Household Leverage and the Recession[*]

Callum Jones[†]   Virgiliu Midrigan[‡]   Thomas Philippon[§]

October 2021

Abstract

We evaluate and partially challenge the household leverage view of the Great Recession. In the data, employment and consumption declined more in U.S. states where household debt declined more. We study a model of a monetary union composed of many regions in which liquidity constraints shape the response of employment and consumption to changes in debt. We estimate the model with Bayesian methods combining state and aggregate data. Changes in household credit explain 40% of the differential rise and fall of employment across states, but a small fraction of the aggregate employment decline in 2007-2010. Nevertheless, since household deleveraging was gradual, credit shocks greatly slowed the recovery.

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

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

[†]Federal Reserve Board, callum.j.jones@frb.gov
[‡]New York University, virgiliu.midrigan@nyu.edu
[§]New York University, tphilipp@stern.nyu.edu

1 Introduction

A salient feature of the Great Recession is that U.S. regions with the largest declines in household debt had the largest declines in employment and consumption. Figure 1 illustrates this pattern, originally documented in a series of papers by Mian and Sufi.[1] One interpretation of this evidence that has received a lot of attention is the *household leverage* view of the recession. According to this view, declines in household debt forced households to reduce consumption and, because of price rigidities and trade frictions, caused a reduction in employment.

Our goal in this paper is to quantitatively evaluate the household leverage view of the recession. We ask: how important were shocks to household credit in generating fluctuations in employment and consumption across regions and in the aggregate? The evidence in Figure 1 does not directly answer this question for at least two reasons. First, a variety of shocks, in addition to credit shocks, affect the economy. Second, the response of employment or consumption in an individual state to credit shocks may be very different than the response in the aggregate. Extrapolating from the cross-sectional to the aggregate elasticity requires a careful consideration of trade patterns and monetary policy.

We therefore use a model to isolate the effects of credit shocks and answer our question. We study a tractable model of a monetary union composed of many regions that trade. Households face liquidity constraints that restrict their ability to use the equity in their homes to finance consumption. These constraints generate a time-varying spread between the equilibrium interest rate and the households' rate of time preference. In our baseline model we assume that credit shocks affect loan-to-value constraints which limit the amount of housing equity that households can extract. In an extension we obtain similar results by assuming instead that household debt fluctuates in response to credit supply shocks. In addition to credit shocks, households are subject to additional disturbances to productivity and preferences.

We estimate the model using Bayesian full-information methods and data on employment, consumption, wages, house prices and household debt at the regional and aggregate levels. An important challenge that arises in our estimation is that the policy interest rate was at the zero lower bound (ZLB) during part of our sample. We account for the non-linearities arising from this bound by using a piece-wise

---

[1]See Mian and Sufi (2011); Mian et al. (2013); Mian and Sufi (2014) for additional evidence.

Figure 1: State Debt, Employment, and Spending
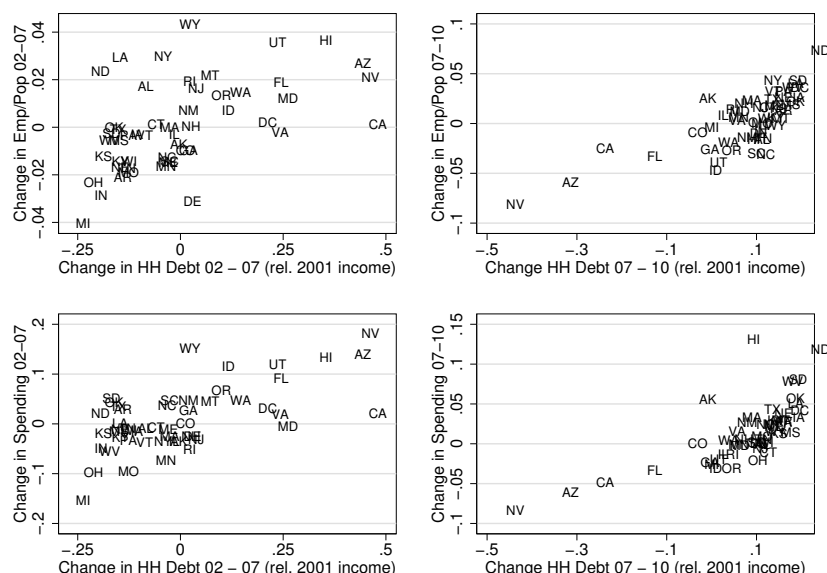
linear solution method.[2] Even with this approximation, a direct application of full-information methods is infeasible in our setting given the large number of regional and aggregate state variables. We therefore develop a methodology that allows us to combine regional and aggregate data to evaluate the likelihood function.

A popular approach to modeling household debt is to assume heterogeneity in agents' rate of time preference, as in Iacoviello (2005) and a large literature that follows.[3] These models feature a single financial asset, which impatient households use to borrow from patient households. A tightening of credit limits reduces the amount that impatient households borrow and translates into a dollar-for-dollar drop in their consumption, but, in the absence of changes in the equilibrium interest rate, leaves patient households unaffected. Such models thus predict that consumption in an individual region changes dollar-for-dollar with changes in household debt, regardless of how many agents are impatient, as long as the region is sufficiently small so it does not affect the economy-wide interest rate. One-asset models therefore imply counterfactually large comovements between consumption and credit across regions, at odds with the evidence in Figure 1.[4]

---

[2]See Eggertsson and Woodford (2003); Guerrieri and Iacoviello (2015); Jones (2017).

[3]See Eggertsson and Krugman (2012); Justiniano et al. (2015); Guerrieri and Iacoviello (2017) for the U.S., and Martin and Philippon (2017); Gourinchas et al. (2016) for the Eurozone.

[4]If there are few borrowers in a region, the percentage change in consumption can be much lower

Motivated by this observation, we pursue an alternative approach. We assume that households are subject to idiosyncratic shocks to their liquidity needs which make it optimal to simultaneously hold liquid assets and borrow. A household can therefore respond to a tightening of credit by drawing on its liquid assets instead of cutting consumption. The extent to which it can do so depends on the volatility of idiosyncratic shocks. If these shocks are volatile, the household maintains its liquid asset position, thus reducing consumption sharply after a credit tightening. If, in contrast, the volatility of idiosyncratic shocks is low, households find it optimal to dip into their liquid assets to insulate consumption from credit shocks. Our model thus flexibly nests both one-asset models in which a region-specific tightening of credit leads to a large drop in consumption, as well as the frictionless model in which credit shocks have no effect on consumption. We view this framework as our main theoretical contribution: it is flexible enough to capture complex credit dynamics but simple enough to be estimated on a full panel of regional economic data.

We next explain why households in our model value liquidity. The typical approach to modeling the notion that housing equity is illiquid is to assume that agents must pay a fixed cost to refinance their mortgage, as in the work of Kaplan et al. (2020) and Boar et al. (2020). Since most homeowners who do not refinance are constrained,[5] they save in a liquid asset in order to smooth consumption. By contrast, considerations of computational tractability lead us to follow an alternative approach inspired by Lucas (1990). Agents in our model can costlessly rebalance their portfolios between periods, but they must do so prior to the realization of an idiosyncratic preference shock. We achieve tractability by assuming a family construct to eliminate the distributional consequences of asset market incompleteness.[6] These assumptions allow us to parsimoniously model an interest-elastic supply of liquid assets of the type that would arise in richer but less tractable models of mortgage refinancing. In particular, the larger the interest rate is relative to the rate of time preference, the more households save in their liquid accounts to smooth the impact of idiosyncratic shocks. Our notion of liquidity therefore closely resembles that in the work of Shi

---

than the percentage change in debt. Notice, however, that in Figure 1 we scale debt by the 2001 income in a region. Since income and consumption are of a similar magnitude, one-asset models would imply much larger percentage changes in consumption compared to the data.

[5]See Amromin et al. (2007) and Adelman et al. (2010) for evidence on this.

[6]See also Lucas and Stokey (2011) who emphasize the role of liquidity frictions and Challe and Ragot (2016) who use a family construct to characterize an economy with uninsurable risks.

(2015), Rocheteau et al. (2018) and Kiyotaki and Moore (2019) in which liquidity is valued because of idiosyncratic investment opportunities or preference shocks.

Consider now the macroeconomic implications of a credit shock. At the economy-wide level, a tightening of credit leads to a reduction in the natural interest rate. The drop in the natural rate depends on the volatility of idiosyncratic preference shocks: more volatile shocks reduce the elasticity of liquid assets and imply a larger fall in the natural rate. If prices are sticky and monetary policy is unable to track the natural rate, the credit shock leads to a decline in employment. The response to economy-wide credit shocks is therefore highly non-linear depending on how long the ZLB is expected to constrain monetary policy. At the regional level, a tightening of credit acts like a sudden stop by forcing an increase in the net foreign asset position, as in Gourinchas et al. (2016), but the extent of the increase in net foreign assets depends, once again, on the volatility of idiosyncratic shocks. Hence, the same parameter governing the elasticity of liquid assets determines both the economy-wide and regional implications of a tightening of credit, a feature that we exploit in estimation.

We estimate three structural parameters that play a key role in shaping the economy's response to a credit shock: the degrees of wage and price stickiness, and the degree of idiosyncratic uncertainty. We also estimate the persistence and volatility of the regional and aggregate shock processes. To make full-information estimation feasible we exploit the structure of the model and express regional variables as deviations from the corresponding aggregates. Up to a first-order approximation, these deviations evolve independent of monetary policy. This structure allows us to separate the likelihood function into independent state-level and aggregate components. 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.

We use the model to generate counterfactual series for employment and consumption by turning off all shocks other than credit shocks. We find that credit shocks alone account for 40% to 50% of the relative movements in state-level employment and consumption during the 2002 to 2010 period. Our findings are thus consistent with those of Mian and Sufi (2011, 2014).

We next discuss the model's economy-wide implications. We distinguish between two periods: the onset of the Great Recession from 2007 to 2010, and the recovery period up to the end of 2012. Given the non-linear nature of the responses at the

ZLB, we provide bounds on the impact of credit shocks. The lower bound assumes no ZLB. The upper bound assumes that the ZLB binds and that the Fed does not pursue forward guidance. Under both bounds we find that household credit shocks have a relatively small impact on employment between 2007 and 2010. In the data, employment fell by 7% during this period. In contrast, the lower and upper bounds on the contribution of credit shocks to the decline in employment are 0.5% and 1.7%, respectively. Thus, even under the stark assumption that the Fed did not react to credit shocks at all, household deleveraging explains at most a quarter of the drop in employment during the acute phase of the crisis. This result is a consequence of the gradual nature of household deleveraging which leads to a gradual decline in the natural rate and reflects the findings of Kaplan et al. (2020) who show that the duration of mortgages greatly matters for the effects of shocks to household credit.[7]

The results are different when we consider the period 2007 to 2012. Since household debt fell more over this longer period, so did the natural rate, and the model predicts a more important role for credit shocks. In the data, employment was 6.2% lower by the end of 2012 compared to 2007. Under the two bounds, employment declined by 1.3% and 4.8%, respectively. The model therefore predicts that credit shocks greatly delayed the recovery.[8] As we show in the Online Appendix, our results are robust to a number of perturbations, including adding a construction sector, allowing for fiscal policy, lowering the duration of mortgage contracts, explicitly modeling mortgage default, credit supply shocks and heterogeneity in state responses to credit shocks.

