Household Leverage and the Recession*

Callum Jones† Virgiliu Midrigan‡ Thomas Philippon§

December 2017

Abstract

During the Great Recession, employment declined more in regions 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 using Bayesian likelihood methods on state-level and aggregate data. Credit shocks account well for the differential rise and fall of employment across individual states. Credit shocks explain a smaller fraction of the initial drop in aggregate employment but the tightening of household credit greatly contributes to the slow recovery in the aftermath of recession.

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

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

*We thank Sonia Gilbukh for excellent research assistance. We are indebted to Fernando Alvarez, Andy Atkeson, David Backus, Patrick Kehoe, Andrea Ferrero, Mark Gertler, Veronica Guerrieri, Erik Hurst, Ricardo Lagos, Guido Lorenzoni, Robert Lucas, Atif Mian, Tom Sargent, Robert Shimer, Nancy Stokey, Amir Sufi, Ivan Werning, Mike Woodford and numerous seminar participants for comments. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management.

†International Monetary Fund, jonescallum@gmail.com

‡New York University, virgiliu.midrigan@nyu.edu

§New York University, tphilipp@stern.nyu.edu

1 Introduction

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

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.[1]

The goal of our paper is to quantitatively evaluate the household leverage view of the recession. The difficulty with interpreting the evidence in Figure 1 is that 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 generate changes in that state's employment and the amount households borrow. We therefore use data on employment, consumption, wages, house prices and household debt at the state level, together with the corresponding series in the aggregate, to isolate the effect of the various shocks. We then ask: how important were exogenous shocks to the households' ability to borrow in generating the variation in employment and consumption observed in the data?

We answer this question by studying a parsimonious and tractable model of a monetary union composed of a large number of states that trade among themselves. Households in each state face a liquidity constraint that restricts 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, akin to a discount rate shock typically assumed exogenously in the New Keynesian literature. 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 is subject to exogenous shocks. Households in each state are subject to a number of additional exogenous shocks to productivity and preferences. All shocks have both a state-specific and an aggregate component.

We estimate the model and the shocks using Bayesian likelihood methods. One complication that arises in our estimation is that the policy interest rate is at the *zero lower bound* (ZLB)

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

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 proposed by Jones (2017), itself a variant of methods developed by Eggertsson and Woodford (2003) and Guerrieri and Iacoviello (2015).

Our paper makes three main contributions. The first contribution is to provide a quantitative general equilibrium assessment of the role of household deleveraging during the Great Recession. In doing so, we discover that deleveraging plays a relatively more important role in 2012 than in 2009. Our second contribution is therefore to the literature on the slow recovery, recently discussed by Fernald et al. (2017).

Our third contribution is methodological. To the best of our knowledge, our paper is the first to combine regional and aggregate data to form the likelihood and identify the structural parameters of the model. Combining cross-sectional and time series information turns out to be critical for the identification of the model, and we believe that our approach – in particular the decomposition of the model's solution into time-varying aggregate matrices and time-independent cross-sectional matrices – can be fruitfully applied to a variety of other contexts. 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).

We find that exogenous state-level shocks to the loan-to-value constraint, in short, *credit shocks*, account for about half of the differential rise and then fall in state-level employment and consumption during the 2001 to 2012 period. Using the model, we generate counterfactual series for employment and consumption using credit shocks alone and find that those series have correlations of between 0.3 to 0.7 with the actual movements in state-level employment and consumption and are almost as volatile as the data. Our findings are thus consistent with those of Mian and Sufi (2011, 2014) who argued that household credit played an important role in determining variations in regional consumption and employment.

What are the aggregate implications of credit shocks? Answering this question is much more difficult due to the non-linearity implied by the ZLB and the Fed's pursuit of unconventional monetary policy. Our parameter estimates imply that if credit shocks were the only source of shocks, the short-term nominal interest rate would have fallen by about 1.5 percent over the 2009 to 2015 period but would have never reached the zero bound. Because monetary policy can adjust to accommodate aggregate credit shocks, they would generate a relatively modest 0.5 percent drop in consumption and employment, and thus account for less than one-tenth of the actual decline observed in the data.

Credit shocks, however, have a much larger impact on employment if they are also accompanied by other shocks that render the monetary authority unable to further reduce the short-term nominal rate to accommodate the credit shocks. Indeed, we find that credit shocks had a much larger impact during the 2009 to 2015 period when the ZLB did bind. Absent these shocks, employment would have been 1.4 percentage points higher in 2009, and 2.8 percentage points higher in 2015. Since employment was about 10 percent below trend in 2009 and 6 percent below trend in 2015, we conclude that credit shocks account for a fairly modest fraction of the initial employment decline, but played an increasingly important role in accounting for the slow recovery of aggregate employment in the aftermath of the Great Recession.

The model we study is a flexible extension of Iacoviello (2005). That paper, and much of the literature that builds on it, generates a motive for borrowing and thus a wedge between the rate of time preference and the interest rate by assuming differences in agents' discount factors. Our paper, in contrast, assumes idiosyncratic shocks to agents' liquidity needs. Aggregating these shocks across individuals gives rise to a wedge between the discount rate and the interest rate that 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 households' needs for liquidity are, the model either collapses to a frictionless benchmark in which credit is simply a veil and 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 an unexpected medical bill say 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 (2016). 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 any time, 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] We view this set of assumptions as a convenient modeling device that allows

[2]See also Lucas and Stokey (2011) who emphasize the role of liquidity frictions in the recent financial crisis and Challe et. al. (2015) who employ a family construct in order to characterize the aggregate properties of an

us to parsimoniously capture the role of household credit in the macroeconomy. Our model is flexible yet sufficiently tractable to be used in a Bayesian estimation in which we also account for the non-linearities associated with the ZLB.

Despite risk sharing at the family level, the quantity of HH 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, which is itself a function of all past and present aggregate shocks. However, up to a first-order approximation, the marginal effect of a shock is uniquely pinned down by the expected duration. We thus augment the data on aggregate variables with survey data on the expected duration of the ZLB 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, in addition to these parameters, the persistence and size 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 2001 to 2012. At the aggregate level, we use the aggregate counterparts of the state-level series, together with the Fed Funds rate and inflation, all observed quarterly from 1984 to 2015. Our data on the expected duration of the ZLB from 2009 to 2015 comes from survey data from the New York Fed.

Directly combining aggregate series with their counterparts from a panel of states is computationally challenging in the presence of non-linearities arising from the ZLB, because the state-level observables are a non-linear function of both state and aggregate shocks. To make

economy with uninsurable unemployment risk.

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 components and an independent aggregate component. 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 of the model. We show that our parameter estimates are well-identified by this data.

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 HH credit crunch. These researchers find, as we do, that a credit crunch has a minor effect on employment away from the ZLB. While they study a closed-economy setting, our model is that of a monetary union composed of a large number of states. Moreover, our focus is on estimating the model using state-level 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 wages appear more flexible in the cross-section than in the time series. Our estimation strategy takes both into account and computes the best intermediate value. The gap between estimates of wage rigidity 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) when regional shocks are more volatile and more persistent than aggregate shocks. 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). This literature focuses on understanding the role of shocks that disrupt financial intermediation, which we argue must accompany household credit shocks for the model to be able to replicate the large decline in U.S. employment.[3]

[3]See also the work Lustig and van Nieuwerburgh (2005), Garriga, Manuelli and Peralta-Alva (2014), Favilukis,

2 A Simplified Closed-Economy Real Model

We first describe a simplified version of the richer 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 interact to generate an endogenous spread between the equilibrium interest rate and the household's rate of time preference. Because the interest rate is lower than the rate of time preference, borrowing constraints bind in equilibrium. A tightening of borrowing constraints or an increase in idiosyncratic volatility reduces the equilibrium interest rate. We also show that the extent to which the equilibrium interest rate responds to a credit tightening depends on the amount of idiosyncratic volatility: the larger is volatility, the larger the drop in the equilibrium interest rate 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

$$y_t = n_t. \tag{1}$$

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

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

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

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

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

Ludvigson and van Nieuwerburgh (2015), Burnside, Eichenbaum and Rebelo (2015), Landvoigt, Piazzesi and Schneider (2015) who study the determinants of house prices.

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.[4] 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 is obligated 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

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

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:

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

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

The budget constraint states that

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

The left-hand side of this expression gives the household's uses of funds: transfers to individual members, x_t, as well as new housing purchases, $e_t(h_{t+1} - h_t)$. The right-hand side is the amount of resources available: the amount of labor income earned in that period, $w_t n_t$; the amount received 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 it owns. Each unit of the

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

security the household owns pays one unit in coupon payments and can be sold at a price γq_t reflecting the geometric decay of the payments.

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

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

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

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

Notice that our specification of the borrowing limit only restricts a household's ability to take on new loans. A tightening of the credit limit precludes agents from taking on new loans, but does not force prepayment of old debt. Had we assumed a limit on the stock of debt, a tightening of credit limits would force agents to deleverage immediately, which would be counterfactual. In our empirical work we resort to piece-wise 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, $q_t b_{t+1} \leq \gamma q_t \gamma b_t + m_t e_t h_{t+1}$. In particular, allowing households to save using additional instruments of shorter maturities would be inconsequential.

