credit shocks alone and find that these account for about 40% of the differential rise and then fall in state-level employment and consumption during the 2002 to 2010 period. Our findings are thus consistent with those of Mian and Sufi (2011, 2014) who argued that household credit played an important role in determining variation in regional consumption and employment.

Although credit shocks have played an important role at the state level, they had a minor role in the aggregate at the onset of the Great Recession. Our model implies that changes in household credit limits generate a 0.8% decline in employment from 2008 to 2010, about one-tenth of the drop observed, despite the fact that monetary policy was constrained by the ZLB. This result reflects the gradual nature of household deleveraging which implies a gradual decline in the natural rate of interest. Nevertheless, since the deleveraging process was long-lasting, the natural rate of interest continued to decline until 2015. Our model implies that credit shocks alone account for about one-half of the employment gap in this latter time period and thus have greatly contributed to the slow recovery.³

Our model's implications for the aggregates are robust to a number of perturbations, including adding a construction sector, changing the duration of mortgage contracts, the parameters governing the strength of real rigidities and estimating the structural parameters using state data only. The model reproduces well the behavior of tradable and nontradable employment in individual states, as well as the medium-term fluctuations in mortgage rates in the data, series that were not directly used in estimation.

Related Work In addition to the work of Mian and Sufi and Iacoviello (2005), our paper is related to Guerrieri and Lorenzoni (2015) and Eggertsson and Krugman (2012) who also study the responses of an economy to a household credit crunch. 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 and aggregate 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). One puzzle in the literature is that

³See the literature on the slow recovery discussed by Fernald et al. (2017).

wages appear more flexible in the cross-section than in the time series. Our estimation strategy takes both into account and settles on an intermediate value. 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.

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).⁴

Relative to these papers, we make three contributions. First, we propose a parsimonious model in which liquidity constraints imply that credit constraints bind in equilibrium and generate a time-varying spread between interest rates and subjective discount rates. Despite its parsimony, the model is flexible enough to encompass a wide range of predictions regarding the role of credit for the macroeconomy and is thus suitable for empirical analysis.

Second, we develop a novel methodology that allows us to combine state-level and aggregate data to estimate a model with occasionally binding constraints using Bayesian techniques. We believe that this methodology can be fruitfully applied in other settings in which regional evidence can help discipline identification of key parameters. We view this as a contribution to the fast growing literature on the use of regional evidence to identify general equilibrium models, recently surveyed by Nakamura and Steinsson (2017).

Finally, we apply our model and methodology to quantify the role of credit shocks in accounting for state-level and aggregate employment during the Great Recession. In doing so, we keep our analysis transparent by solely focusing on the *household leverage* view of the Great Recession and abstract from other forces that have surely played an important role in this episode, such as bank- and firm-level financing constraints, changes in uncertainty, demographics, increase in markups and risk premia and other forces.⁵ Since we omit all these factors, they are implicitly accounted for by our rich set of exogenous shocks. Clearly, household credit may interact with a number of these forces in important ways. For example, a tightening of household credit that depresses home prices may trigger a wave of default, reducing the net worth of financial intermediaries.⁶ Though beyond the scope of this paper, we believe that jointly modeling household credit and other frictions is a fruitful area for future research.

⁴See the work of Lustig and van Nieuwerburgh (2005), Garriga et al. (2014), Favilukis et al. (2015), Burnside et al. (2015), Landvoigt et al. (2015).

⁵See also Alter et al. (2018), and Anderson et al. (2014).

⁶See Faria-e-Castro (2017) for a model along these lines.

2 A Simplified Closed-Economy Real Model

We first describe a simplified version of the model we will use in our empirical analysis below. We focus on the workings of the credit market and explain how liquidity constraints and idiosyncratic shocks generate an endogenous spread between the equilibrium interest rate and the household's rate of time preference. Owing to this spread, borrowing constraints bind in equilibrium. A tightening of borrowing constraints or an increase in the variance of idiosyncratic shocks increases this spread by reducing the equilibrium interest rate. The extent to which the interest rate responds to a credit tightening depends on the variance of idiosyncratic shocks: when these shocks are more dispersed, the equilibrium interest rate falls more after a credit tightening.

2.1 Setup

We describe the assumptions we make on technology and preferences, the nature of securities agents trade and 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 Agents trade 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.⁷ The household trades this security with perfectly competitive financial intermediaries.

It is convenient to describe a household's holdings of the security by recording the amount of coupon payments b_t that the household has to make period t . Letting l_t denote the amount of securities the household sells in period t , the date $t + 1$ coupon payments are

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

We let a_t denote the amount of coupon payments the household is entitled to receive in period t and b_t the amount it must repay. Thus $q_t a_t$ represents the household's total asset holdings (savings), while $q_t b_t$ represents her outstanding debt. As we explain below, the liquidity constraint that we assume pins down the gross positions a_t and b_t in isolation. Absent liquidity constraints, only the net position $a_t - b_t$ would be pinned down.

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

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

The left-hand side of this expression gives the household's uses of funds: transfers to individual members, x_t , and expenditure on housing, $e_t(h_{t+1} - h_t)$. The right-hand side is the amount of resources available: date t labor income, $w_t n_t$; the proceeds from issuing l_t units of new debt at price q_t , net of the required coupon payments on outstanding debt, b_t ; as well as the market value of the a_t securities the household owns. Each unit of the security pays a coupon of one unit of the consumption good and can be sold at a price γq_t .

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:

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

⁷See Hatchondo and Martinez (2009) and Arellano and Ramanarayanan (2012) who describe the properties of these securities in more detail.

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 simplified version of the model:

$$\log m_t = (1 - \rho_m) \log \bar{m} + \rho_m \log m_{t-1} + \sigma_m \varepsilon_{m,t}. \quad (8)$$

Our specification of the borrowing limit only restricts a household's ability to take on new loans and does not force prepayment of old debt. In our empirical work we resort to piecewise first-order perturbation techniques to solve and estimate the model. Given this solution method, the maturity of debt γ only affects equilibrium outcomes through its impact on the borrower's overall credit limit, $q_t b_{t+1} \leq \gamma q_t \gamma b_t + m_t e_t h_{t+1}$ which adjusts gradually in response to a shock to m_t . Introducing additional savings instruments of different maturities would thus be inconsequential.

Our model shares a number of similarities with cash-in-advance models, though there are several important differences. First, we assume that the household can access its date t labor income $w_t n_t$ immediately. Second, we allow agents to save in the interest-bearing assets at the conclusion of their shopping period. The only distortion here is that arising from the household's inability to borrow in order to smooth the marginal utility of consumption *across* its members.

Savings 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' taste for consumption v_{it} is then 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 $c_t(v)$, 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 taste shock v by imposing that x_t is not measurable with respect to v .

Let μ_t denote the shadow value of wealth, that is, the multiplier on the flow budget constraint (6); $\xi_t(v)$ denote the multiplier on the liquidity constraint (5) of a member with realization of the taste shock v ; and λ_t denote the multiplier on the borrowing constraint (7). Let R_{t+1} denote the realized one-period 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

$$\mu_t = \beta \mathbb{E}_t \mu_{t+1} R_{t+1} + \int_0^1 \xi_t(v) dF(v), \quad (11)$$

where $\mathbb{E}_t$ is the expectation operator over the realization of the aggregate credit shock $\varepsilon_{m,t+1}$.

This last 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 at the conclusion of the shopping period, they are 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 idiosyncratic taste shock v , these liquidity services are equal to the expected multiplier of the liquidity constraint of individual members.

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 preferences⁸ the choice of consumption reduces to

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

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

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

The supply of labor is given by

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

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

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

denote the consumption of a member with the lowest realization of the taste shock, $v = 1$. 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 household's subjective 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 to the spread between the discount rate and the interest rate, while the left-hand side is proportional to the fraction of constrained household members with sufficiently large taste shocks, $v > x_t/\underline{c}_t$. As the spread between the discount rate and the interest rate increases, it becomes costlier for households to save, so transfers fall relative to consumption and more members end up constrained.

Consider next the household's total consumption expenditures, $c_t = \int_0^1 c_t(v) dF(v)$. We have

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

⁸See the earlier draft of this paper for an analysis of the more general case of CRRA preferences.

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

Finally, letting

$$\Delta_t = \left(\beta \mathbb{E}_t \frac{\mu_{t+1}}{\mu_t} R_{t+1} \right)^{-1} - 1 \quad (20)$$

denote the spread between the discount rate and the interest rate, the savings to consumption ratio can be written as:

$$\frac{q_t a_{t+1}}{c_t} = \left(\frac{\alpha}{\alpha - 1} [(\alpha - 1) \Delta_t]^{\frac{1}{\alpha}} - \Delta_t \right)^{-1} - 1, \quad (21)$$

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 positive. A comparison of (11) and (12) shows that the multiplier on the borrowing constraint, λ_t , is positive as well. Intuitively, agents want to borrow because uncertainty regarding taste shocks pushes the interest rate below the rate of time preference.