Our goal in this paper is to quantify the role of credit shocks in accounting for regional and aggregate employment during the Great Recession. To keep our analysis as transparent as possible, we focus on the *household leverage* view of the Great Recession and do not explicitly model the other forces that arguably played an important role, such as bank and firm-level financing constraints, changes in uncertainty, demographics, increase in markups and risk premia. These factors are instead captured by our rich set of exogenous shocks. A limitation of our approach is that household credit may interact with a number of these forces in important ways.[9]

Our analysis also focuses on a representative agent economy, which, in general

---

[7]See also Kiyotaki et al. (2011).

[8]See the literature on the slow recovery discussed by Fernald et al. (2017).

[9]For example, a tightening of household credit that depresses income and home prices may trigger a wave of default that reduces the net worth of financial intermediaries. See Faria-e-Castro (2018).

equilibrium, precludes us from capturing housing wealth effects on consumption. Introducing housing wealth effects, which Kaplan et al. (2020) argue were responsible for a large fraction of the expenditure drop during the Great Recession, requires explicitly modelling household heterogeneity and the resulting correlation between wealth effects and marginal propensities to consume across households. To keep our analysis tractable, we abstract from such heterogeneity and thus implicitly capture these forces using changes in the household's rate of time-preference.

Our paper is related to Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2017) who also study the responses of an economy to a household credit crunch. While they study a closed-economy, our model is that of a monetary union composed of a large number of regions. Moreover, our focus is on estimating the model using state and aggregate data. Since we found that one-asset models cannot account for the comovement between state-level consumption and household debt, we explicitly introduce liquidity constraints which make it optimal for households to simultaneously borrow and save in the liquid asset, an ingredient which we argue is critical to match the regional data.

Our use of regional data in estimation is related to Beraja et al. (2019) who apply a limited-information approach to estimate the degree of wage stickiness using state-level data, and then use those estimates as a prior in a full-information estimation with aggregate data. In contrast our methodology simultaneously uses regional and aggregate data in a full-information approach, despite the large number of state variables and non-linearities caused by the ZLB. Our emphasis on cross-sectional evidence is also shared by the work of Nakamura and Steinsson (2014). Finally, our work is related to the literature on financial intermediation, originating with Bernanke and Gertler (1989); Kiyotaki and Moore (1997); Bernanke et al. (1999) and more recently Mendoza (2010); Gertler and Karadi (2011); Gilchrist and Zakrajšek (2012).[10]

2   Model

We first describe the model economy. The economy consists of a continuum of ex-ante identical islands of unit mass that belong to a monetary union and trade among themselves. Consumers derive utility from a final good, leisure and housing. The

---

[10]See also the work Lustig and Van Nieuwerburgh (2005); Garriga et al. (2019); Favilukis et al. (2017); Burnside et al. (2016); Landvoigt et al. (2015) who study the determinants of house prices.

final good is assembled using as inputs traded and non-traded goods. We assume that producers of intermediate goods are monopolistically competitive. 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.

We explicitly model the distinction between liquid and illiquid assets on households' balance sheets. In contrast to Kaplan and Violante (2014) and much of the segmented asset markets literature which assumes fixed costs of portfolio rebalancing, we formulate an alternative approach following Lucas (1990). We assume that consumers allocate their wealth between liquid and illiquid assets prior to the realization of an idiosyncratic preference shock. The more volatile this shock is, the stronger the precautionary savings motive in liquid assets, and the larger the response of consumption to credit shocks. Following Lucas (1990), we assume that consumers belong to families that share idiosyncratic risks.

Households face shocks to their ability to tap home equity, which we refer to as credit shocks. We also allow for shocks to the households' rate of time preference, disutility from work, preference for housing and productivity. Each shock has an island-specific and aggregate component. We also assume aggregate shocks to the interest rate rule and to the aggregate inflation equation. These additional shocks capture other channels that can explain movements in macroeconomic aggregates.

2.1 Households

Securities. Households borrow using mortgages which are long-term perpetuities 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.[11] The household borrows from perfectly competitive financial intermediaries.

It is convenient to describe the outstanding mortgage debt of a household on island s by recording the amount of coupon payments $b_t(s)$ it must make in period t. Letting $l_t(s)$ denote the amount of newly issued mortgages in period t, the date $t+1$ coupon payments are

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

---

[11]See Hatchondo and Martinez (2009) and Arellano and Ramanarayanan (2012).

Letting q_t denote the economy-wide price of one claim to such a perpetuity, the value of a household's mortgage liabilities is $b_t(s)(1 + \gamma q_t)$, reflecting the amount the household owes in coupon payments in period t and the market value of its outstanding debt. On the asset side, we assume that households save in a one-period nominal security at a rate i_t.[12]

Preferences. We introduce wage-setting frictions by assuming that households are organized in unions that set wages. We assume perfect risk-sharing across households that belong to different labor unions on a given island. Because of separability in preferences, risk-sharing implies that all households on an island make identical consumption, housing and savings choices, even though their labor supply differs depending on when the union they belong to last reset its wage. The preferences of a household on island s who belongs to a labor union ι are

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

where $c_{it}(\iota, s)$ denotes the consumption of a member i of household (ι, s), $h_t(\iota, s)$ denotes the total amount of housing owned by that household and $n_t(\iota, s)$ denote the labor it supplies. The shifters $\eta_t^h(s)$ and $\eta_t^n(s)$ determine the preference for housing and the disutility from work, while $\beta_t(s)$ is the household's one-period ahead discount factor. Each of these preference shifters have an island-specific and aggregate component, all of which follow AR(1) processes with independent Gaussian innovations.

The term $v_{it}(\iota, s) \geq 1$ is specific to each consumer i and denotes an idiosyncratic taste shifter, which is an i.i.d random variable drawn from a Pareto distribution

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

The parameter $\alpha > 1$ determines the amount of uncertainty about v. A lower α implies more uncertainty. Below we drop the dependence on ι for notational simplicity.

Budget and Credit Constraints. Let $x_t(s)$ denote the total amount of funds the household transfers to its members. We assume that $x_t(s)$ is chosen prior to the

---

[12]Since we solve the model using first-order perturbation methods, the assumption that assets have a different maturity structure than liabilities is inconsequential.

realization of the idiosyncratic preference shock $v_{it}(s)$. Since individual members are ex-ante identical and of mass 1, $x_t(s)$ also represents the amount of funds received by any individual member. We assume a *liquidity constraint* which limits each member's consumption by the amount of funds it has when entering the goods market:

$$p_t(s)c_{it}(s) \leq x_t(s), \tag{1}$$

where $p_t(s)$ is the price of the final good on island s. The budget constraint is

$$x_t(s)+e_t(s)(h_{t+1}(s)-h_t(s)) = w_t(s)n_t(s)+q_tl_t(s)-b_t(s)+(1+i_{t-1})a_t(s)+T_t(s), \tag{2}$$

where $e_t(s)$ is the price of housing, $w_t(s)$ is the wage rate the household faces, $b_t(s)$ are the coupon payments on outstanding mortgage debt, $a_t(s)$ are the liquid assets it enters the period with and $T_t(s)$ 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 island.

The household also faces a *borrowing constraint* which limits the size of *new* loans:

$$q_tl_t(s) \leq m_t(s)e_t(s)h_{t+1}(s), \tag{3}$$

where the credit limit $m_t(s)$ evolves as the product of an island-specific and aggregate component, both of which follow an AR(1) process. This formulation allows us to capture the slow movement of debt in the data in which only a fraction of households obtain new mortgages in any given period, and so face the current credit conditions and house prices in determining mortgage limits.

We introduce housing preference shocks to capture movements in house prices that are otherwise difficult to rationalize.[13] Because housing is separable in the utility function and the housing stock is in fixed supply, movements in house prices only affect equilibrium allocations via their impact on the households' ability to borrow.[14]

---

[13]See Kiyotaki et al. (2011) for an illustration of the problem and Favilukis et al. (2017), Garriga et al. (2019), and Kaplan et al. (2020) for approaches to resolve it. We think of preference shocks as a parsimonious way of capturing changes in risk premia or beliefs that trigger changes in house prices in these richer models.

[14]In contrast, in partial equilibrium house price changes would have wealth effects on household

Figure 2: Timing of the Model
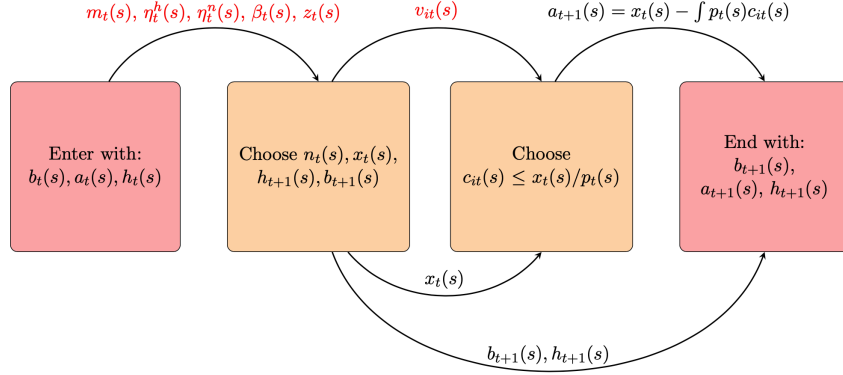
Specifically, a housing preference shock that increases house prices loosens the borrowing constraint (3) and leads to an increase in consumption.

Savings. A household's savings are the unspent funds in the goods market. The total amount of liquid assets the household has at the end of the shopping period is

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

These funds are deposited with financial intermediaries and earn the interest rate i_t.

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

Timing. Figure 2 summarizes our timing assumptions. The household's portfolio composition is described by liquid assets $a_t(s)$, housing $h_t(s)$ and mortgage debt $b_t(s)$. All shocks except $v_{it}(s)$ are realized at the beginning of the period. The household then chooses labor $n_t(s)$, housing $h_{t+1}(s)$, borrowing $b_{t+1}(s)$, and transfers $x_t(s)$. The idiosyncratic preference shocks $v_{it}(s)$ are then realized and individual members choose consumption $c_{it}(s)$. All unspent funds $a_{t+1}(s)$ are then saved in the liquid asset.

consumption, as shown in Berger et al. (2018).

10

2.2 Optimal Savings Choice

Let $\lambda_t(s)$ denote the multiplier on the household's budget constraint (2) and $\xi_{it}(s)$ denote the multiplier on the liquidity constraint (1). The choice of $x_t(s)$ satisfies

$$\lambda_t(s) = \beta_t(s)(1 + i_t)\mathbb{E}_t\lambda_{t+1}(s) + \int_0^1 \xi_{it}(s)\mathrm{d}i. \tag{4}$$

The transfer $x_t(s)$ is valued at $\lambda_t(s)$, the shadow value of wealth. Since unspent funds can be saved at the conclusion of the shopping period, they are valued at $\beta_t(s)(1+i_t)\mathbb{E}_t\lambda_{t+1}(s)$. The last term captures the liquidity services provided by $x_t(s)$, namely the average multiplier on the liquidity constraint of individual members.