Notice that our model differs from cash-in-advance models in that we assume that the household can access its date t labor income $w_t n_t$ immediately. Moreover, 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

$$q_t a_{t+1} = x_t - \int_0^1 c_{it} \mathrm{d}i. \tag{9}$$

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 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 return of the long-term security:

$$R_{t+1} = \frac{1 + \gamma q_{t+1}}{q_t}. \tag{10}$$

The first-order condition that determines x_t is then

$$\mu_t = \beta \mathbb{E}_t \mu_{t+1} R_{t+1} + \int_0^1 \xi_t(v) \mathrm{d}F(v), \tag{11}$$

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, unspent funds have a return R_{t+1} in the following period and 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 average 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

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

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

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

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 preferences[5] the choice of consumption reduces to

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

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

$$a_{t+1} = b_{t+1}. \tag{15}$$

The supply of labor is given by

$$n_t^{\nu} = \mu_t w_t. \tag{16}$$

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

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

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

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

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

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

Intuitively, the right-hand side of (18) is equal (up to a first-order approximation) to the difference between the discount rate and the interest rate, while the left-hand size is proportional to the fraction of constrained household members, that is, those with $v > x_t/\underline{c}_t$. As the gap between the discount rate and the interest rate increases, it becomes costlier for households to save, transfers fall relative to consumption, so more members end up constrained.

Consider next the household's total consumption expenditures, $c_t = \int_0^1 c_t(v)\mathrm{d}F(v)$. We have

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

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

Finally, letting

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

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

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

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 would like to borrow as much as possible in this economy as long as the interest rate is below the rate of time preference, which is indeed the case due to uncertainty about taste shocks.

2.5 Steady State Equilibrium Interest Rate

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

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

Intuitively, 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 ration $\bar{m}$ 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 about the taste shocks. Notice how the intersection of the upward-sloping savings curve and the downward-sloping debt curve determines the equilibrium interest rate. A tightening of the debt limit reduces the demand for debt, thus reducing the interest rate. 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, $\varepsilon_{m,t}$, which, due to the long-term nature of securities, leads to a gradual reduction in household debt. We contrast the responses in economies with a high and low degree of idiosyncratic uncertainty.

Notice that the equilibrium interest rate falls in both economies. The interest rate falls more in the economy with greater idiosyncratic uncertainty, reflecting the steeper savings curve. In contrast, output barely falls. Although a tightening of credit magnifies the consumption-leisure distortions, these are small here, as in cash-in-advance models.[6]

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

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

[6] Cooley and Hansen (1991).

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.[7] An important difference, however, is that in our model the natural rate is determined endogenously, by the availability of credit as well as the strength of the precautionary-savings motive, as determined by the amount of idiosyncratic volatility.

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 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.

Since we estimate the model using Bayesian likelihood techniques, we allow for several additional shocks to be able to fit the aggregate and state-level data, in addition to the shocks to credit limits. 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 impose a priori that credit shocks account for all variation in discount rates in the data, we allow for a separate exogenous source of movements in discount rates unrelated to credit. We also include shocks to households 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 in the previous section. We let s index an individual island and $p_t(s)$ denote the price of the final consumption good. We assume perfect risk-sharing across households belonging to different labor unions on a given island. Because of separability in preferences, risk-

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

sharing implies that all households on an island make identical consumption, housing and savings choices, even though their labor supply differs depending on when the union that represents them last reset its wage. The problem of a household that belongs to labor union ι is to

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

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 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 are equal to the sum of an island-specific component and an aggregate component, all of which follow AR(1) processes with 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

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

There are no barriers to capital flows, so all islands trade securities at a common price q_t. As before, we assume, as with all other shocks, that the credit limit $m_t(s)$ is the sum 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 loan-to-value 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.[8] We note, however, that in this model changes in housing preferences only affect island and economy-wide variables through their effect on the amount of debt households can take on. This follows because housing is separable in the utility function and the housing stock is fixed.

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:

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

where ω determines the share of non-traded goods, σ is the elasticity of substitution between traded and non-traded goods and κ is the elasticity of substitution between varieties of the traded goods produced on different islands. 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

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

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

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

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

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

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

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

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

[8] See 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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:

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

3.2 Monetary Policy

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

$$1 + i_t = \mathbb{E}_t R_{t+1} \tag{41}$$

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 goods producers shows that the aggregate inflation equation satisfies, 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, and $\bar{\beta}$ is the steady state discount factor and $\bar{\pi}$ is the steady-state level of inflation.[9]

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

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

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

[9]We 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.

state level, and output growth, respectively. We assume that $\bar{\imath}$ 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:

$$\int_0^1 a_{t+1}(s)\mathrm{d}s = \int_0^1 b_{t+1}(s)\mathrm{d}s. \tag{42}$$

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 Jones (2017), itself a variant of methods developed by Eggertsson and Woodford (2003) and Guerrieri and Iacoviello (2015).[10] 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

[10] See Jones (2017) who contrasts this method with fully non-linear ones in smaller scale models and shows that that the piece-wise linear method we use here is reasonably accurate.

convergence. Once the algorithm has converged, the solution of the model is

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

where $\mathbf{x}_t$ collects the endogenous variables, ϵ_t collects the shocks, and the matrices $\mathbf{J}$, $\mathbf{Q}$ and $\mathbf{G}$ are the solutions of the model. Note that, unlike in simple models without a ZLB constraint, these matrices are time-varying.[11]

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 Credit Shock Figure 5 reports the impulse response to a state-level credit shock in the baseline parameterization using the modal posterior estimates discussed below. Panel A shows that a one-time shock to the credit limit $m_t(s)$ of an individual island leads to a gradual decline in 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, implying that the island runs a current account surplus and thereby reduces 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 an increase in the markups of firms and unions that do not reset their prices.

Figure 6 shows the responses to the same state-level credit shock in an economy with a lower degree of idiosyncratic uncertainty by setting $\alpha = 4$ compared to 2.7 at our modal estimates. Notice that the responses of wages, employment and consumption are all dampened compared to those in the baseline model. 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 therefore barely change.

The intuition for this greater decline in the island's savings is as follows. Households can respond to a tightening of credit in two ways: either by reducing their savings or by cutting consumption. When idiosyncratic uncertainty is low, it is relatively costless to reduce savings and so an island's savings fall nearly as much as debt. Both sides of the island's balance sheet thus contract, with little impact on other variables. In contrast, when idiosyncratic uncertainty is high, reducing savings is costly since individual members are more likely to be liquidity

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

constrained. The household thus finds 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 by either reducing overall consumption, thus worsening the intertemporal allocations, or by reducing savings, thus worsening the allocation of consumption across agents. The more dispersed the idiosyncratic shocks are, the more the household chooses to reduce the overall level of consumption to avoid the high costs of variation of the marginal utility of consumption across its members. Our Bayesian estimation below identifies the parameter α largely from the comovement of state-level household debt and state-level consumption.

Figure 7 illustrates the role of price rigidities in determining the state-level responses to a credit shock by setting $\lambda_p = \lambda_w = 0$ thereby shutting them down altogether. Notice that a decline in debt is now associated with a sharper decline in wages, an increase in employment, owing to wealth effects on labor supply, and a smaller decline in consumption. Intuitively, wage rigidities in this environment act like a tax on labor supply while price rigidities lead to an increase in firm markups and thus reduce real wages. Both of these forces prevent employment from increasing following a credit tightening.[12]

Impulse Response to an Aggregate Credit Shock To understand the effect of an aggregate credit shock in our model, recall from Equation (23) that credit shocks act like discount rate shocks in simpler New Keynesian models and therefore change the natural interest rate. The degree to which a shock to credit affects real variables therefore depends crucially on the stance of monetary policy and whether the monetary authority can accommodate 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 notable feature of the model is that the nominal interest rate immediately falls to zero rather than gradually declining as the natural rate does. This reflects the anticipation in the future of

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

the contractionary effect of a binding ZLB which reduces the output gap today, triggering an immediate decline in the policy rate.

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.

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.[13] 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

[13]The 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 $(1+r)/(1+r-\gamma)$.

wage markup of 5%.[14] 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 loan-to-value ratio is chosen so that the aggregate debt to housing ratio is equal to 0.29, a number once again computed from the SCF. Since the debt constraint binds in the model, these two last two targets imply an aggregate debt to (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. We first describe the details of our approach, and then report our results.

4.2.1 The Parameters

We estimate a total of 26 parameters, which include the three structural parameters: α, λ_p, λ_w, the persistence and volatility of the 5 state-level components of the shocks to productivity, the discount factor, preference for housing, disutility from work, and the credit limit, as well as the persistence and volatility of the 7 aggregate components of the shocks.