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, the amount of debt in the economy is proportional to the value of houses, $qb = \frac{1}{1-\gamma} \bar{m} e h$. 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 makes borrowing less attractive. To see this, notice that in the steady state the Euler equation for housing (13) reduces to

$$e h = \frac{\bar{\eta}}{\mu} \frac{1}{\left(1 - \frac{\bar{m}}{1-\beta\gamma} \right) \rho + \frac{\bar{m}}{1-\beta\gamma} r}. \quad (22)$$

The value of housing is given by the marginal utility of housing, $\bar{\eta}/\mu$, discounted by a weighted average of the rate of time preference and interest rate with a weight that depends on how much the household can borrow. As long as $\rho < r$, an increase in the loan-to-value ratio $\bar{m}$ reduces the effective discount rate, raising house prices.

Figure 3 illustrate how the interest rate is determined in the steady state of the model. The left panel assumes a relatively large degree of idiosyncratic uncertainty. 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. The right panel 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 economy's responses to a one-time negative shock to credit, $\varepsilon_{m,t}$. This shock leads to a gradual reduction in household debt due to the long-term nature of securities and the assumption that the credit limit applies to new loans, not existing debt. We contrast the responses in economies with a high and low degree of idiosyncratic uncertainty.

Notice that the equilibrium interest rate falls after a credit tightening, more so when idiosyncratic uncertainty is high. In contrast, employment barely falls. Although a tightening of credit magnifies the consumption-leisure distortions, these are small here, as in cash-in-advance models.⁹

We provide further intuition for how the equilibrium interest rate responds to a credit shock by noting that equations (20) and (21) imply that the aggregate Euler equation in this economy is

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

where

$$\left(\frac{\alpha}{\alpha-1}[(\alpha-1)\Delta_t]^{\frac{1}{\alpha}} - \Delta_t\right)^{-1} - 1 = \frac{q_t b_{t+1}}{y_t} = \frac{\text{debt}}{\text{income}} \quad (23)$$

follows from asset market clearing and μ_t is inversely related to consumption. A decline in debt limits raises the average multiplier on liquidity constraints, Δ_t , more so when idiosyncratic volatility (α^{-1}) is high, thus acting like a natural rate shock in more traditional New Keynesian models.¹⁰ An important difference, however, is that in our model the natural rate is determined endogenously and responds to changes in credit limits.

3 An Island Monetary Economy with Price Rigidities

We next embed the credit frictions described above into a richer monetary economy with price and wage rigidities. The economy consists of a continuum of ex-ante identical islands of measure

⁹Cooley and Hansen (1991).

¹⁰See, for example, the work of Christiano, Eichenbaum and Rebelo (2011) and Werning (2012).

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.¹¹

Since we estimate the model using Bayesian likelihood techniques, we allow for several additional shocks. We include an additional shock to the rate of time preference of individual households. Recall that in our model credit shocks act much like discount rate shocks, in that they change the natural rate of interest. Since we do not want to assume that household credit shocks alone account for all variation in discount rates in the data, we allow for a separate exogenous source of movements in discount rates unrelated to household credit. We also include shocks to the household's disutility from work, their preference for housing, as well as productivity. Each shock has an island-specific and aggregate component. In addition, we assume aggregate shocks to the interest rate rule and to the aggregate inflation equation.

3.1 Setup

Household Problem The representative household on each island has preferences identical to those described above. 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 the union they belong to. The problem of a household that belongs to labor union ι is to

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\prod_{j=0}^{t-1} \beta_j(s) \right) \left[\int_0^1 v_{it}(s) \log(c_{it}(s)) di + \eta_t^h(s) \log(h_t(s)) - \frac{\eta_t^n(s)}{1+\nu} n_t(\iota, s)^{1+\nu} \right] \quad (24)$$

subject to the budget constraint

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

where $T_t(\iota, s)$ 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

¹¹See our Robustness section below for an extension of the model with a construction sector.

on island s exclusively own firms on that particular island. Here $\eta_t^h(s)$ and $\eta_t^n(s)$ affect the preference for housing and the disutility from work, while $\beta_t(s)$ is the household's one-period ahead discount factor. We assume that each of these preference shifters have an island-specific component and an aggregate component, all of which follow AR(1) processes with independent Gaussian innovations.

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

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

There are no barriers to capital flows, so all islands trade securities at a common price q_t . The credit limit $m_t(s)$ evolves as the product of an island-specific and aggregate component, both of which are AR(1) processes with Gaussian disturbances.

We introduce shocks to housing preferences because, as is well known, LTV shocks alone cannot generate movements in house prices as large as those in the data in this class of models. We thus assume shocks to both the LTV ratio as well as the consumer's preference for housing in order to allow the model to match the state-level and aggregate data on house prices.¹² We note, however, that changes in the preference for housing only matter through their effect on household credit limits since housing is separable in the utility function and the housing stock is in fixed supply.

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 :

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

¹²See 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. We think of preference shocks as a parsimonious way of capturing changes in risk premia or beliefs that trigger changes in house prices in these models.

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

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

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:

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

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 k is

$$y_t^T(k, s) = (p_t^T(k, s)/p_t^T(s))^{-\vartheta} y_t^T(s).$$

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, but is now subject to a productivity disturbance $z_t(s)$ that is common to both sectors, so that

$$y_t^j(s) = z_t(s) n_t^j(s), \text{ for } j \in \{T, N\}$$

so that the unit cost of production is simply $w_t(s)/z_t(s)$ in both sectors.

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

$$\max_{p_t^{T*}(s)} \sum_{j=0}^{\infty} \left(\lambda_p^j \prod_{i=0}^{j-1} \beta_{t+i}(s) \right) \mu_{t+j}(s) \left(p_t^{T*}(s) - \tau_p \frac{w_{t+j}(s)}{z_{t+j}(s)} \right) \left(\frac{p_t^{T*}(s)}{p_{t+j}^T(s)} \right)^{-\vartheta} y_{t+j}^T(s), \quad (33)$$

where $\mu_{t+j}(s)$ 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 price of traded goods is

$$p_t^T(s) = [(1 - \lambda_p)p_t^{T*}(s)^{1-\vartheta} + \lambda_p p_{t-1}^T(s)^{1-\vartheta}]^{\frac{1}{1-\vartheta}}. \quad (34)$$

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

$$n_t(s) = \left(\int_0^1 n_t(\iota, s)^{\frac{\psi-1}{\psi}} d\iota \right)^{\frac{\psi}{\psi-1}}, \quad (35)$$

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

$$\max_{w_t^*(s)} \sum_{j=0}^{\infty} \left(\lambda_w^j \prod_{i=0}^{j-1} \beta_{t+i}(s) \right) \left[\tau_w \mu_{t+j} w_t^*(s) \left(\frac{w_t^*(s)}{w_{t+j}(s)} \right)^{-\psi} n_{t+j}(s) - \frac{\eta_{t+j}^n(s)}{1+\nu} \left(\left(\frac{w_t^*(s)}{w_{t+j}(s)} \right)^{-\psi} n_{t+j}(s) \right)^{1+\nu} \right], \quad (36)$$

where λ_w is the probability that a given union leaves its wage unchanged and $\tau_w = \frac{\psi-1}{\psi}$ 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

$$w_t(s) = [(1 - \lambda_w)w_t^*(s)^{1-\psi} + \lambda_w w_{t-1}(s)^{1-\psi}]^{\frac{1}{1-\psi}}. \quad (37)$$

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

$$n_t(s) = \int_0^1 \left(\frac{p_t^N(\iota, s)}{p_t^N(s)} \right)^{-\vartheta} y_t^N(s) d\iota + \int_0^1 \left(\frac{p_t^T(\iota, s)}{p_t^T(s)} \right)^{-\vartheta} y_t^T(s) d\iota. \quad (38)$$

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

$$\underline{c}_t(s) = \frac{1}{\beta \mathbb{E}_t \mu_{t+1}(s) R_{t+1}} \frac{1}{p_t(s)}, \quad (39)$$

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:

$$q_t(a_{t+1}(s) - b_{t+1}(s)) = (1 + \gamma q_t)(a_t(s) - b_t(s)) + w_t(s)n_t(s) + T_t(s) - p_t(s)c_t(s). \quad (40)$$

3.2 Monetary Policy

Let $y_t = \int_0^1 p_t(s)y_t(s)/p_t ds$ 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

$$1 + i_t = \mathbb{E}_t R_{t+1} \quad (41)$$

be the expected nominal return on the long-term security, which we refer to as the nominal interest rate. Aggregation over the pricing choices of individual producers implies, up to a first-order approximation,

$$\log(\pi_t/\bar{\pi}) = \bar{\beta}\mathbb{E}_t \log(\pi_{t+1}/\bar{\pi}) + \frac{(1 - \lambda_p)(1 - \lambda_p\bar{\beta})}{\lambda_p} (\log(w_t) - \log(z_t)) + \theta_t,$$

where we add an AR(1) disturbance θ_t to individual firms' desired markups, $\bar{\beta}$ is the steady state discount factor and $\bar{\pi}$ is the steady-state level of inflation.¹³

We assume that monetary policy is characterized by a Taylor rule when the ZLB does not bind:

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

where ε_t^r is a monetary policy shock, α_r determines the persistence and 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 $\bar{\pi}$. When the ZLB binds, then

$$i_t = 0.$$

The interest rate may be at zero either because aggregate shocks cause the ZLB to bind, or because the Fed commits to keeping i_t at 0 for a longer time period than implied by the constraint. We thus implicitly assume that the Fed can manipulate expectations of how the path of interest rates evolves, as in Eggertsson and Woodford (2003) and Werning (2015). In our estimation we use survey data from the New York Federal Reserve to discipline the expected duration of the zero interest rate regime during the 2009 to 2015 period.