Optimal consumption spending is equal to

$$p_t(s)c_{it}(s) = \min\left[\frac{v_{it}(s)}{\beta_t(s)(1 + i_t)\mathbb{E}_t\lambda_{t+1}(s)}, \quad x_t(s)\right]$$

and the multiplier on the liquidity constraint is

$$\xi_{it}(s) = \max\left[0, \quad \frac{v_{it}(s)}{x_t(s)} - \beta_t(s)(1 + i_t)\mathbb{E}_t\lambda_{t+1}(s)\right]. \tag{5}$$

Integrating (5) across all members of a given household implies that

$$\int_0^1 \xi_{it}(s)\,\mathrm{d}i = \frac{1}{\alpha - 1}x_t(s)^{-\alpha}\left[\beta_t(s)(1 + i_t)\mathbb{E}_t\lambda_{t+1}(s)\right]^{1-\alpha}. \tag{6}$$

Let

$$\Delta_t(s) \equiv \frac{\lambda_t(s)}{(1 + i_t)\beta_t(s)\mathbb{E}_t\lambda_{t+1}(s)} - 1$$

denote the wedge between the household's rate of time preference and the return on the liquid asset. Combining (4) and (6) allows us to express transfers as:

$$x_t(s) = \frac{1}{\beta_t(s)(1 + i_t)\mathbb{E}_t\lambda_{t+1}(s)}\left[(\alpha - 1)\Delta_t(s)\right]^{-\frac{1}{\alpha}}.$$

Since liquid savings are equal to the amount of unspent funds in the goods market, $a_{t+1}(s) = x_t(s) - p_t(s)\int c_t(s)\mathrm{d}i$, they can be expressed as:

$$\frac{a_{t+1}(s)}{p_t(s)c_t(s)} = \left(\frac{\alpha}{\alpha - 1}\left[(\alpha - 1)\Delta_t(s)\right]^{\frac{1}{\alpha}} - \Delta_t(s)\right)^{-1} - 1. \tag{7}$$

11

To understand this expression, note that when $\Delta_t(s)$ is high, the household is relatively impatient. Idiosyncratic taste shocks generate a precautionary savings motive, in addition to the intertemporal substitution motive. When uncertainty about taste shocks is high (α is low), the precautionary motive dominates so savings are relatively insensitive to the wedge $\Delta_t(s)$. In contrast, when the uncertainty about taste shocks is low (α is high), the precautionary motive is weak and savings are sensitive to $\Delta_t(s)$. As $\alpha \to \infty$, the supply of liquid assets is infinitely elastic.

The assumption that taste shocks are Pareto distributed, or that idiosyncratic uncertainty takes the form of preference, as opposed to income or expense shocks, is not critical to our arguments. In the Online Appendix we consider alternative types of shocks. Although these alternatives are not as tractable as our baseline formulation, their qualitative implications are similar.

2.3   Optimal Debt and Housing Choice

We next discuss the household's demand for mortgage debt and housing. Let

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

denote the one-period return on mortgages. The first-order condition for $b_{t+1}(s)$ is

$$\lambda_t(s) = \beta_t(s)\mathbb{E}_t\lambda_{t+1}(s)R_{t+1} + \mu_t(s) - \gamma\beta_t(s)\mathbb{E}_t\mu_{t+1}(s)\frac{q_{t+1}}{q_t}, \tag{8}$$

where $\mu_t(s)$ is the multiplier on the borrowing constraint (3). The first-order condition for housing is

$$e_t(s)\lambda_t(s) - \beta_t(s)\mathbb{E}_t\lambda_{t+1}(s)e_{t+1}(s) = \beta_t(s)\mathbb{E}_t\frac{\eta_{t+1}^h(s)}{h_{t+1}(s)} + \mu_t(s)m_t(s)e_t(s). \tag{9}$$

The left hand side is the user cost of housing. The right hand side adds the marginal utility of housing services to the collateral value of housing.

2.4   Technology

We next discuss the assumptions on technology we make.

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 $y_t^M(s)$ units of imported goods. Imported and non-tradable goods are combined using a CES aggregator with elasticity of substitution σ,

$$y_t(s) = \left(\omega^{\frac{1}{\sigma}} y_t^N(s)^{\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} y_t^M(s)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where ω determines the share of non-traded goods. The imported good is a CES composite of imports from all islands with elasticity of substitution κ:

$$y_t^M(s) = \left(\int_0^1 y_t^M(s,s')^{\frac{\kappa-1}{\kappa}} \mathrm{d}s' \right)^{\frac{\kappa}{\kappa-1}}.$$

We assume that there are no trade costs, so all islands face the same import price

$$p_t^T = \left(\int_0^1 p_t^T(s')^{1-\kappa} \mathrm{d}s' \right)^{\frac{1}{1-\kappa}},$$

where $p_t^T(s)$ is the price of imports from island s. Letting $p_t^N(s)$ denote the price of non-tradable goods, the final goods price on an island is

$$p_t(s) = \left(\omega p_t^N(s)^{1-\sigma} + (1-\omega) \left(p_t^T \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

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

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

while demand for its exports is the sum of purchases from all islands:

$$y_t^X(s) = (1-\omega) \left(\frac{p_t^T(s)}{p_t^T} \right)^{-\kappa} \left(\int_0^1 \left(\frac{p_t^T}{p_t(s')} \right)^{-\sigma} y_t(s') \mathrm{d}s' \right).$$

Intermediate Goods Producers. Tradable and non-tradable goods are themselves CES composites of varieties of differentiated intermediate inputs with an elasticity of substitution ϑ. Technology is linear in labor and subject to a productivity

13

disturbance $z_t(s)$ that is common to both the tradable and non-tradable sectors:

$$y_{kt}^j(s) = z_t(s)n_{kt}^j(s), \text{ for } j \in \{X, N\}.$$

Producers of intermediate goods are subject to Calvo pricing frictions and can only reset prices with probability $1 - \theta_p$ each period. A firm that resets its price maximizes the present discounted flow of profits weighted by the probability that the price it chooses at t will still be in effect at any particular date. The government levies a production subsidy $\tau_p = \frac{\vartheta - 1}{\vartheta}$ which eliminates the steady state markup distortion and is financed with lump-sum taxes.

Wage Setting. Households are organized in unions that supply differentiated varieties of labor with elasticity of substitution ψ. Unions reset their wages with probability $1 - \theta_w$ each period. We assume that the government offsets the steady-state wage markup distortion using a subsidy $\tau_w = \frac{\psi - 1}{\psi}$ financed with lump-sum transfers.

2.5   Monetary Policy

Let $y_t = \int_0^1 p_t(s)y_t(s)/p_t \, ds$ be total output in this economy, $p_t = \int_0^1 p_t(s)ds$ be the aggregate price level and $\pi_t = p_t/p_{t-1}$ denote the rate of inflation. Aggregating the prices of individual producers in all islands implies, up to a first-order approximation,

$$\log(\pi_t/\bar{\pi}) = \bar{\beta}\mathbb{E}_t \log(\pi_{t+1}/\bar{\pi}) + \frac{(1 - \theta_p)(1 - \theta_p\bar{\beta})}{\theta_p} \left(\log(w_t) - \log(z_t)\right) + \eta_t^{\pi},$$

where $w_t = \int w_t(s)ds$ is the average wage in the economy, η_t^{π} is an AR(1) disturbance to individual firms' desired markups, common to all islands, $\bar{\beta}$ is the steady state discount factor and $\bar{\pi}$ is the steady-state rate of inflation.[15]

We follow Smets and Wouters (2007) and assume that monetary policy follows a Taylor rule when the ZLB does not bind:

$$1 + i_t = (1 + i_{t-1})^{\alpha_r} \left[(1 + \bar{\imath})(\pi_t/\bar{\pi})^{\alpha_\pi} \hat{y}_t^{\alpha_y}\right]^{1-\alpha_r} (\hat{y}_t/\hat{y}_{t-1})^{\alpha_x} \eta_t^i,$$

---

[15]We assume in our quantitative analysis that $\bar{\pi}$ is equal to 2% per year. We implicitly eliminate the steady-state costs of inflation by assuming that prices and wages are automatically indexed to $\bar{\pi}$. See Coibion and Gorodnichenko (2015) and Blanco (2019) who study the size of these costs in the absence of indexation.

where $\bar{\imath}$ is the steady state nominal interest rate, η_t^i is an i.i.d. 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 deviations of inflation from the target $\bar{\pi}$, the output gap $\hat{y}_t$, and the growth rate of the output gap, respectively. The output gap is the ratio of output to its flexible-price level. When the ZLB binds, $i_t = 0$.

The Fed may set an interest rate of zero not only when the ZLB binds, but also when it follows its *forward guidance*. In particular, we assume that in periods in which the ZLB binds the Fed can commit to keeping i_t at zero for longer than the duration prescribed by the Taylor rule. We thus implicitly assume that the Fed can manipulate expectations of how the path of interest rates evolves when the economy hits the ZLB, as in Eggertsson and Woodford (2003) and Werning (2011). In our estimation we use survey data from the New York Federal Reserve to discipline the expected duration of the zero interest rate regime between 2009 and 2015.

2.6 Asset Market

The monetary union is closed so savings in the liquid asset equals aggregate debt:

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

In contrast, individual states can run current account imbalances. An island's net foreign asset position evolves over time according to

$$a_{t+1}(s) - q_t b_{t+1}(s) = (1 + i_{t-1})a_t(s) - (1 + \gamma q_t)b_t(s) + p_t^T(s)y_t^X(s) - p_t^T y_t^M(s).$$

Perfectly competitive risk-neutral financial intermediaries borrow liquid assets from households and lend in the mortgage market. Perfect competition implies that the expected return to holding mortgages equals the interest rate on the liquid asset:

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

We next discuss the role of liquidity constraints in determining the steady-state equilibrium interest rate. Recall that the household's liquid savings are given by (7). In steady state the wedge between the discount rate and the interest rate is equal to $\Delta = \frac{1}{\bar{\beta}} / \frac{1+i}{\bar{\pi}} - 1 \approx \rho - r$, where $\rho = 1/\bar{\beta} - 1$ is the discount rate and $r = \frac{1+i}{\bar{\pi}} - 1$ is the real

interest rate. The upward-sloping curves in Figure 3 illustrate the households' optimal savings choices as a function of the real interest rate. The two panels of the Figure correspond to two scenarios, one with a relatively high idiosyncratic uncertainty (low α), and another with a relatively low uncertainty (high α).

To derive the demand for assets, we note that the borrowing limit binds at all times. This is because taste shocks are unbounded, which implies that the multiplier on the liquidity constraint $\xi_{it}(s)$ is positive for a positive mass of household members. A comparison of (4) and (8), together with the no-arbitrage condition $\mathbb{E}_t R_{t+1} = 1 + i_t$, implies that the multiplier on the borrowing constraint $\mu_t(s)$ is also positive. Intuitively, since the household anticipates that a fraction of its members will end up liquidity constrained, and since the expected return on the mortgage is equal to the liquid interest rate, it borrows as much as possible. The result that all households are constrained, though stark, is consistent with Boar et al. (2020) who find that four-fifths of homeowners are debt-constrained.[16]

Since the credit limit binds, mortgage debt is proportional to the value of houses,

$$qb = \frac{\bar{m}}{1-\gamma}eh = \frac{\bar{m}}{1-\gamma}\frac{\bar{\eta}^h}{\lambda}\frac{1}{\left(1 - \frac{\bar{m}}{1-\beta\gamma}\right)\rho + \frac{\bar{m}}{1-\beta\gamma}r}, \tag{10}$$

where the last equation follows from the Euler equation for housing (9). The value of housing is given by the marginal utility of housing, $\bar{\eta}^h/\lambda$, discounted by a weighted average of the rate of time preference ρ and the interest rate r, with a weight that depends on how much the household can borrow. As long as $\rho > r$, an increase in the loan-to-value ratio $\bar{m}$ reduces the effective discount rate, raising house prices.