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 2001 to 2012. The household debt information is from the FRB New York Consumer Credit Panel,[15] the data on house prices are from the FHFA, while the data on employment, consumption expenditures and wages are from the BEA.

[14] We show in our robustness exercises below that our conclusions are not sensitive to alternative values of this parameter, in that the estimation procedure compensates for a change in ψ by changing the estimated degree of Calvo price and wage stickiness, with little effect on the model's implications for real and nominal variables.

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

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 follow as closely as possible the sources used by Smets and Wouters (2007) for each of these series. A 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.[16]

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.

A direct approach to estimating the model would be to write down a likelihood function that directly uses state and aggregate data and the particular representation (43). This direct 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 direct approach infeasible.

We thus formulate an alternative, mathematically equivalent, but computationally feasible approach to constructing the likelihood function, one that exploits *relative* variation across individual states' outcomes during the 2001 to 2012 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, namely the amount of savings and debt in Nevada and the Nevada-specific component of the shocks. 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:

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

[16] See the Appendix for a more detailed description of the data we use.

Here, $\mathbf{Q}^s$ and $\mathbf{G}^s$ encodes the dependency of island j's variables 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:

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

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 state-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, changes in the regime do not affect the evolution of state-level variables, conditional on the aggregates. From the perspective of agents in an individual state, say Nevada, 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 $\mathbf{x}_t = \int \mathbf{x}_t^j \mathrm{d}j$ denote the economy-wide average of the state-level variables, we have that the deviation of state-variables from their economy-wide averages,

$$\hat{\mathbf{x}}_t^j = \mathbf{x}_t^j - \mathbf{x}_t, \tag{46}$$

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

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

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

We therefore use in our estimation (45) alongside (47) in forming the state-space rather than the much more computationally intensive (44). To do so, we first express each state-level variable in 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 for each series given that our model assumes that islands are ex-ante identical. 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.

We use a Kalman filter to compute forecast errors that enter the calculation of the likelihood,[17] 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 2001 population shares.[18] We assign an equal weight to the aggregate likelihood and the combined state-level likelihood. We compute two independent chains of 150,000 draws and evaluate their convergence in the Appendix.

4.3 Parameter Estimates

Table 2 reports moments of both the prior and posterior distributions of the parameters we estimate. We start with diffuse priors, similar to those in Smets and Wouters (2007). Consider next the posterior distribution.

The modal estimate of the Pareto tail parameter α is equal to 2.7, implying a 1.8% wedge between the discount rate and interest rate at the annual frequency. This modal estimate implies that credit constraints have a relatively small impact on steady-state employment: the implicit employment tax levied by liquidity constraints generates only a 0.1% decline in employment. We also note that the posterior distribution for α is relatively tight around its mode, with a 10th and 90th percentile of 2.5 and 3.1, respectively. As we discussed earlier, the comovement between spending and household debt at the state level is important for identifying this parameter.

We find, consistent with the work of Del Negro, Giannone and Schorfheide (2015), that

[17]Details of how we construct the likelihood function are discussed in the Appendix.

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

wages and prices are sticky, with a modal estimate of λ_w of 0.88 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 experiments reported below.

Table 2 also reports the posterior estimates of the parameters governing the persistence and volatility of state and aggregate shocks. Since these are more difficult to interpret on their own, in the Appendix we report theoretical forecast error variance decompositions of state and aggregate observables into the various shocks at various horizons. These theoretical decompositions simply allow us to put the estimates of the persistence and volatility of shocks in perspective. 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 to the variance of consumption and employment.

Consider first Panel A of Table 3 which reports our decomposition for state-level variables. Recall that these are expressed in deviations from their aggregate counterparts. State-level credit shocks account for about 40 percent of the relative variation in household spending at the regional level and about 20 percent of the relative variation of employment. State-level productivity shocks account for almost 25 percent of the relative variation in employment and only 1 percent of the relative variation in household spending. Shocks to the disutility from work account for about 30 percent of the volatility of employment and 11% of the volatility in consumption. Finally, shocks to the individual states' discount rates account for about 25% percent of the volatility of employment and 45% of the volatility of consumption. Shocks to the preference for housing have little impact on state-level real variables and are mostly needed to soak up the variation in house prices across individual states.

Consider next Panel B of Table 3 which reports our 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 about 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. Finally, productivity shocks account for about a quarter of the employment volatility and 5 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.

5 Credit Shocks in the Great Recession

We next study role of household credit shocks, as captured by the loan-to-value parameter $m_t(s)$, in shaping the dynamics of employment and consumption in the cross-section and in the aggregate during the Great Recession. 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 by using a Kalman filter to extract their aggregate component. We show that aggregate credit shocks generate about a 2% decline in the natural rate of interest during the Great Recession, a sizable amount but 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. We conclude by gauging the robustness of our findings to alternative parameterizations of the degree of wage and price stickiness and duration of mortgage contracts.

5.1 Role of Credit Shocks at the State Level

We start by studying the 2002 to 2007 period. In Figure 9 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 in the Kalman filter to extract the state-level component of the shocks. We construct a counterfactual series for each variable by shutting down all shocks other than the credit shocks, 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 9 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 with credit shocks only and the change in the actual data is 0.53 for employment and 0.66 for consumption. The slope coefficient of a regression of these counterfactual changes against the data is 0.59 for employment and 0.30 for consumption. Since the slope coefficient is equal to the product of the correlation and the relative standard deviation, these numbers imply that the change in employment explained by credit shocks is as volatile as in the data, while the change in consumption explained by credit shocks is about 40% as volatile as the data.

The lower panels of Figure 9 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 0.72 for employment and 0.58 for consumption, while the slope coefficients are 0.72 for employment and 0.62 for consumption. Thus, the model implies that credit shocks alone account for about 60% 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 9 would simply collapse to 0. Similarly, if we were to estimate a lower degree of wage and price stickiness, as is the case when we use regional data alone to estimate the structural parameters of the model, 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 10 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, 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 credit shocks. Panel A of Figure 10 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 short-term nominal rate. In the data the nominal interest rate increased from 1.75% in 2002 to 5.25% in 2007 and then fell to 0 by the end of 2009. In contrast, in an economy with credit shocks alone the short-term nominal interest rate would have only increased from 2.7% in 2002 to 3.0% in 2007 and would have approached zero only by the end of 2015.

The reason credit shocks generate much more gradual movements in interest rates is the gradual change in the household debt to income ratio observed 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. As Panel B of Figure 10 shows, credit shocks alone imply an increase in the natural interest rate from 2.4% in 2002 to 3.0% in 2007 and then a decline to 1.1% by the end of 2013. Due to the Taylor rule, the short-term nominal interest rate tracks the natural rate quite closely, causing a slow increase and then decline in the short-term nominal rate. Notably, credit shocks alone were not sufficiently strong to trigger the ZLB on nominal interest rates until 2015.

Consider next how employment responds to credit shocks. As the lower-left panel of Figure 10 shows, the Great Recession was associated with a nearly 10% drop in the employment-population ratio from 2007 to 2010. Credit shocks alone account for a much more modest 0.5% drop in this time period, thus accounting for less than one-tenth of the actual decline. The reason employment falls 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 suggest that absent additional shocks around the 2007 to 2008 period, the household leverage cycle observed in the U.S. would have caused a fairly mild recession. 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 larger?

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 the consequences of the

ZLB. 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.[19]

We therefore need to keep forward guidance constant across our experiments. To make our results as transparent as possible, we consider a counterfactual experiment in which the Fed does not pursue forward guidance, with or without credit shocks. The expected duration of the ZLB is then 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 11 shows what the response of employment would have been in the absence of forward guidance. Clearly, the decline in employment would have been worse, about 3–4 percentage points lower, than in the data, especially between 2012 to 2014, the years when the Fed was explicitly pursuing these policies.

Panel A also reports what employment would have been absent the forward guidance policies as well as absent the credit shocks from the middle of 2007 onward. This counterfactual series, reported as a dashed line is above the dotted line that is computed using all shocks extracted by the Kalman filter. The difference between these counterfactual series allows us to gauge the contribution of credit shocks at the ZLB holding forward guidance constant (at zero). Panel B shows the difference between these two series and compares it to the data. Credit shocks alone generate a 1.4 percentage points decline in employment in 2009, almost 20% of the observed decline, but gradually become more important and by 2015 generate a 2.8% drop in employment, almost half of the employment gap in the data.

We thus conclude that credit shocks indeed had an important effect on employment during the Great Recession owing to constraints on monetary policy. Although credit shocks explain a modest fraction of the large immediate drop in employment in 2009, due to the gradual nature of household deleveraging, we find that they did play a substantial role in delaying the recovery in the aftermath of the recession.

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

5.3 Robustness Checks

We considered several robustness checks.[20] 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. Finally, we have assumed a lower elasticity of substitution ψ between varieties of labor supplied by different unions.