Since we assume that an individual island is of measure zero, monetary policy does not react to island-specific disturbances. The monetary union is closed so aggregate savings must equal to aggregate debt:

$$\int_0^1 a_{t+1}(s)ds = \int_0^1 b_{t+1}(s)ds. \quad (42)$$

¹³We 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.

3.3 Solution Method

Since we use Bayesian techniques to estimate the model, we need a fast and efficient solution method. We thus use a piece-wise linear solution method in order to solve the model. The first step in our algorithm is to linearize all equations of the model in the absence of a ZLB on the interest rate.

We accommodate the ZLB using an algorithm developed by Eggertsson and Woodford (2003), Guerrieri and Iacoviello (2015), and Jones (2017). Jones (2017) contrasts this method with fully non-linear ones in smaller scale models and shows that the piece-wise linear method we use here is accurate. Lepetyuk, Maliar and Maliar (2017) contrast these perturbation-based methods with fully non-linear projection methods in larger-scale models and reach similar conclusions regarding their accuracy.

We 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 all aggregate shocks 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

$$\mathbf{x}_t = \mathbf{J}_t + \mathbf{Q}_t \mathbf{x}_{t-1} + \mathbf{G}_t \epsilon_t, \quad (43)$$

where $\mathbf{x}_t$ collects the endogenous variables, ϵ_t collects the shocks, and the matrices $\mathbf{J}$, $\mathbf{Q}$ and $\mathbf{G}$ are varying over time due to the ZLB constraint.¹⁴

3.4 The Workings of the Model

We next provide some intuition for how this richer model works by showing impulse responses to state-level and aggregate credit shocks for a range of parameter values.

Impulse Response to a State-Level LTV Shock Figure 5 reports the impulse response to a state-level shock to the LTV constraint in the baseline parameterization using the modal posterior estimates discussed below. Panel A shows that the shock leads to a gradual decline in

¹⁴See Kulish and Pagan (2016) and Jones (2017) for more details on the recursion underlying this solution.

that island's debt because of the long-term nature of debt contracts. The island's asset position falls as well, yet not as much as debt does, so the island runs a current account surplus, reducing consumption of both tradable and non-tradable goods. Since prices are sticky, the decline in non-tradable consumption leads to a decline in employment owing to price and wage rigidities.

Figure 6 shows the responses to the same state-level LTV shock in an economy with a lower degree of idiosyncratic uncertainty. Notice that the responses of wages, employment and consumption are all dampened compared to those under the baseline estimates. Mechanically, this dampening is due to the much sharper decline in the asset position on the island, shown in Panel A, which now drops almost as much as debt does. As a result, the island no longer runs a large current account surplus and consumption and employment barely change.

The intuition for why assets fall much more when idiosyncratic uncertainty is low is as follows. Households can respond to a tightening of credit limits in two ways: either by reducing their assets or by cutting consumption. When idiosyncratic uncertainty is low, it is relatively costless to reduce asset holdings because the liquidity constraints are not very severe. Both sides of the households' balance sheets thus contract, with little impact on other variables. Credit in this case is a veil that has little consequence for macroeconomic aggregates. In contrast, when idiosyncratic uncertainty is high, reducing asset holdings is costly because individual members are more likely to be liquidity constrained. Households thus find it optimal to respond to the credit tightening by cutting consumption a lot more.

Intuitively, a tightening of debt constraints distorts allocations in two ways: it prevents households from smoothing the marginal utility of consumption both *across members* as well as *across time*. Households thus face a tradeoff: they can respond to a tightening of credit limits by either reducing overall consumption, thus worsening the intertemporal allocations, or by reducing asset holdings, thus worsening the allocation of consumption across agents. The more dispersed the idiosyncratic shocks are, the more the household chooses to reduce the overall level of consumption to avoid the high costs of variation of the marginal utility of consumption across its members.

Figure 7 illustrates the role of price rigidities in determining the state-level responses to a credit shock. We consider two economies, one in which prices and wages are sticky, and another in which they are flexible. When prices are flexible, wages respond immediately to the credit tightening owing to wealth effects on labor supply. Employment thus increases and consumption falls relatively little. In contrast, when prices are sticky, employment falls since wage rigidities

effectively act like a tax on labor supply and price rigidities lead to a rise in firm markups.¹⁵ The decline in employment amplifies the drop in consumption as well.

Impulse Response to an Aggregate Credit Shock To understand the effect of an aggregate credit shock in our model, recall from Equation (23) that changes in household credit lead to changes in the natural interest rate. The degree to which a shock to credit affects real variables therefore depends crucially on the extent to which the monetary authority accommodates these changes in the natural rate.

We illustrate this in Figure 8 where we contrast the effect of a large and gradual credit tightening with and without a binding ZLB. On impact, the credit shock causes the ZLB to bind for five years and employment to fall by 6.5 percent. Absent the ZLB, the Fed would gradually reduce the nominal rate below zero and employment would fall very little.

A broader implication of these impulse responses is that we cannot simply extrapolate state-level correlations to inform about the role of credit shocks in generating fluctuations in the aggregate. Even though credit shocks may have large effects on state-level employment, owing to the presence of price rigidities, price rigidities alone are not sufficient for credit shocks to affect employment in the aggregate – monetary policy must be constrained as well. Thus, while the state-level evidence is important for identifying the key parameters of the model, one needs to use the structure of the model to isolate the role of credit shocks in generating aggregate fluctuations in real variables.

4 Estimation

We next describe how we have chosen parameters for our model. We first discuss the parameters we assign values to based on the literature or steady-state considerations, and then the ones we estimate using Bayesian likelihood methods and state-level and aggregate data.

The parameters that we estimate are those that are critical in determining the model's responses to a credit shock: the degree of wage and price stickiness, λ_w and λ_p , as well as the volatility of idiosyncratic taste shocks, α . We refer to these parameters as our *structural parameters*. In addition, we estimate the AR(1) processes characterizing the island-specific and aggregate components of the various shocks. We next describe the parameters we have assigned, the construction of the likelihood function, and then our results.

¹⁵See 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.

4.1 Assigned Parameters

We report the parameter values we assign in Table 1. The period is one quarter. We assume a Frisch elasticity of labor supply of $1/2$ and thus set ν equal to 2. We set γ , the parameter governing the duration of debt, to 0.985, so that the Macaulay duration of debt in our model is equal to that of mortgage debt in the data, approximately 13 years.¹⁶ We follow the trade literature in setting the weight on non-traded goods in an island's consumption basket, ω , equal to 0.7, the elasticity of substitution between traded and non-traded goods, σ , equal to 0.5, and the elasticity of substitution between varieties of traded goods produced in different islands, κ , equal to 4. We follow Christiano, Eichenbaum and Evans (2005) in choosing ψ to ensure a wage markup of 5%.¹⁷ We use the Justiniano and Primiceri (2008) estimates of the parameters characterizing the Taylor rule.

We pin down three additional parameters using steady state considerations. The steady state discount factor $\bar{\beta}$ is chosen so that the steady state real interest rate is equal to 2% per year. The steady state weight of housing in preferences $\bar{\eta}^h$ is chosen so that the aggregate housing to (annual) income ratio is equal to 2.5, a number that we compute using the 2001 Survey of Consumer Finances (SCF). Finally, the steady state LTV 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 (annual) income ratio of $2.5 \times 0.29 = 0.725$.

4.2 Estimation Procedure

Our approach is motivated by the work of Mian and Sufi (2014) and Beraja, Hurst and Ospina (2015) who illustrate the value of using regional variation in identifying the effect of credit shocks on real activity. Since our goal is to study both the model's regional and aggregate implications, we depart from these papers by combining regional as well as aggregate data in the estimation. Since we study a model with occasionally binding constraints (the ZLB), evaluating the likelihood function using conventional methods is computationally infeasible due to the very large number of state variables and the curse of dimensionality. We therefore propose a novel methodology that exploits the structure of the model and allows us to combine state-level and aggregate data to form the likelihood.