The downward sloping curve in Figure 3 super-imposes the demand for mortgages (10) and illustrates how the equilibrium interest rate is determined in steady state. A tightening of the debt limit mechanically reduces the demand for debt, thus reducing the equilibrium interest rate. The extent to which the equilibrium interest rate falls depends on the degree of idiosyncratic uncertainty. As the two panels of the figure show, the larger the uncertainty, the less elastic the supply of liquid assets, and therefore the larger the reduction in the equilibrium interest rate.

---

[16]In reality, the observation that households are constrained in how much they borrow may reflect higher returns on other assets, such as equity, and not necessarily liquidity constraints.

Figure 3: Equilibrium Real Interest Rate in Steady State
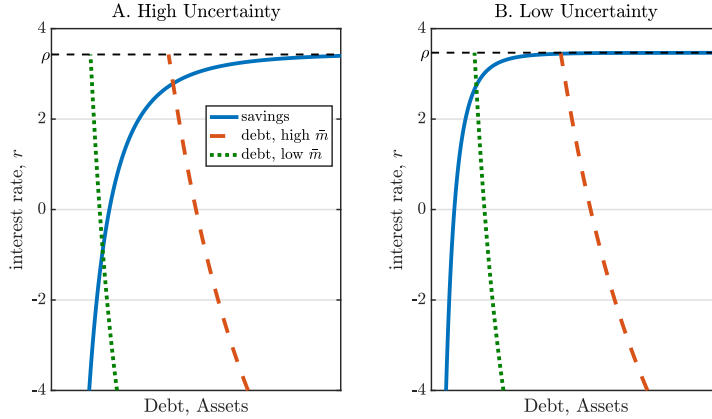
2.7 The Workings of the Model

We next provide some intuition for how the model works. The key parameters in the model are those that determine the degree of idiosyncratic uncertainty, α, and those governing the degree of price and wage stickiness, θ_p and θ_w. We illustrate the role of each of these sets of parameters by perturbing them and describing the impulse responses to shocks to the credit limit $m_t(s)$.

Impulse Response to a State-Level Credit Shock. Figure 4 reports impulse responses to a state-level credit shock. To build intuition, we first report, in the left column, the responses under flexible prices. Panel A of the figure shows that the shock leads to a gradual decline in debt because of the long-term nature of mortgages. Households cannot borrow as much as they used to, and they respond in three ways. They reduce their liquid asset holdings, which limits their ability to smooth idiosyncratic shocks. They also reduce consumption and increase labor supply. Prices and wages fall, employment rises, and the island runs a current account surplus.

The middle column of Figure 4 illustrates how price rigidities affect the response of an island to a credit shock. We refer to this parameterization as the "Baseline" because it relies on the parameters from our estimation below. Wage and price rigidities reverse the response of employment. Price stickiness limits the increase in exports so the increase in the current account takes place via a compression of domestic demand. The fall in employment reduces income and consumption further. Wage rigidities act as a tax on labor supply, while price rigidities increase firms'

17

Figure 4: Impulse Response to State-Level Credit Shock
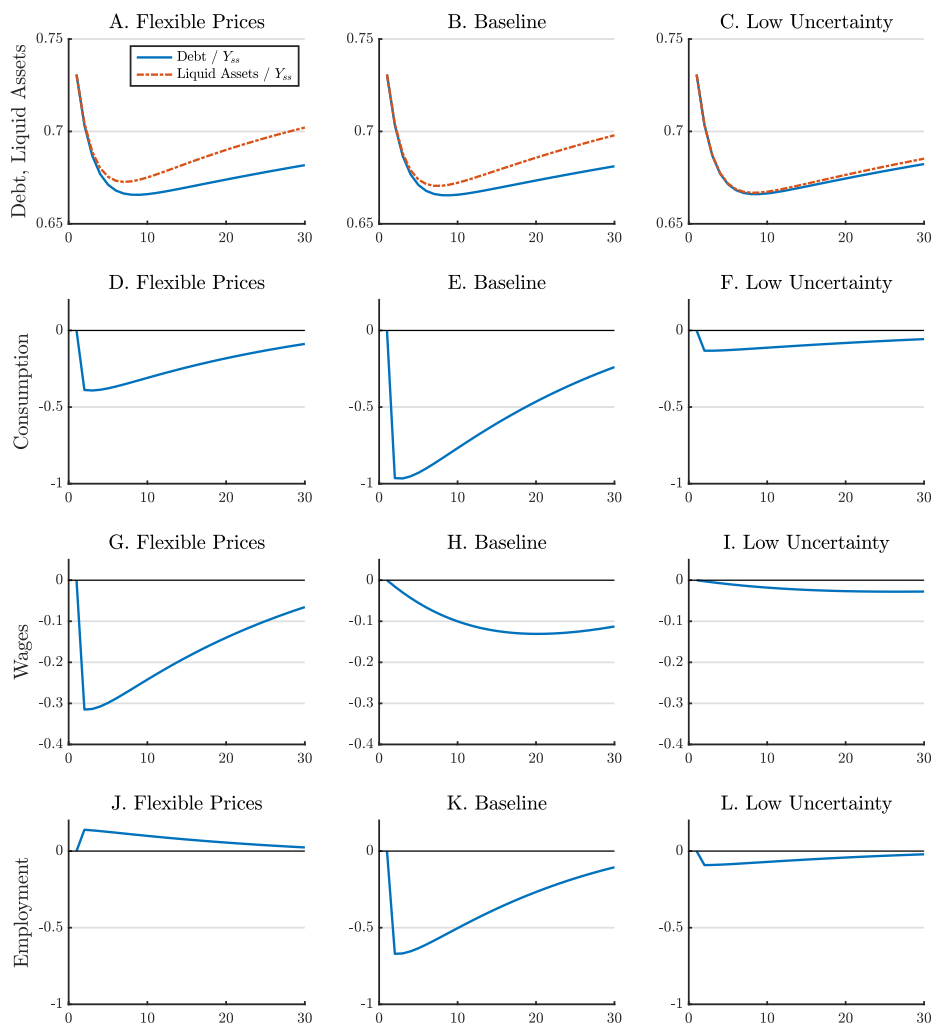
markups, and both effects reduce employment.[17]

The last column of Figure 4 illustrates how liquidity constraints and liquidity demand affect the response of an island to a credit shock. We reduce the volatility of idiosyncratic shocks by increasing the parameter α. With less need to smooth large liquidity risks, households are more willing to reduce their liquid asset holdings in response to a tightening of their borrowing limit. As a result, consumption barely falls and the impact on employment and wages is small. These effects are captured in equation (7) which shows that liquid asset holdings are more sensitive to the wedge between the discount rate and the interest rate when α is high. The discount

---

[17]See Kehoe et al. (2016) for evidence that both of these margins account for the drop in employment in states that have experience the largest declines in household credit.

rate $\frac{\lambda_t(s)}{\beta_t(s)\mathbb{E}_t\lambda_{t+1}(s)}$ increases in response to the credit tightening, while the interest rate is constant since the island is atomistic. When idiosyncratic uncertainty is low, it is relatively costless to reduce liquid asset holdings. Both sides of the household's balance sheet thus contract, with little impact on other variables. In this case credit is a veil with little impact on macroeconomic aggregates. In contrast, when idiosyncratic uncertainty is high, reducing liquid assets is costly. Households therefore find it optimal to respond to the credit tightening by cutting consumption.

A tightening of debt constraints thus distorts allocations in two ways: it prevents households from smoothing the marginal utility of consumption both *across members* and *across time*. Households face a tradeoff: they can respond to a tightening of credit by either reducing overall consumption, distorting the intertemporal allocations, or by reducing liquid assets, distorting the intrafamily allocations. The more dispersed idiosyncratic shocks are, the more the household reduces overall consumption to limit variation in the marginal utility of consumption across its members.

Impulse Response to an Aggregate Credit Shock. To understand the effect of an aggregate credit shock, we note first that the shadow value of wealth λ_t is approximately inversely proportional to nominal consumption spending, $\lambda_t \sim 1/p_t c_t$. Up to a first-order approximation, the Euler equation that determines the growth rate of λ_t is given by

$$(1 + \Delta_t)(1 + i_t)\beta\mathbb{E}_t\frac{\lambda_{t+1}}{\lambda_t} = 1.$$

where

$$\left(\frac{\alpha}{\alpha - 1} \left[(\alpha - 1)\Delta_t \right]^{\frac{1}{\alpha}} - \Delta_t \right)^{-1} - 1 = \frac{q_t b_{t+1}}{y_t} \tag{11}$$

follows from asset market clearing. A tightening of debt limits reduces the right-hand side of (11) and increases the wedge Δ_t, thus reducing the natural rate.

The first column of Figure 5 shows the response to an aggregate credit shock in an economy with flexible prices. The shock reduces the interest rate (Panel A), and has a negligible effect on employment (Panel D). Employment falls slightly because a tightening of credit magnifies consumption-leisure distortions, as in cash-in-advance economies where the impact of these distortions is small (Cooley and Hansen, 1989).

We next turn to our baseline parameterization with sticky prices. As is well understood, the extent to which changes in the natural rate affect macroeconomic aggregates depends on the stance of monetary policy. We illustrate this point in the

19

Figure 5: Impulse Response to an Aggregate Credit Shock
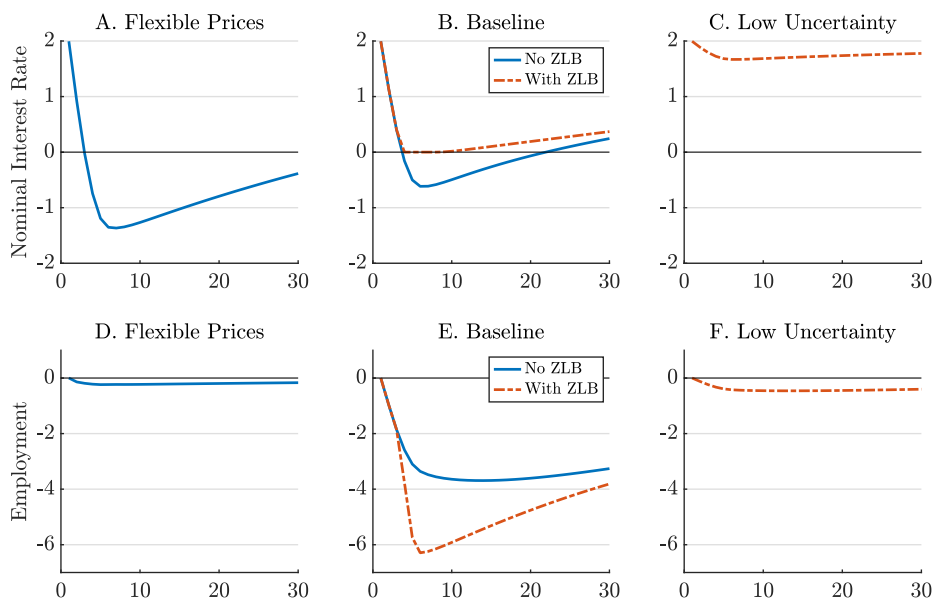
second column of Figure 5 by contrasting the responses to the same credit shock under two scenarios: one in which we impose the ZLB and another one in which we do not. Clearly, the ZLB magnifies the reduction in employment because it prevents the monetary authority from accommodating the decline in the natural rate.