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 4 shows that the posterior distribution of the parameters estimated using regional data alone implies a lower degree of idiosyncratic uncertainty, $\alpha = 3.4$, compared to 2.7 in our baseline estimates that use both aggregate and regional data, a similar degree of price stickiness, $\lambda_p = 0.95$, compared to 0.97 in our baseline and a much lower degree of wage stickiness, $\lambda_w = 0.62$, compared to 0.88 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. Matching the aggregate data thus requires a greater degree of nominal stickiness, while matching 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. (See our Appendix for estimates and counterfactuals derived using aggregate data alone.) 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 12 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 the boom and the recession. Although the counterfactual series for employment constructed with credit shocks alone is quite strongly correlated with the data (0.59 during the boom, and 0.73 during the recession), the slope coefficient is much smaller (0.22 during the boom and

[20] We report the main checks here, with further robustness exercises left to the Appendix.

0.33 during the recession), implying that credit shocks alone account for about 30–40% of the volatility of relative changes in employment in the data.

Panel A of Figure 13 shows that credit shocks alone imply a maximal drop of employment of about 2.5%, compared to 3.8% under our baseline estimates in an economy. As earlier, we compute the contribution of credit shocks by eliminating forward guidance and comparing the employment outcomes in the presence of all shocks extracted from the Kalman filter to those obtained when we shut down the credit shocks and let the expected duration of the ZLB adjust accordingly. Thus, the smaller degree of wage stickiness and idiosyncratic uncertainty favored by the regional data dampens the importance of credit shocks in driving both state-level and aggregate employment.

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 2.7 in our baseline, and re-estimating all other parameters of the model.

As Panel B of Figure 12 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 13 shows, credit shocks explain virtually none of the employment drop in the aggregate: the maximal drop in employment accounted for by credit shocks is only 0.7% compared to 3.8% in our baseline. This robustness check illustrates that our findings that credit shocks play a relatively important role in determining state-level and aggregate employment changes are not an artefact of our modelling choices, but rather a result of the estimation.

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). As Panel C of Figure 12 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 13 shows that credit shocks now account for a slightly larger drop in aggregate employment: the maximal drop in employment is now 4% compared to 3.8% in our baseline.

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 4 shows, the estimation now favors an even greater degree of nominal wage and price stickiness to compensate for the removal of the real rigidity. None of our model's implications change, however, as Figures 12 and 13 show. Intuitively, up to a log-linear approximation, it makes little difference whether one uses real or nominal wage rigidities to reproduce the patterns in the data.

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 Recesssion 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, we find very similar estimates of the Taylor as those of Justiniano and Primiceri (2008). As a result, our model's implications are largely unchanged.

6 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 comoves strongly with changes in household debt in the cross-section of U.S. regions. This paper proposes a theory that captures this view by introducing a role for credit in alleviating household liquidity constraints. The theory is a parsimonious extension of Iacoviello (2005) that allows us to flexibly parameterize the importance of household credit frictions in determining consumption-savings choices. The more uncertain households are about their individual liquidity needs, the stronger their precautionary savings motive and thus 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 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.

We use Bayesian likelihood methods together with state-level and aggregate data, to estimate the parameters of the model and the process for a rich set of shocks that drive regional and aggregate fluctuations. We find that state-level credit shocks account for a substantial fraction of state-level relative movements in employment and consumption, about 20 to 40 percent, during the 2001 to 2012 period. In contrast, aggregate credit shocks account for a much smaller fraction of the variation in employment during the 1984-2008 period, the years prior to the Great Recession. Though our estimates imply that aggregate credit shocks generate substantial movements in the natural rate of interest, estimates of the Taylor rule suggests that monetary policy tracks these movements quite well, implying modest effects on real variables when monetary policy is unconstrained.

In contrast, when monetary policy is constrained, as was the case during the 2009 to 2015 period, household credit shocks have a much larger impact on real activity. Our baseline estimates imply that credit shocks alone caused a maximal 3.8% drop in employment, a sizable fraction of the drop observed in the data. While credit shocks alone account for less than 20% of the large immediate drop in employment during the 2009 to 2010 period, the persistent nature of household deleveraging implies that credit shocks account for about a half of the employment gap in 2015. Changes in household-level credit limits thus help account for the slow recovery of employment in the aftermath of the Great Recession.

References

[1] 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.

[2] Beraja, Martin, Erik Hurst and Juan Ospina. 2015. "The Aggregate Implications of Regional Business Cycles," Chicago Booth Working Paper.

[3] Bernanke, Ben and Mark Gertler. 1989 "Agency Costs, Net Worth and Business Fluctuations," *American Economic Review,* 79, 14-31.

[4] 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.

[5] Blanco, Andres. 2015. "Optimal Inflation Target in an Economy with Menu Costs and an Occasionally Binding Zero Lower Bound," University of Michigan Working Paper.

[6] 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.

[7] Burnside, Craig, Martin Eichenbaum and Sergio Rebelo. 2015. "Understanding Booms and Busts in Housing Prices." Northwestern University Working Paper.

[8] Burstein, Ariel, Martin Eichenbaum and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113(4), 742-784.

[9] Challe, Edouard, Julien Matheron, Xavier Ragot and Juan Rubio-Ramirez. 2016. "Precautionary Savings and Aggregate Demand," Working Paper.

[10] Christiano, Lawrence, Martin Eichenbaum and Sergio Rebelo. 2011. "When is the Government Spending Multiplier Large?" *Journal of Political Economy*, 119(1) 78-121.

[11] 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.

[12] Cooley, Thomas and Gary Hansen. 1991. "The Inflation Tax in a Real Business Cycle Model," *American Economic Review,* 79(4), 733-748.

[13] 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.

[14] Eggertsson, Gauti and Michael Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 2003(1), 139-233.

[15] 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.

[16] 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.*

[17] Fernald, John, Robert Hall, James Stock, and Mark Watson. 2017. "The Disappointing Recovery of Output after 2009," *Brookings Papers on Economic Activity.*

[18] 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.

[19] 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.

[20] Gertler, Mark and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics,* 58(1), 17-34.

[21] Gilchrist, Simon and Zakrajšek, Egon. 2012. "Credit Spreads and Business Cycle Fluctuations," *American Economic Review,* 102(4), 1692-1720.

[22] Gilchrist, Simon and Schoenle, Raphael and Sim, Jae and Zakrajek, Egon. 2015. "Inflation Dynamics During the Financial Crisis," *Working Paper*

[23] Gorea, Denis and Virgiliu Midrigan. 2015. "Liquidity Constraints in the U.S. Housing Market," NYU Working Paper.

[24] 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.

[25] Guerrieri, Veronica and Guido Lorenzoni. 2015. "Credit Crises, Precautionary Savings, and the Liquidity Trap," Chicago Booth Working Paper.

[26] Hatchondo, Juan Carlos and Leonardo Martinez. 2009. "Long-Duration Bonds and Sovereign Defaults," *Journal of International Economics*, 79(1), 117-125.

[27] Guvenen, Fatih and Anthony Smith. 2014. "Inferring Labor Income Risk and Partial Insurance from Economic Choices," *Econometrica*, 82(6), 2085-2129.

[28] Iacoviello, Matteo. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95(3): 739-764.

[29] Jones, Callum. 2017. "Unanticipated Shocks and Forward Guidance At the Zero Lower Bound." NYU Working Paper.

[30] Justiniano, Alejandro and Giorgio Primiceri. 2008. "The Time Varying Volatility of Macroeconomic Fluctuations," *American Economic Review*, 98(3), 604-641.

[31] Justiniano, Alejandro, Giorgio Primiceri and Andrea Tambalotti. 2015. "Household Leveraging and Deleveraging," *Review of Economic Dynamics*, 18(1), 3-20.

[32] Kaplan, Greg, and Gianluca Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82(4) 1119-1239.

[33] Kaplan, Greg, Kurt Mitman, and Gianluca Violante. 2015. "Consumption and House Prices in the Great Recession: Model Meets Evidence," NYU Working Paper.

[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] 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.

[37] Kulish, Mariano, and Adrian Pagan. 2017. "Estimation and Solution of Models with Expectations and Structural Changes," *Journal of Applied Econometrics,* 32(2), 255-274.

[38] Landvoigt, Tim, Monika Piazzesi and Martin Schneider. 2015. "The Housing Market(s) of San Diego," *American Economic Review,* 105(4), 1371-1407.

[39] Lucas, Robert Jr., 1990. "Liquidity and Interest Rates," *Journal of Economic Theory*, 50(2) 237-264.

[40] Lucas, Robert Jr. and Nancy Stokey, 2011. "Liquidity Crises." Federal Reserve Bank of Minneapolis Economic Policy Papers.

[41] Lustig, Hanno and Van Nieuwerburgh, Stijn. 2005. "Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective," *Journal of Finance*, 60(3), 1167-1219.

[42] Mendoza, Enrique 2010. "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100(5), 1941-1946.

[43] 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.

[44] Mian, Atif, Kamalesh Rao and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128(4), 1687-1726.

[45] Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Quarterly Journal of Economics*, 82(6), 2197-2223.

[46] Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99(3): 769-803.