¹⁶The Macaulay duration is the weighted average maturity of the flows, with weights given by the present value of the flows accruing at each date. In our model with geometrically decaying perpetuities this duration is given by $(1+r)/(1+r-\gamma)$.

¹⁷We show in our Robustness section that our conclusions are not sensitive to alternative values of this parameter.

We first describe the details of our approach, and then report our results.

4.2.1 The Parameters

We estimate the three structural parameters, α , λ_p , λ_w , as well as the persistence and volatility of the state and aggregate components of the shocks to productivity, the discount factor, preference for housing, disutility from work, and the credit limit.

4.2.2 The Data

Regional Data We use a panel of employment, household spending, wages, household debt and house prices in the cross-section of 51 U.S. states from 1999 to 2015. The household debt information is from the FRB New York Consumer Credit Panel.¹⁸ Since the measure of debt in the data is the book value of outstanding debt, we construct a corresponding series for our model using

$$\text{debt}_t(s) = \gamma \text{debt}_{t-1}(s) + q_t l_t(s),$$

where $q_t l_t(s)$ is the market value of mortgage debt issued in period t on island s and $\gamma \text{debt}_{t-1}(s)$ is the book value of outstanding debt. We found that our results are largely unchanged when we alternatively use the market value of debt $q_t b_{t+1}(s)$ to proxy for household credit in the data.

We use data on house prices from the FHFA, and data on employment, consumption expenditures and wages from the BEA. Our measure of employment is the employment-population ratio in a given state. Our measure of wages is total employee compensation divided by the number of workers. Since we do not model the construction sector, we subtract construction employment and compensation in computing measures of employment and wages.¹⁹

Aggregate Data We use aggregate data on employment, household consumption, wages, household debt, house prices, inflation and the Fed Funds rate from 1984 to 2015. We construct this data in a similar way to the state-level data. An additional critical input in the estimation is the sequence of expected durations of the ZLB between 2009 and 2015, which we take from the Blue Chip Financial Forecasts survey from 2009 to 2010 and the New York Federal Reserve's Survey of Primary Dealers from 2011 to 2015.²⁰

¹⁸We include credit card debt, auto loans and student loans, in addition to mortgage debt in our measure of household credit. Since mortgage debt makes up a large fraction of household debt, our results are virtually unchanged when we only use mortgage debt in our estimation.

¹⁹See our Robustness section for an extension that explicitly models the construction sector and uses data on construction employment in estimation. Our results are robust to this modification.

²⁰See the Appendix for a more detailed description of the data we use.

4.2.3 The Likelihood Function

Consider first the state-space representation of the piece-wise linear approximation of our model discussed in Section 3.3. Recall that 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 variables, both state and aggregate and ϵ_t collects the shocks, both state and aggregate.

The conventional approach to estimating the model would be to write down a likelihood function that directly uses state and aggregate data and the above solution. This approach is computationally infeasible in part because of the non-linearity induced by the ZLB, captured by the time-varying reduced-form matrices, $\mathbf{J}_t$, $\mathbf{Q}_t$, $\mathbf{G}_t$, which change depending on how long the ZLB is expected to last. This non-linearity, combined with the curse of dimensionality which arises from our use of 51 regions, each of which has 10 individual state variables (the island's savings, debt positions as well as the past wage and price composites, together with the 5 shocks) renders the conventional approach infeasible.

We thus formulate an alternative approach to constructing the likelihood function, one that exploits *relative* variation across individual states' outcomes during the 1999 to 2015 period. Intuitively, our approach recognizes that, up to a first-order approximation, the difference between employment in, say, Nevada and in the aggregate is a linear function of the Nevada state variables only. This observation allows us to separate the likelihood into state-level components and an aggregate component.

Formally, let $\mathbf{x}_t^j$ denote the vector of variables for each island j , expressed in log-deviations from the steady state. Given our piece-wise linear approximation, we can write:

$$\mathbf{x}_t^j = \underbrace{\mathbf{Q}^s \mathbf{x}_{t-1}^j + \mathbf{G}^s \epsilon_t^j}_{\text{state-level component}} + \underbrace{\mathbf{J}_t^a + \mathbf{Q}_t^a \mathbf{x}_{t-1}^* + \mathbf{G}_t^a \epsilon_t^*}_{\text{aggregate component}}. \quad (44)$$

Here, $\mathbf{Q}^s$ and $\mathbf{G}^s$ encodes how island j 's variables depend on its own state variables and island-specific shocks ϵ_t^j , while the vector $\mathbf{x}_t^*$ collects the aggregate variables and evolves according to:

$$\mathbf{x}_t^* = \mathbf{J}_t^* + \mathbf{Q}_t^* \mathbf{x}_{t-1}^* + \mathbf{G}_t^* \epsilon_t^*. \quad (45)$$

Here, ϵ_t^* are the aggregate shocks. Notice that the piecewise-linear approximation gives our problem a very specific structure, whereby the ZLB affects individual islands through a time-varying matrix of coefficients $\mathbf{Q}_t^a$ and $\mathbf{G}_t^a$ multiplying aggregate variables only. Since these

matrices depend on the underlying structural parameters in a non-linear way, computing them for each parameter draw in the estimation is too time-consuming given the dimensionality of the model.

In contrast, the matrix of coefficients $\mathbf{Q}^s$ and $\mathbf{G}^s$ multiplying the island-level variables is time-invariant. Intuitively, since each island is of measure zero, shocks specific to an island do not change the expected date at which the ZLB will stop binding for the rest of the economy. Moreover, agents on each island take aggregate prices, including the interest rate, as given. Therefore, conditional on the aggregate state variables, whether or not the ZLB binds does not affect the evolution of island-level relative variables. From the perspective of agents in a given island, the presence of the ZLB acts like any other aggregate shock which, up to a first-order approximation, does not change how that island responds to its own history of idiosyncratic shocks.

Given this structure our model, letting $\mathbf{x}_t = \int \mathbf{x}_t^j dj$ denote the economy-wide average of the island-level variables, the deviation of island-level variables from their economy-wide averages,

$$\hat{\mathbf{x}}_t^j = \mathbf{x}_t^j - \mathbf{x}_t, \quad (46)$$

can be written as a time-invariant function of island-level variables alone:

$$\hat{\mathbf{x}}_t^j = \mathbf{Q}^s \hat{\mathbf{x}}_{t-1}^j + \mathbf{G}^s \epsilon_t^j, \quad (47)$$

where we use the assumption $\int \epsilon_t^j dj = 0$, that island-level shocks have zero mean in the aggregate.

We therefore use the representation in (45) and (47), rather than the computationally infeasible representation in (44) to estimate the model using state-level and aggregate U.S. data. To do so, we first express each state's observable variable as deviations from its aggregate counterpart by subtracting a full set of time effects, one for each year and each variable. We also subtract a state-specific fixed effect and time trend for each observable since in our model all islands are ex-ante identical.²¹ Since the island-level shocks in (47) are independent and do not affect aggregate outcomes, we can write the likelihood of the model as the product of each individual state's likelihood, computed from (47) and the aggregate likelihood, computed from (45). To summarize, our estimation exploits the differential rise and fall of individual states' spending, debt, wages, house prices and employment, in addition to the aggregate comovement of these series, to identify the structural parameters of the model.

²¹In practice, subtracting the state-specific fixed effects only, as we have done in an earlier version of the paper, does not change our substantive results.

We use a Kalman filter to compute forecast errors that enter the calculation of the likelihood,²² and a Metropolis-Hastings algorithm to sample from the posterior distribution, which, by Bayes' rule, is proportional to the product of the prior and the likelihood. We initialize the Kalman filter for each state and for each proposal of the parameters at the steady-state implied by those parameters. The state-level data is observed at an annual frequency, while the model is quarterly, so we conduct a *mixed-frequency* estimation, by computing forecast errors every four quarters for state-level variables. The aggregate data, in contrast, is observed quarterly and for a longer time-period, from 1984 to 2015. To account for differences in the size of different states and ensure that smaller states exert a relatively smaller influence on the shape of the likelihood, we weight the likelihood contribution of each state by their 1999 population shares.²³ We assign an equal weight to the aggregate likelihood and the combined state-level likelihood.

4.3 Parameter Estimates

Table 2 reports moments of the prior and posterior distributions of the structural parameters we estimate. We used diffuse priors. The modal estimate of the Pareto tail parameter α is equal to 3.68, implying a 0.3% spread between the subjective discount rate and interest rate at the annual frequency. The posterior distribution for α is relatively tight around its mode, with a 10th and 90th percentile of 3.2 and 4.1, respectively. As we discussed earlier, this parameter is largely identified from the comovement of state-level household debt, consumption and employment which, as Figure 1 shows, are highly correlated in the data. When α is too high, fluctuations in credit have no consequences for real variables since an island's asset holdings adjust one-for-one with credit, leaving its net asset position unchanged. In this case the model lacks an endogenous mechanism to replicate the observed correlation between credit and real variables in the data and the value of the likelihood is low. Similarly, when α is too low, the model implies too strong a relationship between household debt and consumption and employment, which also reduces the value of the likelihood.