As with the island-level responses, the volatility of idiosyncratic shocks is critical in determining how the economy reacts to a credit shock. As the last column of Figure 5 shows, if idiosyncratic uncertainty is low, interest rates and employment fall little. In this case liquid asset holdings are sensitive to interest rates, so a small reduction in the equilibrium rate is required to clear the asset market. The natural rate of interest thus falls little, leading to a negligible impact on macroeconomic outcomes.

This exercise shows that the parameters that shape the responses of island-level variables to credit shocks also determine the aggregate-level responses. These parameters govern the slopes of two curves: that of the Phillips curve that relates inflation to changes in marginal costs and that of the savings curve which relates liquid asset holdings to the gap between the discount rate and the interest rate. Since these are both functions of the primitive parameters, they are identical in the aggregate and at the state level. In contrast, the equilibrium responses of, say, employment to credit shocks are determined by factors such as a state's openness to trade and the stance of monetary policy and one cannot extrapolate state-level elasticities to draw conclusions about the aggregate. Indeed, if monetary policy perfectly tracked the natural

rate, it would eliminate aggregate employment fluctuations, yet would be unable to respond to disturbances on an individual island. We can therefore use state-level data to identify the parameters that determine the slopes of these curves, but need to use the structure of the model to identify the aggregate impact of credit shocks.

Our modeling of how households allocate their wealth between housing, liquid assets and mortgage debt is purposefully simple and designed to capture aggregate comovements, not the behavior of individual households. Just as the Calvo pricing protocol is a crude description of how individual firms price, our family construct is a crude description of how households save. Fully micro-founding household portfolio choices would require a rich three-state model of the housing market similar to that studied by Kaplan et al. (2020) and Boar et al. (2020). Such a model would entail significant micro-level non-convexities that would make full-information estimation infeasible. Since our goal is to use state-level data to isolate the impact of monetary policy in stabilizing the response of real variables to credit shocks, our conjecture is that our simplification provides a useful approximation, analogous to the Calvo pricing approach in modeling the dynamics of inflation.

2.8 Solution Method

To use Bayesian techniques in estimation, we need an efficient solution method. We thus use a piece-wise linear approximation method developed by Eggertsson and Woodford (2003), Guerrieri and Iacoviello (2015) and Jones (2017) to deal with the ZLB. We next describe the method, which we adapt for our application combining state-level and aggregate variables.

The linearized system of equations defining the equilibrium of the model is

$$\mathbf{A}\mathbf{x}_t(s) = \mathbf{B}\mathbf{x}_{t-1}(s) + \mathbf{D}\mathbb{E}_t\mathbf{x}_{t+1}(s) + \mathbf{F}\epsilon_t(s) + \bar{\mathbf{A}}^r\mathbf{x}_t + \bar{\mathbf{B}}^r\mathbf{x}_{t-1} + \bar{\mathbf{D}}^r\mathbb{E}_t\mathbf{x}_{t+1} + \bar{\mathbf{F}}^r\epsilon_t,$$

where $\mathbf{x}_t(s)$ is a vector that collects state-level variables, $\epsilon_t(s)$ are the state-level shocks, $\mathbf{x}_t$ are the aggregate variables, ϵ_t are the aggregate shocks, and the bold matrices are functions of the parameters of the model.

The economy may be in one of two regimes. Let $r = 0$ denote periods when the ZLB does not bind, and $r = 1$ periods when it does. Because each individual island is atomistic, whether the ZLB binds does not, up to a first-order approximation, affect the dependence of an island's variables on its own lagged or future variables, or its

shocks. Thus the matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{D}$, and $\mathbf{F}$ are not regime-dependent.

Integrating over the individual states under the normalization $\int \epsilon_t(s)\mathrm{d}s = 0$, and adding the Taylor rule, we derive the dynamics of aggregate variables as

$$\mathbf{A}^r \mathbf{x}_t = \mathbf{C}^r + \mathbf{B}^r \mathbf{x}_{t-1} + \mathbf{D}^r \mathbb{E}_t \mathbf{x}_{t+1} + \mathbf{F}^r \epsilon_t,$$

where $\mathbf{x}_t$ collects all the aggregate variables, $\int \mathbf{x}_t(s)\mathrm{d}s$, and the nominal interest rate i_t. Notice that the coefficient matrices depend on the monetary policy regime.

The algorithm is based on a piece-wise linear solution of the equilibrium conditions in the two monetary policy regimes $r = \{0, 1\}$, under the assumption that agents observe all shocks in each period, but believe that no other shocks are possible in the future. Under the regime when the ZLB does not bind, we use standard methods to derive the time-invariant VAR representation of the aggregate variables

$$\mathbf{x}_t = \mathbf{J}^0 + \mathbf{Q}^0 \mathbf{x}_{t-1} + \mathbf{G}^0 \epsilon_t. \tag{12}$$

When the ZLB binds, we iterate on the date $T\left(\mathbf{x}_{t-1}, \epsilon_t, f_t\right)$ at which the ZLB will stop binding conditional on no further aggregate shocks hitting the economy. Here f_t encodes forward guidance announcements, if any. For periods after T, the economy evolves following (12). Prior to T, the solution of the model is

$$\mathbf{x}_t = \mathbf{J}_t + \mathbf{Q}_t \mathbf{x}_{t-1} + \mathbf{G}_t \epsilon_t, \tag{13}$$

where the reduced-form matrices are updated iteratively according to

$$\begin{aligned}
\mathbf{Q}_t &= \left[\mathbf{A}^1 - \mathbf{D}^1 \mathbf{Q}_{t+1}\right]^{-1} \mathbf{B}^1 \\
\mathbf{J}_t &= \left[\mathbf{A}^1 - \mathbf{D}^1 \mathbf{Q}_{t+1}\right]^{-1} \left(\mathbf{C}^1 + \mathbf{D}^1 \mathbf{J}_{t+1}\right) \\
\mathbf{G}_t &= \left[\mathbf{A}^1 - \mathbf{D}^1 \mathbf{Q}_{t+1}\right]^{-1} \mathbf{E}^1,
\end{aligned} \tag{14}$$

These recursions are the same as those in Guerrieri and Iacoviello (2015) and Kulish and Pagan (2017). In short, we iterate backwards from the conjectured date of exit from the ZLB regime by using the equilibrium conditions consistent with $i_t = 0$. We iterate on the conjectured date T until convergence.[18]

---

[18]See Kulish and Pagan (2017) and Jones (2017) for more details on the recursion underlying this solution. Jones (2017) contrasts this method with fully non-linear ones in smaller scale models and

Having solved for the aggregate outcomes we solve for the island-level variables,

$$\mathbf{x}_t(s) = \underbrace{\mathbf{Q}\mathbf{x}_{t-1}(s) + \mathbf{G}\epsilon_t(s)}_{\text{state-level component}} + \underbrace{\bar{\mathbf{J}}_t + \bar{\mathbf{Q}}_t\mathbf{x}_{t-1} + \bar{\mathbf{G}}_t\epsilon_t}_{\text{aggregate component}}. \tag{15}$$

The matrices $\bar{\mathbf{J}}_t$, $\bar{\mathbf{Q}}_t$, and $\bar{\mathbf{G}}_t$ encode how an island responds to aggregate-level variables and vary over time because of the ZLB. In contrast, the matrices $\mathbf{Q}$ and $\mathbf{G}$ which determine how an island's variables depend on their own lags and island-specific shocks are time-invariant. Intuitively, since each island is of measure zero, shocks to an individual island do not change the expected date at which the ZLB will stop binding for the rest of the economy. Moreover, agents on each island take aggregate prices, including the interest rate, as given. Therefore, conditional on the aggregate state variables, the ZLB regime does not affect the evolution of island-level relative variables. From the perspective of agents on a given island, the presence of the ZLB acts like any other aggregate shock which, up to a first-order approximation, does not change how that island responds to its own history of idiosyncratic shocks.

Equation (15) reveals that our model has a special structure. First, the mapping from state-level histories of shocks to state-level outcomes is time-invariant and common to all states. Second, the mapping from aggregate histories of shocks to state-level variables, though time-varying, is also common to all states. We exploit this structure to construct the likelihood function, as we describe below.

3 Estimation

We next explain how we chose parameters for our model. We first discuss the parameters we assign values to, and then the ones we estimate. 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, θ_w and θ_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.

---

shows that that the piece-wise linear method we use here is accurate. Lepetyuk et al. (2017) contrast these perturbation-based methods with fully non-linear projection methods in larger-scale models and reach similar conclusions regarding their accuracy.

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 et al. (2005)
α_r	0.86	Taylor rule persistence	Justiniano et al. (2011)
α_π	1.71	Taylor coefficient inflation	Justiniano et al. (2011)
α_y	0.05	Taylor coefficient output	Justiniano et al. (2011)
α_x	0.21	Taylor coefficient output growth	Justiniano et al. (2011)
Parameters chosen to match steady-state target			
$-\ln(\beta)$	2.46%	Annual discount rate	2% real rate
η	0.084	Weight on housing	Housing-to-income ratio of 2.5
$\bar{m}$	0.0044	Credit limit	Debt-to-housing ratio of 0.29

3.1 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 $1/2$. 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 30-year mortgages, approximately 13 years.[19] 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 tradable and non-tradable goods, σ, equal to 0.5; and the elasticity of substitution between varieties of tradable goods produced in different islands, κ, equal to 4, the estimate of Simonovska and Waugh (2014). We follow Christiano et al. (2005) in choosing ψ to ensure a wage markup of 5%.[20] We use the Justiniano et al. (2011) estimates of 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

---

[19]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 $(1+r)/(1+r-\gamma)$.

[20]We show in our Online Appendix that our conclusions are not sensitive to the value of ψ.

state LTV ratio is chosen so that the aggregate debt to housing ratio is equal to 0.29, a number once again computed from the SCF. Since the debt constraint binds, these last two targets imply an aggregate debt to (annual) income ratio of $2.5 \times 0.29 = 0.725$.

3.2 Estimation Procedure

Our goal is to simultaneously use regional and aggregate data to estimate the model and identify the impact of credit shocks on real activity. In contrast to Nakamura and Steinsson (2014) and Beraja et al. (2019), who also use regional data to understand the effect of various shocks on macroeconomic aggregates, our approach simultaneously uses state and aggregate-level data in a full-information estimation. One important technical challenge we need to overcome in order to compute the likelihood function is the presence of the occasionally binding ZLB, which renders conventional methods infeasible due to the very large number of island-level state variables and the resulting curse of dimensionality. We therefore exploit the special structure of our model and develop an efficient algorithm that circumvents these difficulties. We first describe the details of our approach, and then report our results.