[47] Nakamura, Emi and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions." *American Economic Review,* 104(3), 753-792.

[48] Nakamura, Emi and Jón Steinsson. 2017. "Identification in Macroeconomics." forthcoming *Journal of Economic Perspective,*.

[49] 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			
$-\ln(\beta)$	3.62%	Annual discount rate	2% real rate
η	0.124	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 Parameters

	Prior				Posterior			
Parameter	Dist	Median	10%	90%	Mode	Median	10%	90%
A. Structural Parameters								
α	N	2.6	1.5	3.8	2.66	2.74	2.50	3.08
λ_p	B	0.5	0.2	0.8	0.97	0.96	0.95	0.98
λ_w	B	0.5	0.2	0.8	0.88	0.88	0.85	0.91
B. Regional Shock Processes								
ρ_z	B	0.5	0.2	0.8	0.90	0.76	0.39	0.92
ρ_m	B	0.5	0.2	0.8	0.72	0.67	0.44	0.82
ρ_h	B	0.5	0.2	0.8	0.89	0.85	0.75	0.91
ρ_n	B	0.5	0.2	0.8	0.73	0.73	0.68	0.77
ρ_b	B	0.5	0.2	0.8	0.92	0.91	0.84	0.95
$100 \times \sigma_z$	IG	0.6	0.3	1.9	0.64	0.93	0.50	1.58
σ_m	IG	0.6	0.3	1.9	0.75	0.82	0.55	1.28
σ_h	IG	0.6	0.3	1.9	0.47	0.72	0.35	1.64
σ_n	IG	0.6	0.3	1.9	1.13	1.51	0.84	2.81
$1000 \times \sigma_b$	IG	0.6	0.3	1.9	0.96	1.08	0.57	2.32
C. Aggregate Shock Processes								
ρ_z	B	0.5	0.2	0.8	0.97	0.96	0.94	0.98
ρ_m	B	0.5	0.2	0.8	0.97	0.97	0.95	0.98
ρ_h	B	0.5	0.2	0.8	0.95	0.94	0.92	0.96
ρ_n	B	0.5	0.2	0.8	0.06	0.09	0.03	0.20
ρ_b	B	0.5	0.2	0.8	0.89	0.88	0.86	0.90
ρ_p	B	0.5	0.2	0.8	0.53	0.55	0.45	0.70
$100 \times \sigma_z$	IG	0.6	0.3	1.9	0.61	0.62	0.55	0.70
σ_m	IG	0.6	0.3	1.9	0.19	0.20	0.16	0.27
σ_h	IG	0.6	0.3	1.9	0.15	0.17	0.12	0.22
$\frac{1}{100} \times \sigma_n$	IG	0.6	0.3	1.9	0.20	0.23	0.14	0.40
$1000 \times \sigma_b$	IG	0.6	0.3	1.9	1.50	1.58	1.33	1.90
$1000 \times \sigma_p$	IG	0.6	0.3	1.9	0.88	0.88	0.64	1.11
$100 \times \sigma_r$	IG	0.6	0.3	1.9	1.17	1.20	1.07	1.38

Table 3: Variance of Consumption and Employment Due to Each Filtered Shock, %

Variable \ Shock	Credit	Housing	Productivity	Leisure	Discount	Fed Funds	Markup
A. State-level, 2001–2012							
Employment	21.4	0.3	23.3	30.6	24.4	–	–
Consumption	42.1	0.6	0.9	11.1	45.3	–	–
B. Aggregate-level, 1984–2008							
Employment	1.6	0.0	25.8	1.2	36.0	27.9	7.6
Consumption	2.0	0.1	4.6	1.6	46.2	35.8	9.7

Table 4: Estimated Structural Parameters: Robustness

	A. Regional Data			B. $\alpha = 5$		
Parameter	Mode	10%	90%	Mode	10%	90%
α	3.36	2.93	4.04	-	-	-
λ_p	0.95	0.91	0.97	0.96	0.94	0.97
λ_w	0.62	0.50	0.71	0.87	0.84	0.91

	C. $\gamma = 0.965$			D. $\psi = 5$		
Parameter	Mode	10%	90%	Mode	10%	90%
α	2.38	2.27	2.70	2.75	2.58	3.21
λ_p	0.97	0.95	0.98	0.96	0.95	0.98
λ_w	0.88	0.84	0.89	0.95	0.93	0.96