We illustrate this point in Table 3 where we simulate a panel of 51 states using artificial data generated by our baseline estimates and calculate the correlation between changes in employment and household debt and changes in consumption and household debt. Notice that in the data these correlations are equal to 0.17 and 0.23, respectively, numbers which our baseline model reproduces reasonably well. In contrast, when we impose $\alpha = 2$, this correlation

²²Details of how we construct the likelihood function are discussed in the Appendix.

²³See Agostinelli and Greco (2012) who show that this simple weighting scheme preserves the asymptotic properties of the true likelihood function.

increases to 0.58 and 0.7, respectively. Similarly, when we simulate artificial data from a model with $\alpha = 10$, these correlations fall to zero.

We find, consistent with the work of Del Negro, Giannone and Schorfheide (2015), that wages and prices are sticky, with a modal estimate of λ_w of 0.85 and a modal estimate of λ_p of 0.97. As is well known, accounting for the stability of inflation around the Great Recession requires a great deal of price stickiness. We show below that estimating the model using state-level data alone implies a much lower degree of price stickiness, consistent with the evidence in Beraja, Hurst and Ospina (2015). We explore the implications of this alternative set of parameter estimates in our Robustness section below.

We report in the Appendix the posterior estimates of the parameters governing the persistence and volatility of state and aggregate shocks, as well as theoretical forecast error variance decompositions of state and aggregate observables into the various shocks at different horizons. These theoretical decompositions allow us to interpret the estimates of the persistence and volatility of shocks. In a short sample, however, these decompositions do not necessarily reveal which particular shock explains movements in the actual variables in the data. Because of this, we next use the filtered shocks to calculate the contribution of each shock to the variance of observed consumption and employment.

Consider first Panel A of Table 4 which reports this variance decomposition for state-level variables. Recall that these are expressed in deviations from their aggregate counterparts. State-level credit shocks account for about 10% of the variation of relative employment and 18% of the variation in relative household spending. State-level productivity shocks account for about 29% of the variation in relative employment and a negligible fraction of the volatility of consumption. Shocks to the disutility from work account for about 14% of the volatility of employment and 3% of the volatility of consumption. Finally, shocks to the individual states' discount rates have the largest impact and account for about 50% percent of the volatility of relative employment and 80% of the volatility of relative consumption. Shocks to the preference for housing have a negligible impact on state-level real variables.

Consider next Panel B of Table 4 which reports a similar decomposition for aggregate variables prior to the ZLB period, from 1984 to 2008. Due to non-linearities, this decomposition is not as straightforward to interpret during the years in which the ZLB binds, so we discuss it in more detail in the next section. In the pre-ZLB period household credit shocks account for only 0.2% of the volatility of aggregate consumption and employment. Intuitively, credit shocks matter much more at the state-level than in the aggregate because individual states are

part of a monetary union and cannot use monetary policy to offset their effects. Productivity shocks account for about 30% of the employment volatility and 1 percent of the volatility of consumption. Discount rate, monetary policy and markup shocks each account for a substantial fraction of movements in consumption and employment, while shocks to the disutility from work or preference for housing have little effect on these real variables.

4.4 Additional Implications

We finally study our model's ability to reproduce several additional variables that we have not directly used in estimation. Consider first Figure 9 which shows how in our model state-level tradable and nontradable employment comove with consumption during the Great Recession. As Mian and Sufi report, most of the decline in employment at the state level was due to a decline in nontradable employment. Our model reproduces this fact well: the elasticity of nontradable employment to consumption is equal to 0.75 in the model (0.55 in the data, as reported by Kehoe, Midrigan and Pastorino, 2017). Similarly, tradable employment comoves little with state-level consumption, both in the model (the elasticity of tradable employment to consumption is -0.22) and in the data (an elasticity of -0.03).

Consider next the model's implications for the mortgage interest rate. Since in our model mortgages are long-term perpetuities with decaying coupon payments, the return on such securities is not directly comparable to the interest rate on 30-year mortgages in the data. Nevertheless, we can derive the implied rate at which the flows underlying these securities are discounted as the rate i_t^m that rationalizes the price of the security q_t . This rate is defined by

$$q_t = \frac{1}{1 + i_t^m} \sum_{k=0}^{\infty} \left(\frac{\gamma}{1 + i_t^m} \right)^k,$$

which gives

$$1 + i_t^m = \frac{1}{q_t} + \gamma.$$

Figure 10 compares this implied long-term rate in our model with the average interest rate on 30-year mortgages in the data. Since the latter series has a trend, we detrend both the model and the data. The model does a reasonable job at reproducing medium-term movements in the mortgage rate in the data, though it misses the high-frequency variation. Since the model abstracts from several sources of risk embedded in mortgages rates, such as default and prepayment risk, we do not view the model's failure to match these high frequency fluctuations as critical. Indeed, since our Kalman filter isolates the LTV shocks from the dynamics of mortgage

debt in the data, changes in mortgage spreads in the data are captured in a reduced-form way as shifts in the LTV constraint.

5 Credit Shocks in the Great Recession

We next study the role of household credit shocks in shaping the dynamics of employment and consumption in the cross-section and in the aggregate during the Great Recession. Recall that in our model the borrowing limit always binds, so household credit fluctuates due to exogenous changes in the LTV parameter $m_t(s)$, as well as due to changes in house prices, $e_t(s)$, which are mainly driven by shocks to the preference for housing. Since housing is in fixed supply and separable in preferences, housing preference shocks only affect equilibrium outcomes through their effect on the credit limit. We thus refer to both shocks to the LTV ratio $m_t(s)$ and housing preferences, $\eta_t^h(s)$ as *credit shocks*, though our estimates imply that the bulk of the movements in real variables are due to the changes in the LTV ratio $m_t(s)$, as the variance decompositions in Table 4 demonstrate.

We first show that state-level credit shocks account for a large fraction of the differential rise and fall in state-level employment and consumption during the boom and the bust years. Our findings thus reinforce the conclusions of Mian and Sufi (2013, 2016) that household credit played an important role in accounting for the heterogeneity in state-level outcomes during the Great Recession. We then study the aggregate implications of credit shocks. We show that aggregate credit shocks generate a modest decline in the natural rate of interest during the Great Recession and are not sufficient to trigger the ZLB on their own. Absent the ZLB, therefore, aggregate credit shocks have a relatively small effect on employment since monetary policy effectively offsets them. We finally study how the effect of credit shocks on employment are magnified when monetary policy is constrained, as was the case during the Great Recession.

5.1 Role of Credit Shocks at the State Level

We start by studying the 2002 to 2007 period. In Figure 11 we show how an individual state's employment and consumption, expressed in deviations from their aggregate counterparts, changed from 2002 to 2007. The x-axes in these figures show the actual data we used to extract the state-level component of the shocks. We construct a counterfactual series for each variable by shutting down all of the shocks other than those to credit, and plot the 2002 to 2007 changes in these series on the y-axes. The solid line in each of these plots is the 45 degree line.

The upper panels of Figure 11 show that credit shocks account for a substantial fraction of

the differential changes in consumption and employment during the 2002 to 2007 period. The correlation between the change in the counterfactual series driven by credit shocks only and the change in the data is 0.7 for employment and 0.6 for consumption. The slope coefficient of a regression of these counterfactual changes against the data is 0.43 for employment and 0.25 for consumption. Credit shocks thus account for about 40% of the differential rise in employment and 25% of the differential rise in consumption across states during the boom.

The lower panels of Figure 11 illustrates the corresponding patterns during the 2007 to 2010 recession. Credit shocks account for an even larger fraction of the differential change in consumption and employment during the recession. The correlations between the counterfactuals and the actual data are 0.7 for employment and 0.75 for consumption, while the slope coefficients are 0.48 for employment and 0.43 for consumption. Thus, the model implies that credit shocks alone account for about 40% of the differential decline in consumption and employment across individual states.

We emphasize that these results are not simply an artefact of the modeling choices we have made, but rather are a result of the estimation and the path for shocks extracted by the Kalman filter from the data. As we report in our robustness section below, if instead the posterior estimates were centered around a higher value of α , credit shocks would have had much smaller effects and the scatterplots shown in Figure 11 would simply collapse to 0. Similarly, if we were to estimate a lower degree of wage and price stickiness, state-level credit shocks account for a smaller fraction of the differential rise and fall in employment and consumption across states.

5.2 Role of Credit Shocks at the Aggregate Level

Figure 12 presents the dynamics of the key aggregate time series during the 1995 to 2015 period in the data. We HP-filter the debt-to-income series and report it as deviations from the trend in Panel A. Panel B reports the evolution of the Fed Funds Rate and the implied natural rate predicted by the model, Panel C reports the data on employment, measured as the total number of employees on non-farm payrolls, scaled by the U.S. population and expressed as percent deviations from a linear trend, and Panel D shows the time series for inflation. We used these series, together with aggregate data on wages, house prices and consumption, to extract the aggregate component of the exogenous shocks used in estimation.