3.2.1 The Data

Regional Data. We use a panel of employment, household spending, wages, household debt and house prices in the cross-section of 51 U.S. states from 1999 to 2015. The household debt information is from the FRB New York Consumer Credit Panel.[21] Since the measure of debt in the data is the book value of outstanding debt, we construct a corresponding series for our model using

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

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

We use data on house prices from the FHFA, and data on employment, consumption expenditures and wages from the BEA. Our measure of employment is

---

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

the employment-population ratio in a given state. Our measure of wages is total employee compensation divided by the number of workers. Since we do not model the construction sector, we subtract construction employment and compensation in computing measures of employment and wages.[22]

Aggregate Data. We use aggregate data on employment, household consumption expenditures, wages, household debt, house prices, inflation, and the Fed Funds rate from 1984 to 2015. An additional critical input in the estimation is the sequence of expected durations of the ZLB between 2009 and 2015, which we take from the Blue Chip Financial Forecasts survey and the New York Federal Reserve's Survey of Primary Dealers. The Online Appendix provides more detail on the data we use.

3.2.2 The Likelihood Function

The conventional approach to estimating the model would be to use a likelihood function that directly combines state and aggregate data. However, this approach is infeasible because of the non-linearity induced by the ZLB and the curse of dimensionality which arises from our use of 51 regions, each of which has 11 state variables.

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

To understand our approach, recall that the evolution of variables in any given state is given by (15), which we reproduce here for convenience:

$$\mathbf{x}_t(s) = \mathbf{Q}\mathbf{x}_{t-1}(s) + \mathbf{G}\epsilon_t(s) + \bar{\mathbf{J}}_t + \bar{\mathbf{Q}}_t\mathbf{x}_{t-1} + \bar{\mathbf{G}}_t\epsilon_t.$$

Since the last three reduced-form matrices are time-varying functions of the underlying structural parameters whenever the ZLB binds, computing them for each parameter draw in the estimation is infeasible. The structure of the model allows us, however,

---

[22]See our Online Appendix for an extension that explicitly models the construction sector and uses data on construction employment in estimation. Our results are robust to this modification.

to write the deviation of island-level variables from their economy-wide averages,

$$\hat{\mathbf{x}}_t(s) = \mathbf{x}_t(s) - \int \mathbf{x}_t(s)\mathrm{d}s,$$

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

$$\hat{\mathbf{x}}_t(s) = \mathbf{Q}\hat{\mathbf{x}}_{t-1}(s) + \mathbf{G}\epsilon_t(s). \tag{16}$$

This structure significantly speeds up the computation of the likelihood function.

We therefore use the representation in (13) and (16) to estimate the model using state-level and aggregate U.S. data. To do so, we first express each state's observable variable as deviations from its aggregate counterpart by subtracting a full set of time effects, one for each year and each variable. We also subtract a state-specific fixed effect and time trend for each observable since in our model all islands are ex-ante identical. Since island-level shocks are independent and do not affect aggregate outcomes, we can write the log-likelihood as the sum of each individual state's likelihood, computed from (16) and the aggregate likelihood, computed from (13):

$$\log \mathcal{L} = \sum_{s=1}^{51} \bar{\omega}(s) \log \mathcal{L}(s) + \log \mathcal{L}^a,$$

where $\mathcal{L}(s)$ is the contribution to the likelihood of an individual state's data, $\mathcal{L}^a$ is the contribution to the likelihood of the aggregate data, and $\bar{\omega}(s)$ is the relative weight of state s. This additive structure follows from the block structure implied by representation (16) which allows us to express the determinant of the variance-covariance matrix of the forecast errors as the product of the determinants of the variance-covariance matrices of the individual states, as we show in the Online Appendix.

To see how we construct the likelihood for an individual state, denote by $\bar{N} = 5$ the number of observable variables and $\bar{T} = 17$ (1999 to 2015) the number of years of data that are available. With Gaussian errors, the log-likelihood function is

$$\log \mathcal{L}(s) = -\left(\frac{\bar{N}\bar{T}}{2}\right) \log 2\pi - \frac{1}{2}\sum_{t=1}^{\bar{T}} \log \det \mathbf{S}_t - \frac{1}{2}\sum_{t=1}^{\bar{T}} \tilde{\mathbf{x}}_t'(s)\mathbf{S}_t^{-1}\tilde{\mathbf{x}}_t(s).$$

where $\tilde{\mathbf{x}}_t(s)$ is the vector of forecast errors, calculated using the Kalman filter and the structural matrices $\mathbf{Q}$ and $\mathbf{G}$ from (16), and $\mathbf{S}_t$ is the covariance matrix of $\tilde{\mathbf{x}}_t(s)$. We

compute the aggregate likelihood component in a similar fashion, taking into account that the sequence of forecast errors and their covariance matrices are functions of the underlying time-varying reduced-form matrices $\mathbf{Q}_t$ and $\mathbf{G}_t$ from (13).

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. Notice that by simultaneously using state and aggregate data to construct the likelihood, our approach differs from that of Beraja et al. (2019), who first apply limited information methods to estimate the degree of wage stickiness using state-level data. They then use that estimate to form a prior in a second step, which uses only aggregate data to form the likelihood function.

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 quarterly and for a longer time-period (1984 to 2015). To account for differences in the size of different states and ensure that smaller states exert a relatively smaller influence on the shape of the likelihood, we weight the likelihood contribution of each state by their 1999 population shares. In a robustness exercise reported in the Online Appendix, we also compute the likelihood under equal weighting. In our baseline, we assign an equal weight to the aggregate likelihood and the combined state-level likelihood.

3.3   Parameter Estimates

Table 2 reports moments of the prior and posterior distributions of the structural parameters we estimate. We use diffuse priors. We find, consistent with the work of Del Negro et al. (2015), that wages and prices are sticky, with a modal estimate of θ_p of 0.96 and a modal estimate of θ_w of 0.86. 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 lower degree of price stickiness, consistent with Beraja et al. (2019). We show the implications under these alternative parameters in our Online Appendix.

The mean estimate of the Pareto tail parameter α is equal to 3.44, implying a 0.47% spread between the subjective discount rate and the interest rate at the annual frequency. Notice that this parameter is well identified with a relatively tight poste-

Table 2: Estimated Structural Parameters

Parameter	Prior				Posterior		
	Dist	Median	10%	90%	Mean	10%	90%
θ_p	B	0.5	0.2	0.8	0.96	0.94	0.98
θ_w	B	0.5	0.2	0.8	0.86	0.82	0.89
α	N	2.5	1.2	3.8	3.44	3.05	3.93

rior distribution. The 10th and 90th percentiles are 3.0 and 3.9, respectively. This estimate implies that the fraction of household members who are liquidity constrained is relatively low, 0.3% in steady-state. Moreover, as we show in the Online Appendix, the model's implications for the marginal propensity to consume out of a transitory income shock are also very similar to those of the frictionless model. In this sense, our results do not rely on assuming severe liquidity constraints.

The Appendices report the posterior estimates of the parameters governing the persistence and volatility of state and aggregate shocks, and the time series of shocks.

3.4   Identification of α

We next provide intuition for how the key parameter α is identified from the comovement of state-level household debt, consumption, and employment. Table 3 reports coefficient estimates from univariate regressions of changes in employment and consumption on changes in household debt in the boom (2002 to 2007) and the bust (2007 to 2010). As pointed out by Mian and Sufi (2014), state-level data show a strong correlation between these variables. For example, as the first column of the table shows, an increase in debt-to-income of 100 percent is associated with a 7 percent increase in employment and 15 percent increase in consumption from 2002 to 2007.

We next discuss the model's predictions for these comovements. We use our mean estimates in Table 2 to simulate a large panel of states and run identical regressions in the model. As the second column of Table 3 shows, our baseline model reproduces the correlations in the data very well. Consider next what happens if we simulate data from the model using higher and lower values of α. When idiosyncratic uncertainty is high ($\alpha = 2$), employment and consumption comove much more strongly with changes in household debt than in the data. For example, a 100 percent increase in household debt-to-income is associated with a 22 percent increase in employment and

Table 3: Regressions with Simulated Panel Datasets From Model

	Data	Baseline	$\alpha = 2^{\dagger}$	$\alpha = 10^{\dagger}$	$\alpha = 10^{\ddagger}$
A. Dependent Variable: Δ Debt. Boom, 2002 to 2007					
Δ Employment	0.07	0.06	0.22	0.02	0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Δ Consumption	0.15	0.11	0.31	0.02	0.02
	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)
B. Dependent Variable: Δ Debt. Bust, 2007 to 2010					
Δ Employment	0.10	0.06	0.22	0.03	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Δ Consumption	0.17	0.10	0.31	0.03	0.02
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)

$\dagger$: Other parameters as in baseline. $\ddagger$: Other parameters reestimated.
Standard errors in parentheses

31 percent increase in consumption. In contrast, when idiosyncratic uncertainty is low ($\alpha = 10$), employment and consumption comove little with changes in household debt. For example, a 100 percent increase in household debt-to-income is associated with a 2 percent increase in employment and consumption.

In the Online Appendix, we report the predictions of our model for the comovement between employment, consumption, and household debt for longer leads and lags using local projections. We show that simulated series from the model with too high or too low values of α generate impulse responses that greatly differ from those computed in the data. We also report in the Online Appendix that a GMM-based procedure produces similar estimates of α as in our baseline.

3.5 Separating Credit Shocks from Other Demand Shocks

How does our framework identify the separate contribution of shocks that change household credit limits from that of discount factor shocks, which also lead to comovements between consumption and household savings? To see this, recall that we assume that there is no spread between the rate at which households save in the liquid asset and the mortgage interest rate. Since liquid assets provide liquidity services, agents find it optimal to always borrow up to the limit.[23] Thus, though both shocks

---

[23]We relax this assumption in the Online Appendix when studying a model with default.

to credit limits and the discount factor lead to comovements between consumption and household net asset positions, only shocks to the credit limit affect household debt. Discount factor shocks change liquid asset holdings, but not mortgage debt.

The result that all homeowners are against their borrowing limit is admittedly extreme, but it is broadly consistent with the data. To see this, notice that the majority of household debt in the data is mortgage debt, which is refinanced infrequently. The vast majority of households are therefore against the limit implicitly built-in their mortgage contract, namely the requirement that they make the minimum principal and interest payments on the mortgage. As is well-known, most households in the data do not exceed these minimum payments. For example Amromin et al. (2007) and Adelman et al. (2010) show that only approximately 15–20% of borrowers have ever made mortgage payments greater than the minimum required ones. Boar et al. (2020) use this and other evidence and a rich model of the housing market to show that the majority of homeowners are indeed borrowing constrained.

To the extent to which some shocks lead unconstrained households to reduce debt, say for precautionary reasons, our estimation will interpret them as credit shocks. Our exercise can thus disentangle the contribution of shocks that change the amount households borrow from other shocks that only change liquid savings.