Figure 1: State Debt, Employment, and Spending

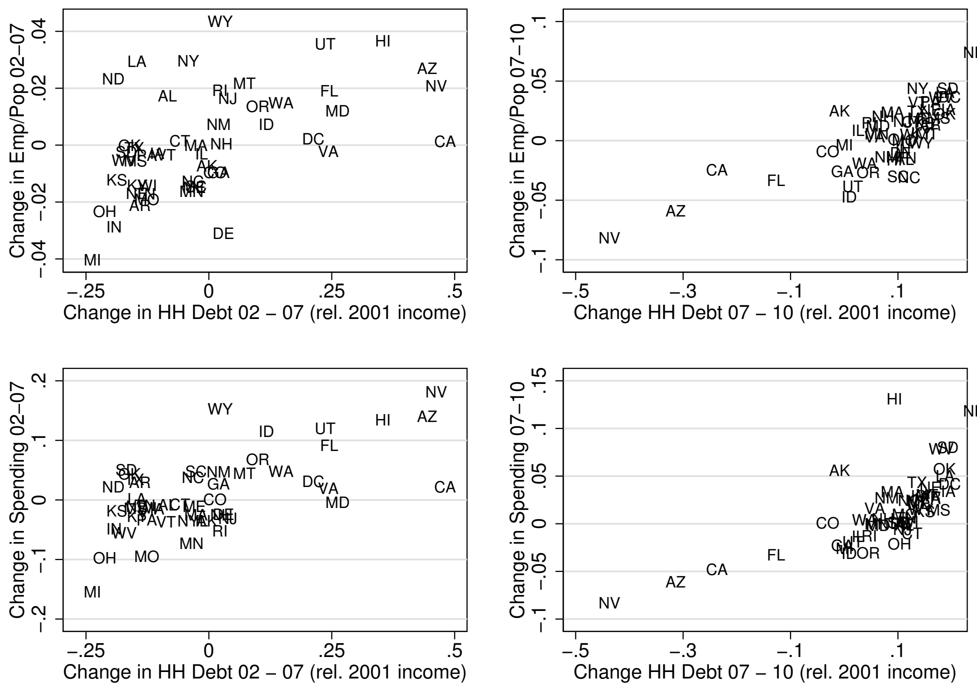

Figure 2: Timing of the Model

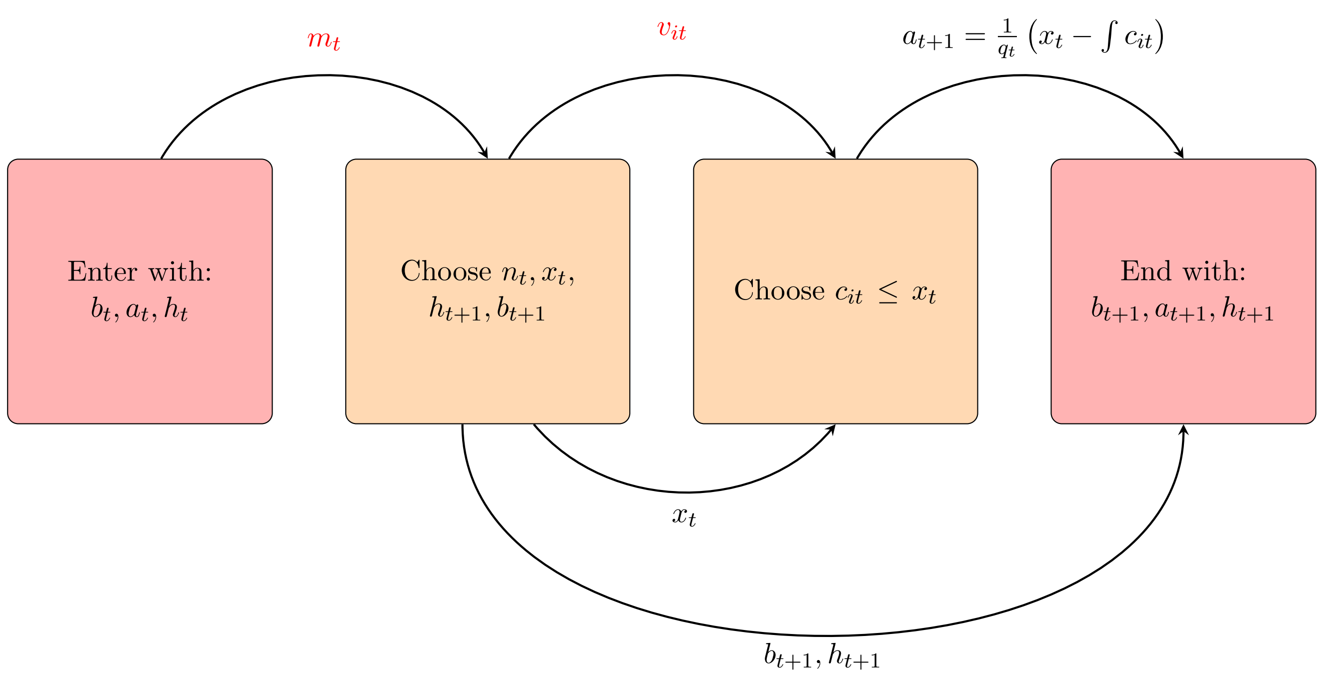

Figure 3: Equilibrium Real Interest Rate in Steady State

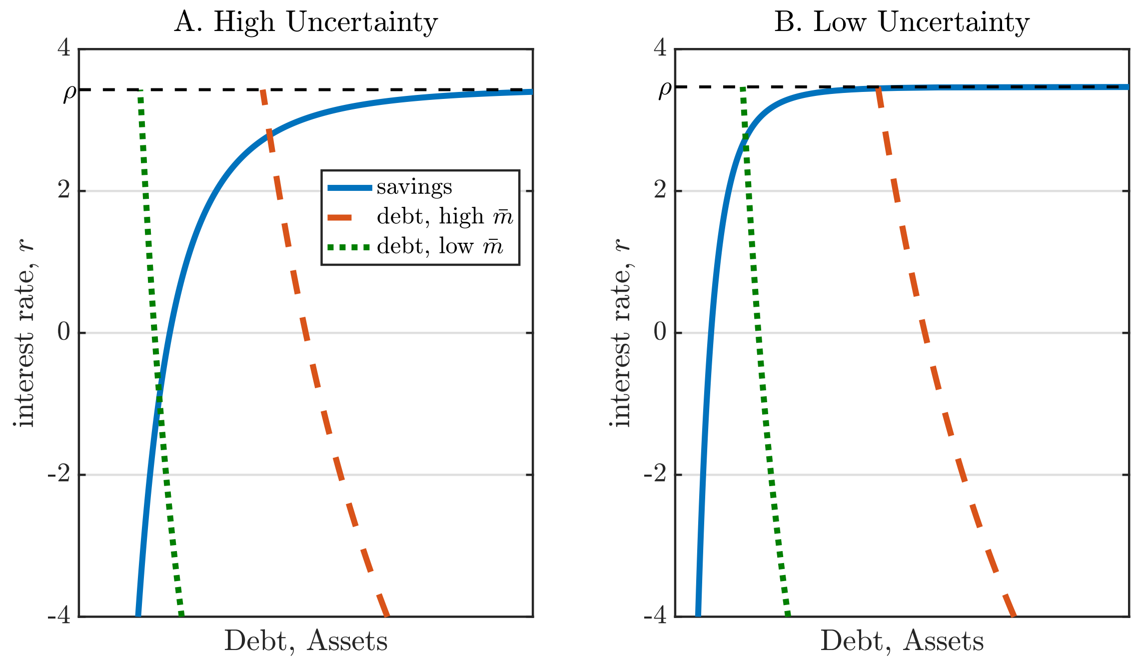

Figure 4: Impulse Response to a Credit Tightening. Simple Model

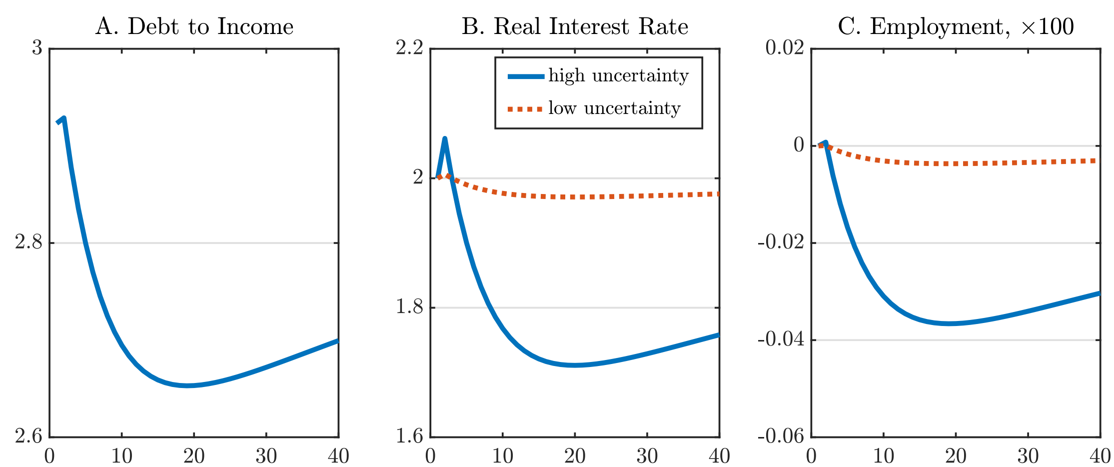

Figure 5: Response to State-Specific Credit Shock. Modal Estimates