We construct a counterfactual series for each of these variables by shutting down all shocks other than the two credit shocks. Panel A of Figure 12 shows that credit shocks alone drive much of the movements in household credit. Panel B shows, however, that credit shocks alone

generate much more gradual changes in the natural rate of interest. All shocks combined give rise to an increase in the natural rate of interest from 2.2% in 2002 to 4.3% in 2007 and then a fall to -1% by the end of 2009. In contrast, in an economy with credit shocks only the natural rate would have increased from 2.1% in 2002 to 2.3% in 2007 and then would have fallen to only 1.8% by the end of 2015.

The reason credit shocks generate much more gradual movements in the natural rate is that household debt evolved gradually in the data. Since household debt and the natural rate of interest are closely related, as equation (23) shows, the gradual increase and then decline in household debt was associated with a corresponding gradual increase and then decrease in the natural interest rate. Moreover, our estimate of α implies an elasticity of the natural rate to changes in debt that is relatively low. Thus, even though our model implies that credit shocks account for a large fraction of state-level variation in consumption and employment, they account for a much smaller fraction of the movements in the natural rate in the aggregate.

Consider next how employment responds to credit shocks. As the lower-left panel of Figure 12 shows, the Great Recession was associated with a nearly 7% drop in the employment-population ratio from 2007 to 2010. Credit shocks alone account for a much more modest 0.3% drop in this time period, less than 5% of the actual decline. The reason employment falls so little in response to credit shocks alone is that the decline in the natural rate of interest in response to these shocks is not sufficiently large to trigger the ZLB. Absent additional shocks monetary policy is not constrained and mimics the dynamics of the natural rate quite well, implying modest movements in employment and inflation.

Our findings thus suggest that absent additional shocks around the 2007 to 2008 period, the household leverage cycle observed in the U.S. would have had little effect on aggregate employment. This result reinforces the point that one cannot simply extrapolate the state-level estimates of how credit and employment comove to make inferences about the aggregate. Even though credit shocks account for a substantial fraction of employment fluctuations in individual states, they explain little of the aggregate employment decline. The decline in the natural rate of interest predicted by our model is fairly persistent, however. Thus, if household credit shocks are accompanied, as they were in the data, 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. How much greater?

Answering this question is challenging because of the non-linearities arising from the ZLB and the Fed's pursuit of *forward guidance* policies aimed at alleviating its consequences. In our

estimation we have taken the Fed's forward guidance announcements as given, by setting the expected ZLB duration each quarter equal to that observed in the data. Since these announcements were responding to the path of all shocks, including credit shocks, we cannot turn off the effect of credit shocks alone without taking a stand on how forward guidance would have been conducted absent such shocks. In short, the presence of the ZLB presents an additional identification challenge, one of decomposing the expected ZLB durations into an endogenous component due to shocks, and one due to forward guidance policies.²⁴

We isolate the contribution of credit shocks when monetary policy is constrained by considering a counterfactual experiment in which the Fed *does not* pursue forward guidance. In this case the expected duration of the ZLB is solely determined by the size and the persistence of the aggregate shocks, as well as the structure of the model, including the parameters of the Taylor rule that govern the dynamics of interest rates upon exit from the ZLB regime. The dotted line in Panel A of Figure 13 shows what the response of employment would have been in the absence of forward guidance. Clearly, the decline in employment would have been much worse than those observed in the data.

Panel A also reports what employment would have been absent the credit shocks from the middle of 2007 onward in this environment without forward guidance. The difference between these two counterfactual series allows us to gauge the contribution of credit shocks alone at the ZLB holding forward guidance constant (at zero). Panel B shows the difference between these two series and compares it to the employment decline in the data. Interestingly, the contribution of credit shocks remains small even at the zero lower bound: they generate a 0.5% drop in employment relative to trend in 2009, and a 0.8% decline in employment relative to trend in 2010. Thus, credit shocks alone account for a small fraction of the observed employment declines during the onset of the Great Recession. This result once again reflects the gradual decline in household debt observed in the data, as well as our parameter estimates.

Notice, however, that the contribution of credit shocks to the employment decline gradually increases over time as household leverage continues to fall. Credit shocks alone generate a much more sizable 2.5% drop in employment by 2015, about half of the employment gap in the data. We thus conclude that credit shocks did play a substantial role in delaying the recovery in the aftermath of the recession, even though they had a modest impact in the years immediately preceding 2007.

²⁴See Jones (2017) for a more detailed discussion of this identification problem and a solution to it.

6 Robustness

We conclude by gauging the robustness of our findings to alternative parameterizations of the degree of wage and price stickiness and the duration of mortgage contracts.

We considered several robustness checks.²⁵ First, we estimated the structural parameters of the model using regional data alone, in the spirit of Beraja, Hurst and Ospina (2015). Second, we imposed a much higher value of α , the parameter governing the degree of idiosyncratic uncertainty, and re-estimated all other parameters as earlier. Third, we have assumed a lower duration of mortgage contracts, as captured by the parameter γ . Fourth, we have assumed a lower elasticity of substitution ψ between varieties of labor supplied by different unions. Finally, we introduce a construction sector and add construction employment as an additional observable in estimation.

6.1 Using Regional Data Alone

Here we estimate the model's structural parameters, α , λ_w and λ_p using regional data alone. We then fix these parameters and use aggregate data to only estimate the parameters of the aggregate shock processes. Table 5 shows that the posterior distribution of the parameters estimated using regional data alone implies a similar degree of idiosyncratic uncertainty, $\alpha = 3.6$, a slightly lower degree of price stickiness, $\lambda_p = 0.93$, compared to 0.97 in our baseline and a much lower degree of wage stickiness, $\lambda_w = 0.58$, compared to 0.85 in our baseline. Intuitively, as Beraja, Hurst and Ospina (2015) point out, state-level wages are quite volatile. In contrast, aggregate inflation has changed little during the Great Recession. Reproducing the aggregate data thus requires a greater degree of nominal stickiness, while reproducing the regional data requires a lower degree of nominal stickiness. Our baseline estimates that use both aggregate and regional data thus fall somewhere in between. Since we allow aggregate and regional shocks to differ in their persistence and volatility, our results suggest that the puzzle is not explained by differences in shocks. Instead, our results are consistent with the idea that prices respond more to large shocks than to small shocks. This is consistent with menu cost models (e.g., Alvarez, Le Bihan and Lippi, 2016, Midrigan and Kehoe, 2014) and with models of rational inattention (Mackowiak and Wiederholt, 2009). We emphasize regional differences, but similar evidence exists for sectoral differences (Boivin, Giannoni, and Mihov, 2009).

Figure 14 shows that the model estimated using regional data alone implies a smaller role for credit shocks in generating the differential rise and fall in employment across U.S. states during

²⁵We report the main checks here, with further robustness exercises left to the Appendix.

the boom and the recession. Although the counterfactual series for employment constructed with credit shocks alone is quite strongly correlated with the data (0.69 during the boom, and 0.72 during the recession), the slope coefficient is much smaller (0.26 during the boom and 0.40 during the recession), implying that credit shocks alone account for about 25–40% of the volatility of relative changes in employment in the data.

Panel A of Figure 15 shows that credit shocks alone imply a similar decline in employment as our baseline model when the economy is at the ZLB and absent forward guidance. As earlier, the model predicts a minor role for credit shocks at the onset of the Great Recession, but these shocks have a substantial role in delaying the recovery. Intuitively, the degree of price stickiness has an ambiguous impact on how employment reacts to credit shocks at the ZLB. More flexible prices may, in fact, exacerbate the response of employment to credit shocks by implying a larger deflation and a sharper increase in the real interest rate.

6.2 Lower Degree of Idiosyncratic Uncertainty

Matching the comovement between household credit and consumption in the data requires a relatively high degree of idiosyncratic uncertainty. Here we impose a much lower degree of idiosyncratic uncertainty, by changing $\alpha = 5$ compared to 3.68 in our baseline, and re-estimating all other parameters of the model.

As Panel B of Figure 14 shows, credit shocks now produce almost no relative movements in employment across states. As discussed earlier, absent a large degree of idiosyncratic uncertainty, agents in a state subject to a credit tightening can simply adjust the asset side of their balance sheet, with little consequence for consumption and employment. Similarly, as Panel B of Figure 15 shows, credit shocks alone generate virtually no employment drop in the aggregate even at the ZLB. This experiment shows that our results regarding the role of credit shocks at the state-level and in the aggregate are not an artefact of our modeling choices, but rather come from the richness of the data used to estimate the model.