3.6   Variance Decompositions

To interpret the estimates of the persistence and volatility of shocks, Table 4 reports the forecast error variance decompositions of state and aggregate observables.

Consider first Panel A of Table 4 which reports this variance decomposition for state-level variables, expressed as deviations from their aggregate counterparts. State-level credit shocks account for 7% of the variation of *relative* employment and 30% of the variation in relative household spending. State-level productivity shocks account for about 10% of the variation in relative employment and a negligible fraction of the volatility of consumption. Shocks to the disutility from work account for about 71% of the volatility of employment and 35% of the volatility of consumption. Finally, shocks to the individual states' discount rates account for 12% percent of the volatility of relative employment and 35% of the volatility of relative consumption, while shocks to the preference for housing have a negligible impact on state-level real variables. The variance decompositions for wages, household debt-to-income, and house prices

31

Table 4: Unconditional Variance Decomposition, Infinite Horizon

Variable \ Shock	Collateral	Housing	Productivity	Leisure	Discount	Policy	Markup
A. State-level							
Spending	30.3	0.4	0.0	34.6	34.7	-	-
Employment	7.1	0.1	9.6	71.1	12.2	-	-
Wages	1.9	0.0	0.0	97.9	0.1	-	-
Debt-to-income	98.8	1.1	0.0	0.1	0.0	-	-
House prices	2.3	95.0	0.0	2.5	0.2	-	-
B. Aggregate-level							
Consumption	1.5	0.1	22.7	8.8	47.6	10.7	8.7
Employment	1.8	0.1	4.0	10.9	59.2	13.2	10.8
Wages	0.3	0.0	3.7	69.6	3.0	0.3	23.1
Debt-to-income	78.8	5.9	0.5	2.4	12.0	0.1	0.4
House prices	0.0	87.3	3.3	1.2	5.3	1.6	1.3
Fed Funds rate	3.4	0.3	11.2	2.6	76.5	3.6	2.2
Inflation	0.8	0.1	2.9	12.2	2.6	0.2	81.3

show that these series are explained by leisure shocks, credit shocks, and housing preference shocks respectively. We also report in the Online Appendix the variance decompositions at shorter horizons. At shorter horizons, productivity shocks are a more important determinant of employment, while discount factor shocks are a more important determinant of relative consumption fluctuations across states.

Consider next Panel B of Table 4 which reports a similar decomposition for aggregate variables. Unconditionally, absent the ZLB, household credit shocks account for only 1.5% and 1.8% of the volatility of aggregate employment and consumption, respectively. Intuitively, credit shocks matter much more at the state-level than in the aggregate because individual states are part of a monetary union and cannot use monetary policy to offset their effects. Productivity shocks account for about 4% of the employment volatility and 23% of the volatility of consumption. Discount rate shocks account for a substantial fraction of movements in employment and consumption, while shocks to the disutility from work or preference for housing have little effect. Monetary policy shocks account for about 10% of the variation in employment and consumption and a negligible amount of the variation in other observables. Consistent with the variance decompositions obtained using the model of Smets and

Wouters (2007), we find that the behavior of inflation is mostly explained by markup shocks, while wages are accounted for by labor disutility shocks.

3.7   Additional Implications

We show in the Online Appendix that our model reproduces variables that were not directly used in estimation. We find that most of the decline in state-level employment was due to declines in non-tradable employment, consistent with Mian and Sufi. Our model also reproduces well the medium term movements in mortgage rates.

4   Credit Shocks in the Great Recession

We next study the role of household credit shocks in shaping the dynamics of employment and consumption in the cross-section and in the aggregate during the Great Recession. Recall that in our model the credit limit always binds, so household credit fluctuates due to exogenous changes in the credit limit $m_t(s)$, and house prices, $e_t(s)$. The latter fluctuate mostly due to shocks to the preference for housing. Since housing is in fixed supply and separable in preferences, housing preference shocks only affect equilibrium outcomes through their effect on the credit limit. We thus refer to both shocks to the credit limit $m_t(s)$ and housing preferences $\eta_t^h(s)$ as *credit shocks*.

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 et al. (2013) and Mian and Sufi (2014). We then study the aggregate implications of credit shocks. We show that aggregate credit shocks alone generate a modest decline in the natural interest rate during the Great Recession and are not sufficient to trigger the ZLB on their own. Absent other shocks, aggregate credit shocks would therefore have had a relatively small effect on employment. We finally study how the effect of credit shocks on employment is magnified when monetary policy is constrained.

4.1   Role of Credit Shocks at the State Level

We report the effects of state-level credit shocks on employment and consumption in Figure 6. The horizontal axes show the actual data we used to extract the state-level component of the shocks. We construct a counterfactual series for each variable by

Figure 6: Effect of Credit Shocks on State Employment and Consumption

shutting down all shocks other than those to credit, and plot the changes in these series on the vertical axes. The solid line in each of these plots is the 45 degree line.

Consider first the 2002 to 2007 period, which we label the *boom*. The upper panels of Figure 6 show that credit shocks account for a substantial fraction of the differential changes in employment and consumption during the boom. The correlation between the change in the counterfactual series driven by credit shocks only and the change in the data is 0.68 for employment and 0.58 for consumption. The slope coefficient of a regression of these counterfactual changes against the data is 0.52 for employment and 0.30 for consumption. Credit shocks thus account for about one-half of the relative movements in employment and one-third of the relative movements in consumption across states during the boom.

The lower panels of Figure 6 illustrate the corresponding patterns during the 2007 to 2010 recession, which we term the *bust*. Credit shocks account for an even larger fraction of the differential change in employment and consumption during the bust. The slope coefficients are 0.56 for employment and 0.50 for consumption, implying that credit shocks account for one-half of the differential decline in employment and consumption across states. We emphasize that these results are not simply an artefact

Table 5: Contribution of Individual Shocks to Observed Change

|	Credit Limit	Housing	Productivity	Leisure	Discount
Δ Employment					
– 2002 to 2007	0.47	0.05	−0.09	−0.09	0.58
– 2007 to 2010	0.52	0.04	0.14	−0.08	0.38
Δ Consumption					
– 2002 to 2007	0.27	0.03	0.00	−0.08	0.74
– 2007 to 2010	0.48	0.03	0.00	−0.03	0.55

of the modeling choices we have made, but rather a result of the estimation and the path for shocks extracted by the Kalman smoother from the data. If instead we used a higher value of α in simulating the model, credit shocks would have had much smaller effects, as we report in the Online Appendix.

Table 5 illustrates the role of all additional shocks in explaining the relative changes in employment and consumption during these two periods. We use the Kalman smoother to extract the shocks and construct counterfactual employment and consumption panels by turning on one shock at a time. Table 5 reports the slope coefficients of the projection of the counterfactual series on the actual data.[24] In addition to the credit shocks, discount factor shocks account for the bulk of the remaining fluctuations in employment and consumption. These shocks capture other sources of state-level movements in consumption, such as perhaps heterogeneity in fiscal transfers across states, expectations of future growth, precautionary savings motives, and other factors. In our Online Appendix, we show that explicitly introducing differential changes in government spending across states reduces the contribution of discount factor shocks, but leaves the contribution of credit shocks unchanged.

We next zoom in on several states and show in Figure 7 how the time-series of relative employment is shaped by credit, discount factor, and other shocks. The top row of the figure plots relative employment for two states – Arizona and Nevada – that experienced large changes in credit and employment. The bottom row of the figure shows two contrasting states – Virginia and Illinois – that experienced mild fluctuations. The figure reinforces the point that credit shocks were an important driver of the relative change in employment across states over this period.

---

[24] These coefficients do not sum to one because part of the variation of the observable series is attributed to initial conditions.

Figure 7: State Level Decomposition, Employment: Four States
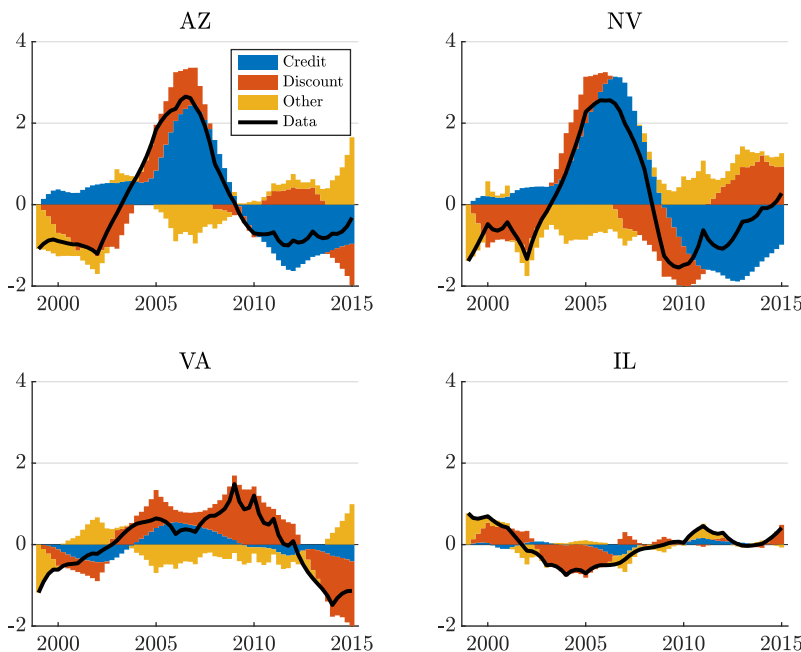
Our findings thus corroborate the conclusion of Mian and Sufi (2014) that credit shocks were an important determinant of the variation in employment and consumption across U.S. regions in the years before and after the Great Recession. Two caveats are in order. First, even though in our model changes in household credit are driven by shocks to the credit limit, in reality a tightening of household credit may have been induced by lower credit supply, as emphasized by Justiniano et al. (2019). Our approach is unable to differentiate between these sources of the decline in household leverage: changes in credit limits in our framework thus implicitly capture both supply and demand-side considerations. What is critical for our study is the comovement between changes in household credit and consumption, which our model reproduces well. We show in our Online Appendix that our results are robust to modeling changes in household credit as arising from credit supply shocks.[25]

Second, our model attributes all relative changes in state-level variables to idiosyncratic state-level shocks. In practice, as pointed out by Mian and Sufi (2014), an important source of state-level variation was the asymmetry with which individ-

---

[25]Additional factors that may have contributed to the observed correlation between changes in household credit and real variables include wealth effects from changes in house prices, as in Berger et al. (2018), or a precautionary motive associated with the rise in uncertainty. Our use of a representative agent framework and a first-order solution implicitly abstracts from these effects.

ual states responded to aggregate shocks. The approach we pursue is motivated by tractability considerations: if the matrices $\bar{\mathbf{J}}_t$, $\bar{\mathbf{Q}}_t$ and $\bar{\mathbf{G}}_t$ of equation (15) that determine how an individual island responds to aggregate shocks were state-specific, say due to heterogeneity in housing supply elasticities, the likelihood function would not be additively separable into individual state-level likelihood contributions. Our estimates of the state-level shocks should therefore be interpreted more broadly as capturing not only state-specific idiosyncratic disturbances, but also heterogeneity in the extent to which state-level variables responded to aggregate shocks. In our Online Appendix, we relax the assumption that all states react in an identical way to state and aggregate shocks by separating states into three groups that differ in their housing supply elasticities. In this extension we can only use state-level data to estimate the parameters, but find that our estimates and the model's implications for the relative importance of credit shocks are similar to those in the baseline.