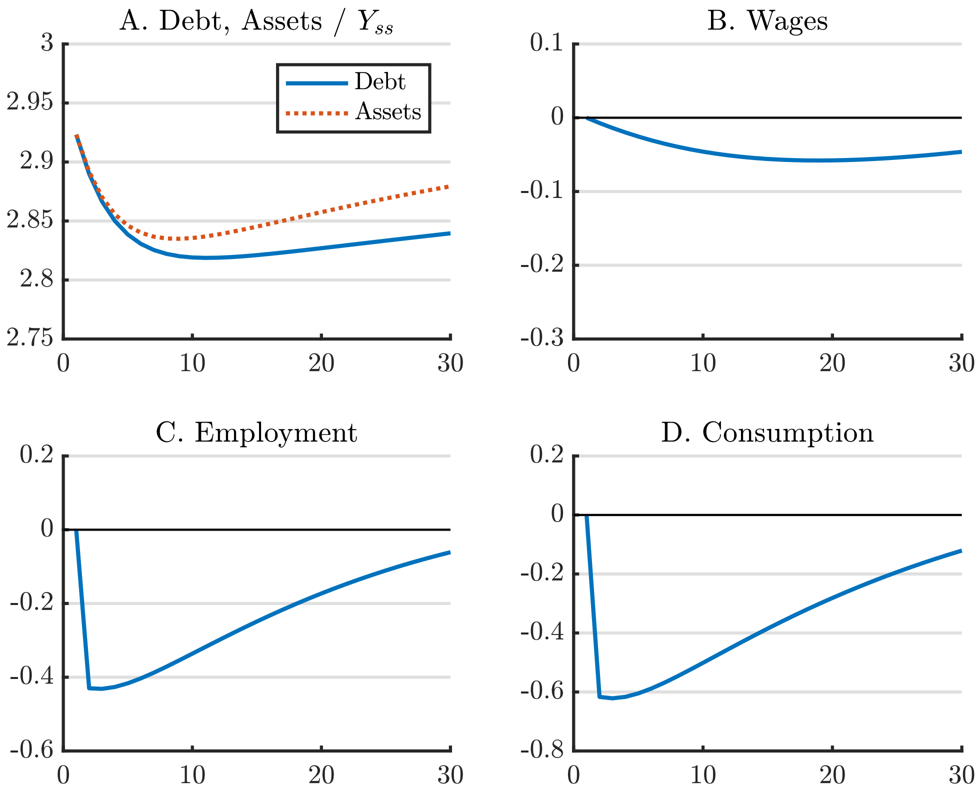

Figure 6: Response to State-Specific Credit Shock. Low Idiosyncratic Uncertainty

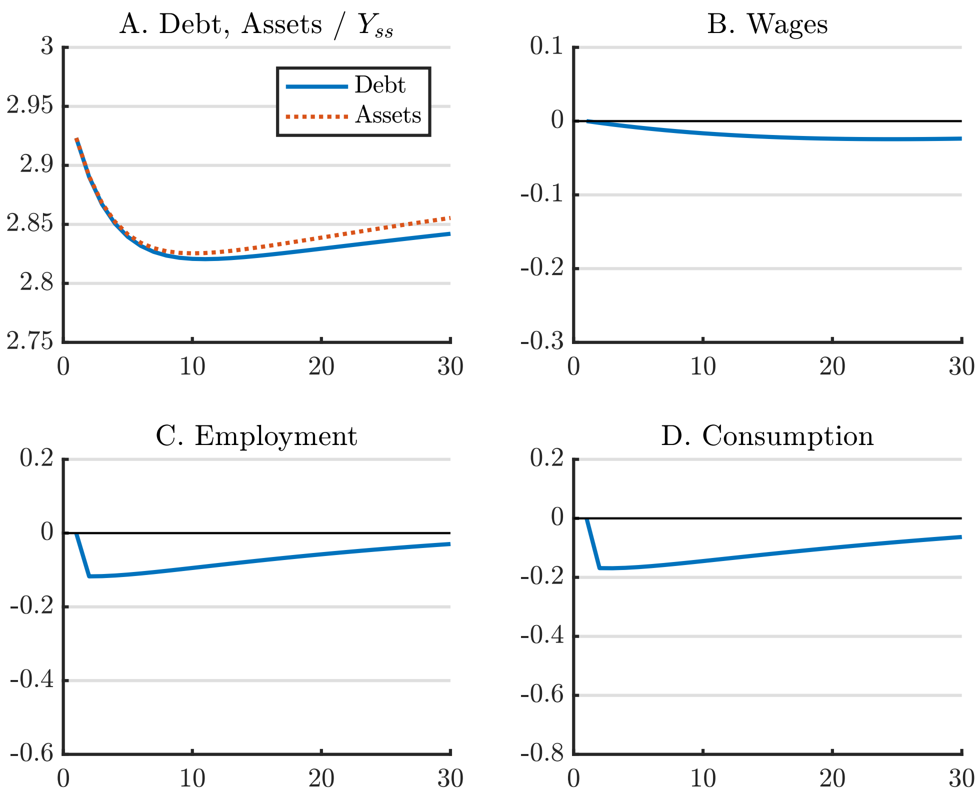

Figure 7: Response to State-Specific Credit Shock. Flexible Prices and Wages

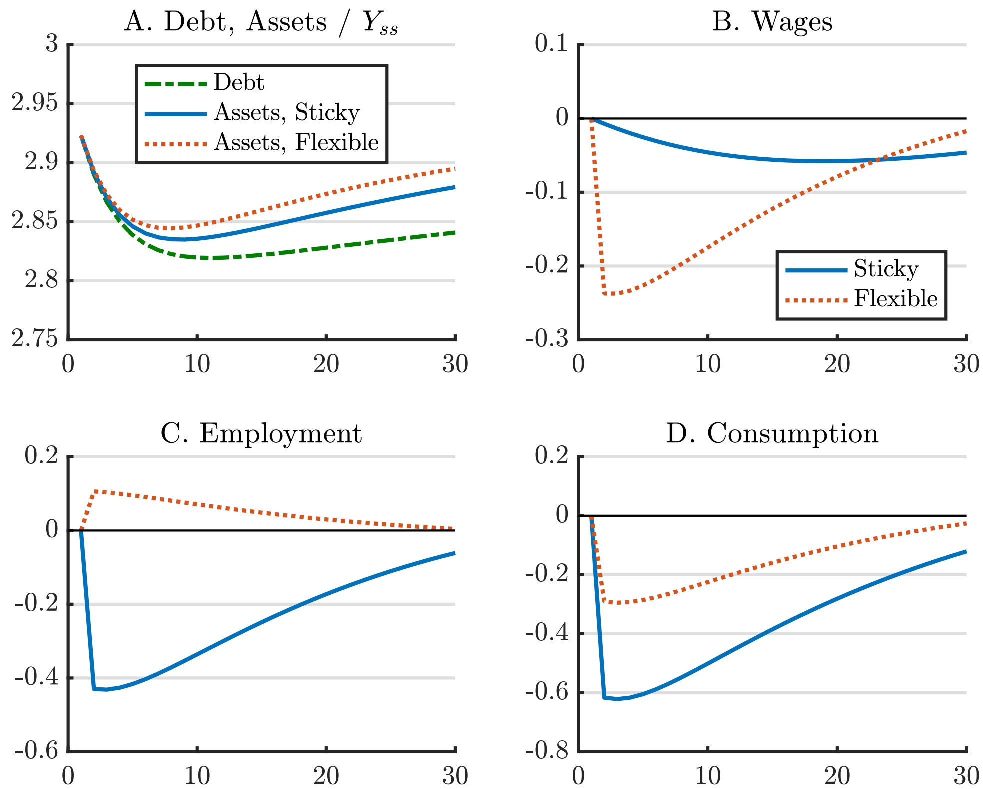

Figure 8: Aggregate Impulse Response to Credit Shock

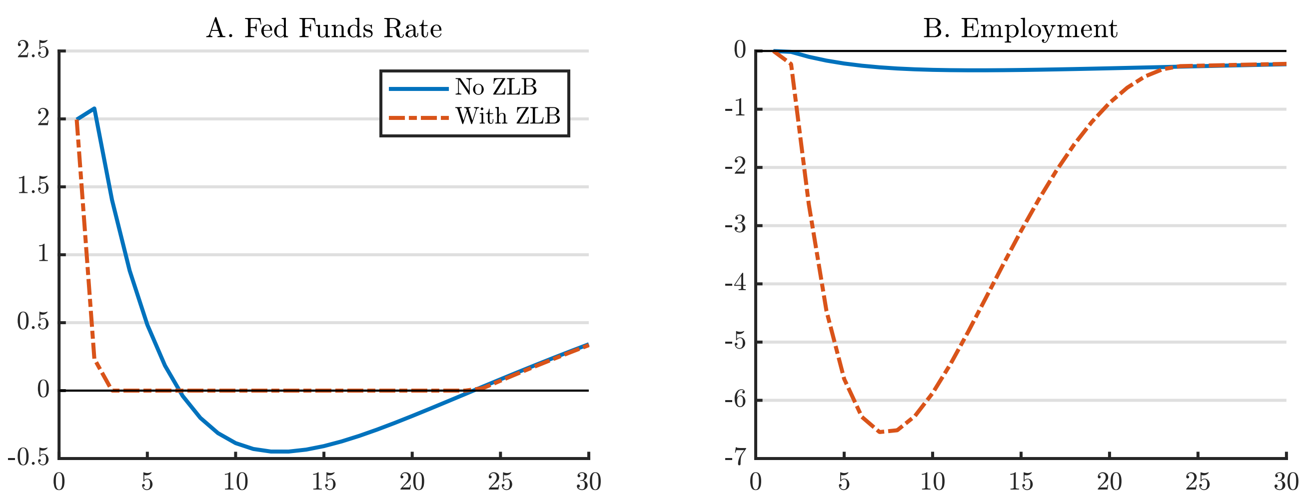

Figure 9: Effect of Credit Shocks on State Employment and Consumption