6.3 Lower Duration of Mortgage Contracts

So far we have imposed a value of γ , the parameter determining the decay rate of coupon payments in the mortgage contract, equal to 0.985, consistent with the maturity of mortgage contracts in the data. One could argue, however, that the effective duration of mortgages in the data is lower, due to households' ability to refinance their mortgages or take on home equity loans. Here we reduce γ to 0.965, implying a duration of mortgages about half that in our

baseline (6 years versus 13 years) and reestimate the model. Table 5 shows that our estimates of the degree of price and wage rigidity are similar to those in our baseline model, while the estimate of α is somewhat lower, 2.97 vs. 3.68. As Panel C of Figure 14 shows, credit shocks now explain almost as much of the variation in employment across regions as in our baseline estimation. Mechanically, the volatility and persistence of the shocks, especially the credit shocks, adjusts when we reduce γ , so that the model can reproduce the dynamics of household credit across regions, with little consequence for the behavior of the other variables. Panel C of Figure 15 shows that credit shocks now account for a slightly larger drop in aggregate employment during the recovery.

6.4 Lower Elasticity of Substitution Between Labor Varieties

In our baseline estimation, we assigned a value of this elasticity, ψ , equal to 21, following Christiano, Eichenbaum and Evans (2005). This parameter acts like a real rigidity, in that it prevents reset wages from responding too much to a given shock. Here we reduce this parameter to 5, and re-estimate all other parameters of the model. As Table 5 shows, the estimation now favors an even greater degree of nominal wage and price stickiness to compensate for the removal of the real rigidity. Since we now estimate a greater degree of idiosyncratic uncertainty (the modal estimate of α falls to 3.2 from its baseline value of 3.7), the model now assigns a greater role to credit shocks in accounting for the state-level variation in employment (Panel D of Figure 14) as well as in delaying the recovery in aggregate employment (Panel D of Figure 15). Nevertheless, the model continues to assign a small role to household credit shocks in generating the drop in aggregate employment at the onset of the recession.

6.5 Construction Sector

So far we have assumed that the housing stock is in fixed supply and for consistency have removed construction employment from state and aggregate data. Here we introduce a construction sector and add construction employment as an additional observable in our estimation. We assume that the housing stock on each island evolves according to

$$h_{t+1}(s) = (1 - \delta_h) h_t(s) + y_t^H(s),$$

where δ_h is the rate at which housing depreciates and $y_t^H(s)$ is new housing investment, produced using a technology that uses a fixed factor and construction employment, $n_t^H(s)$:

$$y_t^H(s) = z_t^H(s) (n_t^H(s))^\chi,$$

where $z_t^H(s)$ is the productivity of the construction sector on island s . The problem of a firm in the the construction sector is to maximize its profits, given by

$$e_t(s)z_t^H(s) \left(n_t^H(s)\right)^\chi - w_t(s)n_t^H(s) - w_t(s)\frac{\xi}{2} \left(n_t^H(s) - \bar{n}^H\right)^2,$$

where the last term is an adjustment cost that capture frictions that prevent movement of employment across sectors.

We set $\chi = 0.37$ as in Garriga and Hedlund (2017) and $\delta_h = 0.012$ to match the 4.9% share of construction employment in total employment. We estimate the process for $z_t^H(s)$ as well as the other parameters of the model using the original variables as well as state and aggregate data on construction employment. We report our estimates in Panel E of Table 5. Our estimates of the structural parameters are very similar to those in the baseline model. Figure 16 shows the model's implications for the role of credit shocks in determining relative movements in state-level non-construction employment, consumption and construction employment in the boom and the bust. Our estimates imply that credit shocks alone account for about 40% of state-level employment variation, 20-35% of consumption variation and about 70-80% of variation in construction employment. Construction employment is much more volatile than non-construction employment, but since it accounts for a small fraction of overall employment, movements in it do not greatly change our baseline model's predictions. As Panel E of Figure 15 shows, the model's implications for the dynamics of aggregate employment are unchanged.

6.6 Alternative Estimates of the Taylor Rule

In our baseline estimates we have used the parameters of the Taylor rule estimated by Justiniano and Primiceri (2008) using pre-Great Recession data. We have also estimated these policy parameters ourselves using a longer sample inclusive of the 2009 to 2015 period. As reported in the Appendix, our estimates of the Taylor Rule are similar to those of Justiniano and Primiceri (2008), so our model's implications are largely unchanged.

7 Conclusions

A popular account of the U.S. Great Recession is the view that declines in 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 household debt in the cross-section of U.S. regions. This paper asks: what are the *aggregate* implications of these regional correlations?

We develop a model that captures this view by introducing a role for credit in alleviating household liquidity constraints. The more uncertain households are about their individual liquidity needs, the stronger their precautionary savings motive and the more sensitive their consumption is to changes in credit limits. Aggregating across individual households implies that the natural rate of interest fluctuates more in response to changes in credit limits when idiosyncratic uncertainty is higher. We use both state-level and aggregate data on the comovement between consumption and household credit to identify the key parameter governing the strength of the precautionary savings motive and thus the sensitivity of macroeconomic variables to credit shocks.

We use Bayesian likelihood methods together with state-level and aggregate data, to estimate the parameters of the model and the processes of a rich set of shocks that drive regional and aggregate fluctuations. To make the estimation feasible, we propose a novel methodology that exploits the structure of the model and identifies parameters from the relative variation of state-level variables, as well as from the U.S. aggregates. Our approach allows us to efficiently calculate the likelihood function in an economy with a large number of aggregate and regional state variables and with non-linear dynamics caused by the ZLB constraint on monetary policy.

Our estimates imply that credit shocks account for about 40% of the differential rise and fall in *state-level* employment and consumption during the boom and bust years of the past decade. In contrast, credit shocks alone generate a modest 0.8% drop in *aggregate* employment from 2008 to 2010, about one-tenth of the 7% observed drop, even though monetary policy was constrained due to the ZLB. This result reflects the fact that household credit declined very gradually in the data, causing a gradual decline in the natural rate of interest.

The persistent nature of household deleveraging implies, however, that credit shocks became more important over time as the quantity of household debt declined and the natural rate fell even further. Our estimates suggest that credit shocks account for about one half of the employment gap in 2015. The tightening in household-level credit limits thus helps account for the slow recovery of employment in the aftermath of the Great Recession.

We chose to keep our analysis transparent by focusing on one mechanism, the ‘household leverage’ view of the recession. In doing so, we have abstracted from other forces, such as constraints on financial intermediaries or demographic trends, that played an important role during the last decade. We captured these additional forces in a reduced-form way with various other shocks. Though beyond the scope of this paper, explicitly modeling the interaction between household credit limits and these other factors is an important area for future research.

References

- [1] Alter, Adrian, Alan Xiaochen Feng and Nico Valckx. 2018 “Understanding the Macro-Financial Effects of Household Debt: A Global Perspective,” IMF Working Paper No 18/76.
- [2] Alvarez, Fernando, Herve Le Bihan and Francesco Lippi. 2016 “The real effects of monetary shocks in sticky price models: a sufficient statistic approach,” *American Economic Review*, 106(10), 1-37.
- [3] Anderson, Asger Lau, Charlotte Duus and Thais Laerkholm. 2014 “Household Debt and Consumption During the Financial Crisis: Evidence from Danish Micro Data,” Danmarks Nationalbank Working Papers, No 89.
- [4] Beraja, Martin, Erik Hurst and Juan Ospina. 2015. “The Aggregate Implications of Regional Business Cycles,” Chicago Booth Working Paper.
- [5] Bernanke, Ben and Mark Gertler. 1989 “Agency Costs, Net Worth and Business Fluctuations,” *American Economic Review*, 79, 14-31.
- [6] 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.
- [7] Blanco, Andres. 2015. “Optimal Inflation Target in an Economy with Menu Costs and an Occasionally Binding Zero Lower Bound,” University of Michigan Working Paper.
- [8] Boivin, Jean, Marc Giannoni, and Ilian Mihov. 2009. “Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data.” *American Economic Review*, 99(1): 350-84.
- [9] Burnside, Craig, Martin Eichenbaum and Sergio Rebelo. 2015. “Understanding Booms and Busts in Housing Prices.” Northwestern University Working Paper.
- [10] Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo. 2005. “Large Devaluations and the Real Exchange Rate,” *Journal of Political Economy*, 113(4), 742-784.
- [11] Challe, Edouard, Julien Matheron, Xavier Ragot and Juan Rubio-Ramirez. 2016. “Precautionary Savings and Aggregate Demand,” Working Paper.
- [12] Christiano, Lawrence, Martin Eichenbaum and Sergio Rebelo. 2011. “When is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119(1) 78-121.
- [13] 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.
- [14] Cooley, Thomas and Gary Hansen. 1991. “The Inflation Tax in a Real Business Cycle Model,” *American Economic Review*, 79(4), 733-748.
- [15] 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.
- [16] Eggertsson, Gauti and Michael Woodford. 2003. “Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity*, 2003(1), 139-233.

- [17] 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.
- [18] Faria-e-Castro, Miguel. 2017. “Fiscal Multipliers and Financial Crises.” NYU Working Paper.
- [19] 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*.
- [20] Fernald, John, Robert Hall, James Stock, and Mark Watson. 2017. “The Disappointing Recovery of Output after 2009,” *Brookings Papers on Economic Activity*.
- [21] 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.
- [22] 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.
- [23] Gertler, Mark and Peter Karadi. 2011. “A Model of Unconventional Monetary Policy,” *Journal of Monetary Economics*, 58(1), 17-34.
- [24] Gilchrist, Simon and Zakrajšek, Egon. 2012. “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102(4), 1692-1720.
- [25] Gorea, Denis and Virgiliu Midrigan. 2015. “Liquidity Constraints in the U.S. Housing Market,” NYU Working Paper.
- [26] 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.
- [27] Guerrieri, Veronica and Guido Lorenzoni. 2015. “Credit Crises, Precautionary Savings, and the Liquidity Trap,” Chicago Booth Working Paper.
- [28] Hatchondo, Juan Carlos and Leonardo Martinez. 2009. “Long-Duration Bonds and Sovereign Defaults,” *Journal of International Economics*, 79(1), 117-125.
- [29] Iacoviello, Matteo. 2005. “House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle,” *American Economic Review*, 95(3): 739-764.
- [30] Jones, Callum. 2017. “Unanticipated Shocks and Forward Guidance At the Zero Lower Bound.” NYU Working Paper.
- [31] Justiniano, Alejandro and Giorgio Primiceri. 2008. “The Time Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 98(3), 604-641.
- [32] Justiniano, Alejandro, Giorgio Primiceri and Andrea Tambalotti. 2015. “Household Leveraging and Deleveraging,” *Review of Economic Dynamics*, 18(1), 3-20.
- [33] Kaplan, Greg, and Gianluca Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4) 1119-1239.
- [34] Kehoe, Patrick and Virgiliu Midrigan. 2014. “Prices Are Sticky After All ” *Journal of Monetary Economics*.

- [35] Kehoe, Patrick, Virgiliu Midrigan and Elena Pastorino. 2016. “Debt Constraints and the Labor Wedge,” forthcoming, *American Economic Review Papers & Proceedings*.
- [36] Kehoe, Patrick, Virgiliu Midrigan and Elena Pastorino. 2017. “Debt Constraints and Employment,” forthcoming, *Journal of Political Economy*.
- [37] 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.
- [38] Kulish, Mariano, and Adrian Pagan. 2017. “Estimation and Solution of Models with Expectations and Structural Changes,” *Journal of Applied Econometrics*, 32(2), 255-274.
- [39] Landvoigt, Tim, Monika Piazzesi and Martin Schneider. 2015. “The Housing Market(s) of San Diego,” *American Economic Review*, 105(4), 1371-1407.
- [40] Lepetyuk, Vadym, Lilia Maliar and Serguei Maliar. 2017. “Should Central Banks Worry About Nonlinearities of their Large-Scale Macroeconomic Models?” Bank of Canada Working Paper.
- [41] Lucas, Robert Jr., 1990. “Liquidity and Interest Rates,” *Journal of Economic Theory*, 50(2) 237-264.
- [42] Lucas, Robert Jr. and Nancy Stokey, 2011. “Liquidity Crises.” Federal Reserve Bank of Minneapolis Economic Policy Papers.
- [43] Lustig, Hanno and Van Nieuwerburgh, Stijn. 2005. “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 60(3), 1167-1219.
- [44] Mendoza, Enrique 2010. “Sudden Stops, Financial Crises, and Leverage,” *American Economic Review*, 100(5), 1941-1946.
- [45] 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.
- [46] Mian, Atif, Kamalesh Rao and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128(4), 1687-1726.
- [47] Mian, Atif, and Amir Sufi. 2014. “What Explains the 2007-2009 Drop in Employment?” *Quarterly Journal of Economics*, 82(6), 2197-2223.
- [48] Mackowiak, Bartosz, and Mirko Wiederholt. 2009. “Optimal Sticky Prices under Rational Inattention.” *American Economic Review*, 99(3): 769-803.
- [49] Nakamura, Emi and Jón Steinsson. 2014. “Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions.” *American Economic Review*, 104(3), 753-792.
- [50] Nakamura, Emi and Jón Steinsson. 2017. “Identification in Macroeconomics.” forthcoming *Journal of Economic Perspectives*.
- [51] Werning, Ivan. 2012 “Managing A Liquidity Trap: Monetary and Fiscal Policy.” MIT Working Paper.

Table 1: Assigned Parameters

Parameter	Value	Description	Source/Target
ν	2	Inverse labor supply elasticity	
γ	0.985	Persistence coupon payments	13 year mortgage debt duration
ω	0.7	Weight on non-traded goods	
σ	0.5	Elasticity traded/non-traded	
κ	4	Elasticity traded goods	Simonovska and Waugh (2014)
ψ	21	Elasticity labor aggregator	Christiano, Eichenbaum and Evans (2005)
α_r	0.86	Taylor rule persistence	Justiniano and Primiceri (2008)
α_π	1.71	Taylor coefficient inflation	Justiniano and Primiceri (2008)
α_y	0.05	Taylor coefficient output	Justiniano and Primiceri (2008)
α_x	0.21	Taylor coefficient output growth	Justiniano and Primiceri (2008)
Parameters chosen to match steady-state target			
$-\ln(\beta)$	2.31%	Annual discount rate	2% real rate
η	0.077	Weight on housing	Housing-to-income ratio of 2.5
$\bar{m}$	0.0044	Credit limit	Debt-to-housing ratio of 0.29

Table 2: Estimated Structural Parameters

Parameter	Prior				Posterior			
	Dist	Median	10%	90%	Mode	Median	10%	90%
α	N	2.6	1.5	3.8	3.68	3.60	3.18	4.11
λ_p	B	0.5	0.2	0.8	0.97	0.96	0.94	0.98
λ_w	B	0.5	0.2	0.8	0.85	0.85	0.82	0.88

Table 3: Moments from State-Level Simulations

Moment	Data	$\alpha = 3.68$	$\alpha = 2$	$\alpha = 10$
Correlation ($\Delta \text{empl}_t, \Delta \frac{\text{debt}_t}{\text{income}_t}$)	0.17	0.19	0.58	0.00
Correlation ($\Delta \text{cons}_t, \Delta \frac{\text{debt}_t}{\text{income}_t}$)	0.23	0.36	0.70	0.06

Table 4: Variance of Consumption and Employment Due to Each Filtered Shock, %

Shock Variable	LTV	Housing	Productivity	Leisure	Discount	Fed Funds	Markup
A. State-level, 1999–2015							
Employment	10.3	0.2	28.8	13.5	47.2	—	—
Consumption	18.0	0.3	0.1	3.2	78.5	—	—
B. Aggregate-level, 1984–2008							
Employment	0.2	0.0	32.3	0.3	24.4	32.5	10.3
Consumption	0.2	0.0	0.9	0.5	35.7	47.6	15.1

Table 5: Estimated Structural Parameters: Robustness

Parameter	A. Regional Data			B. $\alpha = 5$		
	Mode	10%	90%	Mode	10%	90%
α	3.60	3.35	4.27	-	-	-
λ_p	0.93	0.89	0.96	0.97	0.94	0.98
λ_w	0.58	0.46	0.66	0.86	0.83	0.89
Parameter	C. $\gamma = 0.965$			D. $\psi = 5$		
	Mode	10%	90%	Mode	10%	90%
α	2.97	2.69	3.68	3.15	2.95	3.88
λ_p	0.96	0.94	0.98	0.97	0.94	0.98
λ_w	0.86	0.82	0.88	0.93	0.91	0.94
Parameter	E. Construction Model					
	Mode	10%	90%			
α	3.60	2.98	3.80			
λ_p	0.97	0.94	0.98			
λ_w	0.86	0.82	0.88			

Figure 1: State Debt, Employment, and Spending

Figure 2: Timing of the Model

Figure 3: Equilibrium Real Interest Rate in Steady State

Figure 4: Impulse Response to a Credit Tightening. Simple Model

Figure 5: Response to State-Specific Credit Shock. Modal Estimates

Figure 6: Response to State-Specific Credit Shock. Low Idiosyncratic Uncertainty ($\alpha = 5$)

Figure 7: Response to State-Specific Credit Shock. Flexible Prices and Wages

Figure 8: Aggregate Impulse Response to Credit Shock

Figure 9: Dynamics of Sectoral Employment During the Recession

Figure 10: Dynamics of the Mortgage Rate

Figure 11: Effect of Credit Shocks on State Employment and Consumption

Figure 12: Effect of Credit Shocks Only on Aggregate Variables

Figure 13: Effect of Credit Shocks on Aggregate Employment

Figure 14: Effect of Credit Shocks on State Employment: Robustness

Figure 15: Effect of Credit Shocks on Aggregate Employment: Robustness

Figure 16: Effect of Credit Shocks on State Employment: Model with Construction