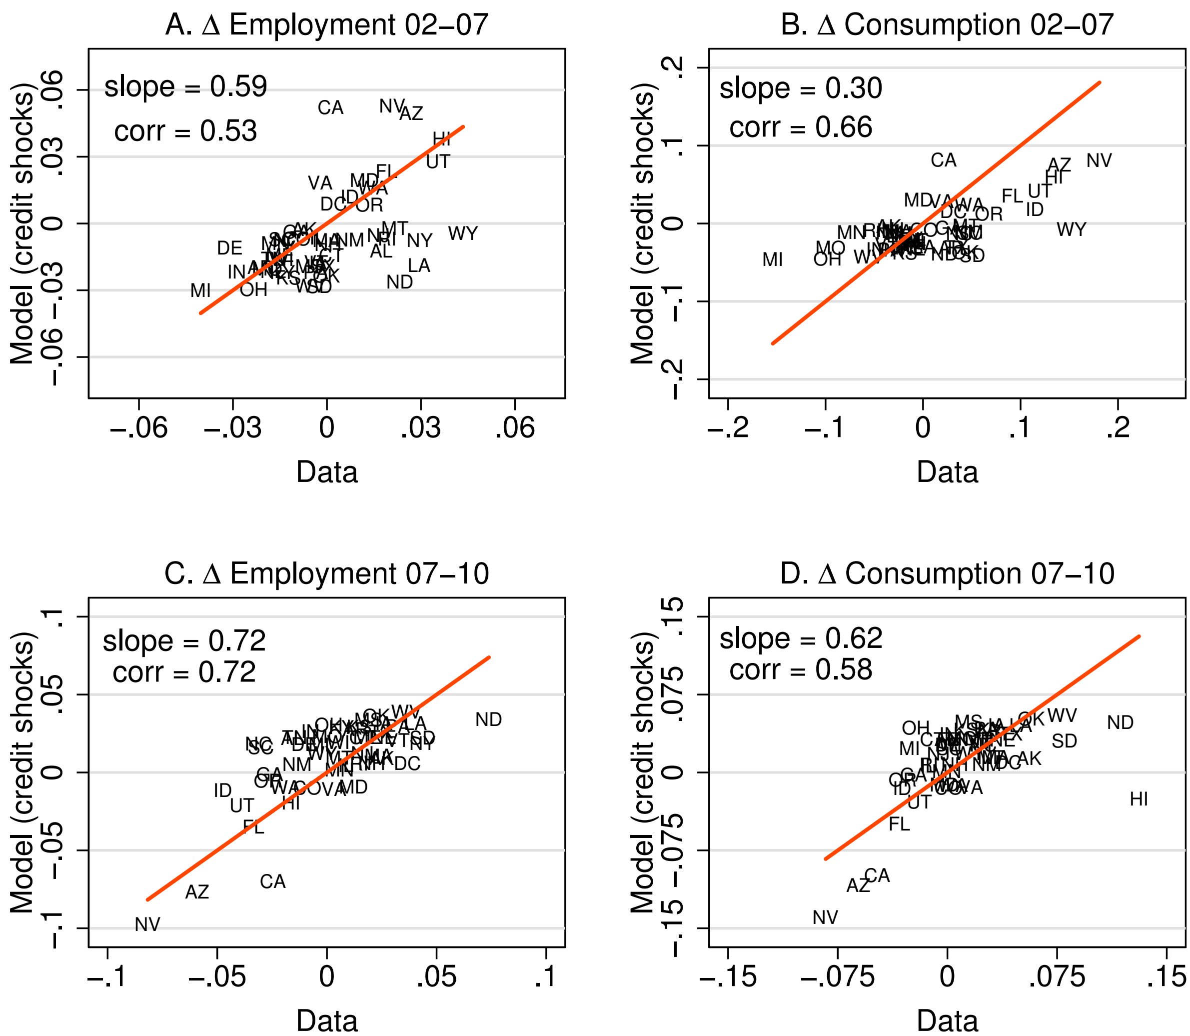

Figure 10: Effect of Credit Shocks Only on Aggregate Variables

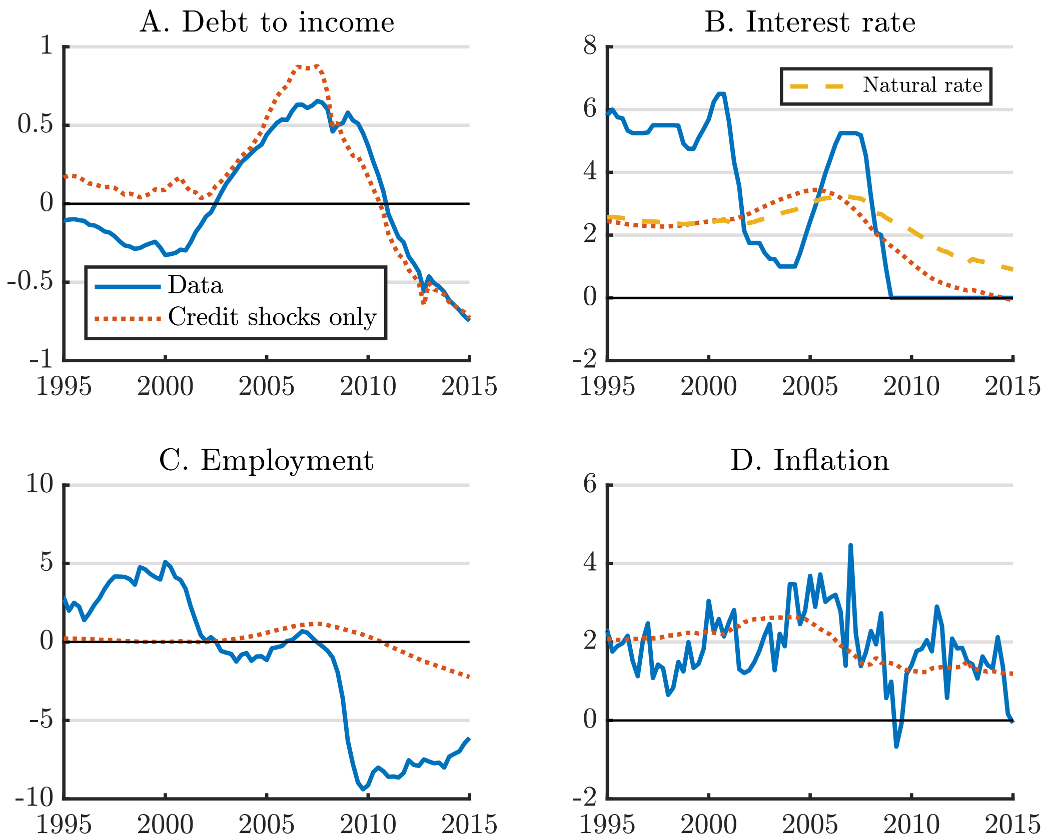

Figure 11: Effect of Credit Shocks on Aggregate Employment

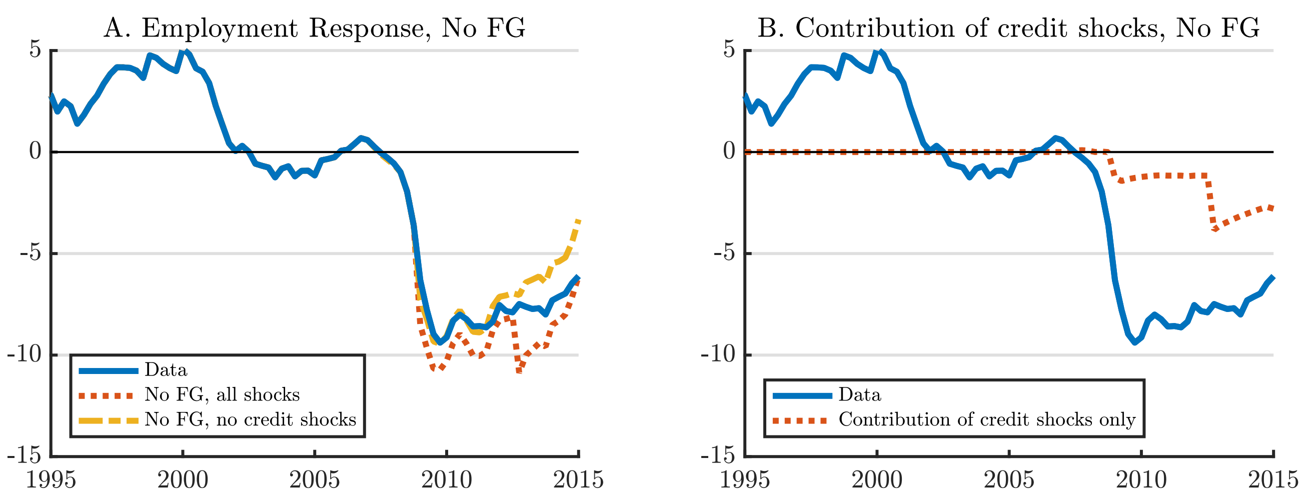

Figure 12: Effect of Credit Shocks on State Employment: Robustness

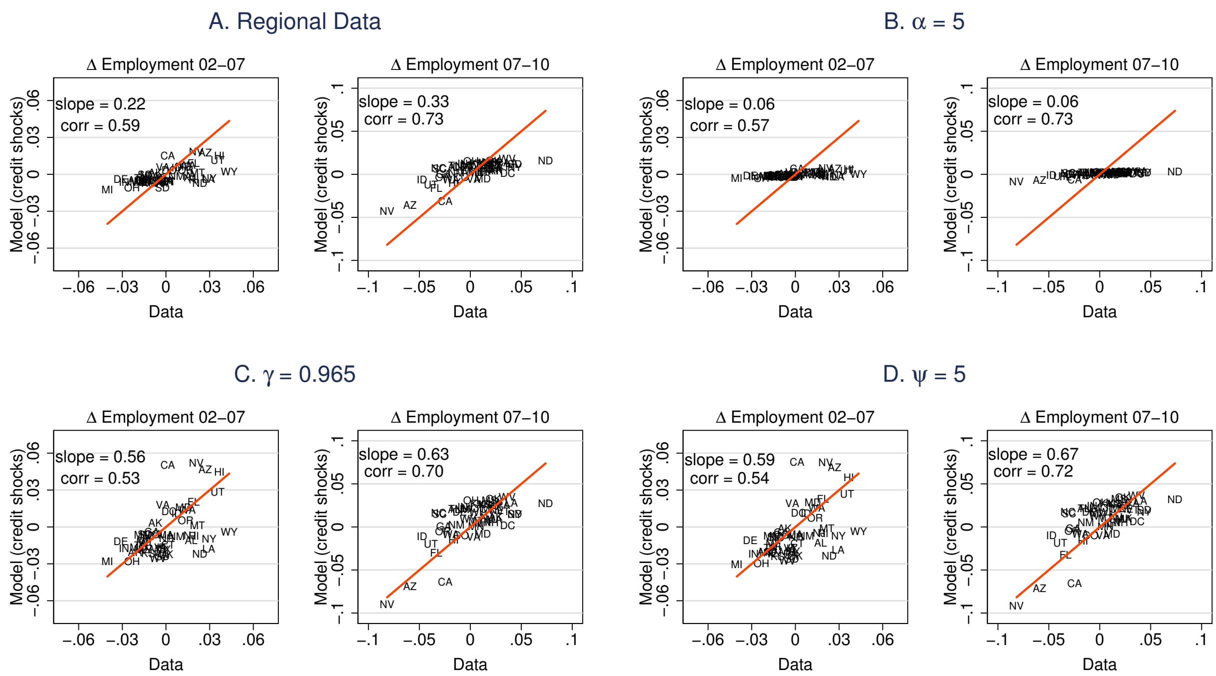

Figure 13: Effect of Credit Shocks on Aggregate Employment: Robustness