4.2 Role of Credit Shocks at the Aggregate Level

We next turn to discussing the model's implications for the role of credit shocks in explaining fluctuations in the aggregate. Figure 8 presents the dynamics of the key aggregate time series during the 1995 to 2015 period in the data. We HP-filter the debt-to-income series and report it as deviations from the trend in Panel A. Panel B reports the evolution of the Fed Funds Rate and the implied natural rate predicted by the model. Panel C reports the data on employment, measured as the total number of employees on non-farm payrolls, scaled by the U.S. population and expressed as percent deviations from a linear trend. 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 shocks.

We construct a counterfactual series for each of these variables by shutting down all shocks other than credit shocks. Panel A of Figure 8 shows that credit shocks drive most of the movements in household credit. As Panel B shows, however, credit shocks alone generate modest movements in the natural rate of interest. Specifically, our estimates imply that the natural rate declined from 4.7% at the beginning of 2007 to -1.4% by the end of 2012. Credit shocks account for only one-tenth of this decline: the natural rate would have fallen from 2.4% to 1.8% in the absence of other shocks.

Consider next how employment responds to credit shocks. As the lower-left panel

Figure 8: Effect of Credit Shocks Alone on Aggregate Variables
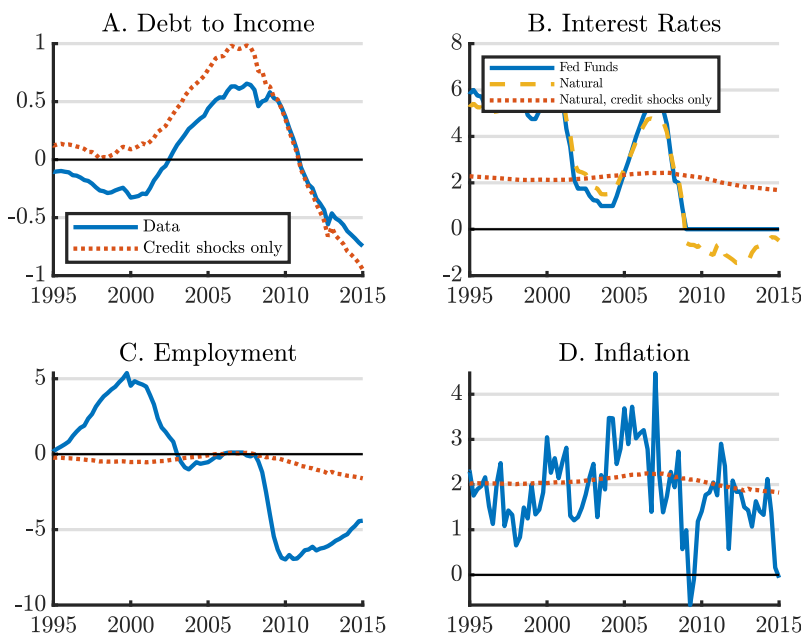
of Figure 8 and Table 6 shows, the Great Recession was associated with a nearly 6.2% drop in the employment-population ratio from 2007 to 2012. Credit shocks alone account for a modest 1.3% drop, one-fifth of the actual decline. Employment falls little because credit shocks generate a small decline in the natural rate, insufficient to trigger the ZLB. Absent additional shocks monetary policy is unconstrained and mimics the dynamics of the natural rate well, implying minor movements in employment and inflation. As is well-understood, however, the ZLB brings about important non-linearities in how aggregate variables respond to shocks. If household credit shocks are accompanied, as they were in the data, by other shocks that reduce the Fed's ability to cut interest rates because of the ZLB, the resulting effects on output can be much greater. How much greater?

Answering this question is challenging because of the Fed's pursuit of *forward guidance* policies. In our estimation we have taken the Fed's forward guidance announcements as given, by setting the expected ZLB duration each quarter equal to the expectations in the data. Since these announcements were responding to the path of all shocks, including credit shocks, we cannot isolate the effect of credit shocks alone without taking a stand on how forward guidance would have been conducted absent such shocks. In short, the pursuit of forward guidance poses an identification challenge, one of decomposing the expected ZLB durations into an endogenous

38

Figure 9: Effect of Credit Shocks on Aggregate Employment
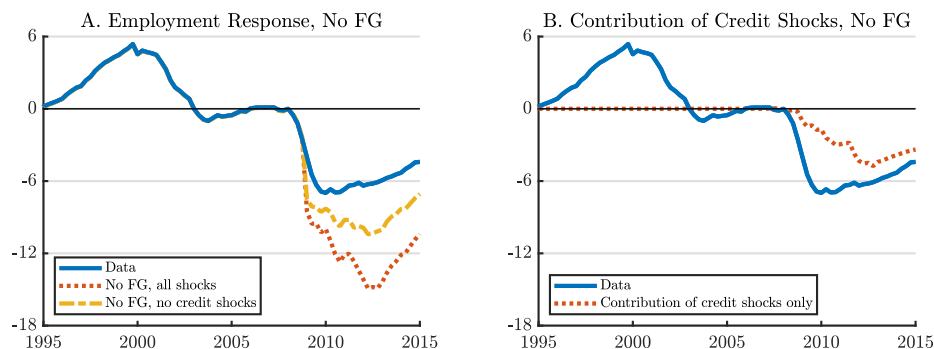
component due to shocks, and one due to forward guidance policies.[26]

We address this challenge by providing bounds on the role of credit shocks under two extreme sets of assumptions. The lower bound discussed above is that arising in the absence of the ZLB. We compute the upper bound by considering a counterfactual experiment in which the Fed *does not* pursue forward guidance. This is an upper bound because the Fed provides no additional monetary stimulus once interest rates hit zero. We compute this upper bound by first calculating two counterfactual series shown in Panel A of Figure 9. The dotted line shows that the recession would have been twice more severe absent forward guidance. The dashed line turns off the credit shocks and computes the impact of all other shocks on employment. The difference between these series provides a measure of how credit shocks would have shaped employment outcomes in the absence of forward guidance. We plot this difference as the dotted line in Panel B of Figure 9 and summarize it in the last column of Table 6.

As the table shows, even under the extreme assumption of no forward guidance, credit shocks would have caused a modest 1.7% drop in employment from 2007 to 2010, one-fifth of the 7.1% decline in the data. In contrast, the contribution of credit shocks to employment is substantially larger by the end of 2012. Credit shocks would have generated a 4.8% drop in employment from 2007 to 2012 in the absence of forward guidance, three-fifths of the actual drop.

Credit shocks contribute little to employment at the onset of the recession because they are relatively modest initially, reflecting the gradual adjustment of household debt in the first years of the recession. For example, household debt-to-income fell by only 0.19 from 2007 to 2010. In contrast, it experienced a much larger 0.57 drop from 2007 to 2012, leading to a larger decline in the natural rate of interest. We

---

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

39

Table 6: Aggregate Effects of Credit Shocks in the Bust

| | | Credit Shocks Only | |
	Model and Data	No ZLB	No FG
2007Q1 to 2010Q1			
– Δ HH Debt to Income	−0.19	−0.24	−0.20
– Δ Employment, %	−7.09	−0.47	−1.72
2007Q1 to 2012Q4			
– Δ HH Debt to Income	−0.57	−0.63	−0.57
– Δ Employment, %	−6.20	−1.34	−4.76

thus conclude that credit shocks may have played a substantial role in delaying the recovery, but had a modest impact in the years immediately following 2007.

To summarize, we find that credit shocks account for a large fraction of the relative variability in consumption and employment across U.S. states, but had a modest impact in the first few years of the Great Recession. Does this result reflect differences in how credit shocks are transmitted to employment in the cross-section as opposed to the aggregate, or differences in the size of these shocks? We answer this question in Tables 6 and 7, in which we compare the response of employment to credit shocks in the aggregate and in two states – Arizona and Nevada – that have experienced relatively large boom and bust cycles in credit and employment.

Consider first Arizona. As Table 7 shows, this state experienced a 0.95 drop in the relative debt-to-income ratio from 2007 to 2010, and an even larger 1.6 drop from 2007 to 2012. This decline was much larger than the corresponding decline in the aggregate debt-to-income ratio and led to a sizable drop in relative employment. As the third column of the table shows, credit shocks generate a 3.1 percent drop in employment in Arizona from 2007 to 2010, and a 3.7 percent drop from 2007 to 2012. The patterns in Nevada are similar.

A comparison of the aggregate and state-level outcomes in Tables 6 and 7 reveals that credit shocks generate larger movements in employment in the cross-section of U.S. states compared to the aggregate because they were much larger. For example, from 2007 to 2013, the elasticity of credit shocks to employment in the aggregate was 2.1% (-1.34%/-0.63) at our lower bound estimate which assumes away the ZLB and 8.4% (-4.76%/-0.57) at our upper bound estimate which imposes the ZLB and assumes away forward guidance. In contrast, the corresponding ratio in Arizona is 2.2% (-

Table 7: Relative State-Level Effects of Credit Shocks in the Bust

	Model and Data		Credit Shocks Only	
	Arizona	Nevada	Arizona	Nevada
2007 to 2010				
– Δ HH Debt to Income	−0.95	−1.39	−1.01	−1.46
– Δ Employment, %	−3.11	−3.47	−3.13	−4.02
2007 to 2013				
– Δ HH Debt to Income	−1.60	−2.15	−1.71	−2.29
– Δ Employment, %	−3.07	−2.41	−3.73	−4.84

3.73%/1.71), similar to the 2.1% in Nevada (-4.84%/-2.29). Thus, the transmission of credit shocks to employment is similar in the cross-section as in the aggregate. The larger contribution of credit shocks in the cross-section reflects their larger size.

5 Conclusions

A popular account of the U.S. Great Recession is that declines in households' ability to borrow led to a reduction in consumption and employment due to price rigidities and constraints on monetary policy. This view is motivated, in part, by the observation that employment co-moves strongly with changes in household debt in the cross-section of U.S. regions. This paper asks: what are the *aggregate* implications of these regional correlations?

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

To make full-information Bayesian estimation feasible, we propose a novel methodology that exploits the structure of the model and identifies parameters from the relative variation of state-level variables, in addition to fluctuations in the aggregate.

Our approach allows us to efficiently calculate the likelihood function in an economy with a large number of aggregate and regional state variables and with non-linear dynamics caused by the ZLB constraint on monetary policy.

Our estimates imply that credit shocks account for about 40% of the differential rise and fall in *state-level* employment and consumption during the boom and bust years. In contrast, credit shocks alone generate modest movements in *aggregate* employment from 2007 to 2010, about one-fifth of the observed decline, even though monetary policy was constrained by the ZLB. This result reflects the gradual nature of household deleveraging, which caused a gradual decline in the natural interest rate.

The persistent nature of household deleveraging implies, however, that credit shocks became more important over time as the quantity of household debt continued to decline and the natural rate fell even further. Our estimates suggest that credit shocks account for a sizable fraction of the employment gap in 2012. The tightening in household-level credit limits thus helps account for the slow recovery of employment in the aftermath of the Great Recession.

A caveat is in order. We have focused on the *household leverage* view of the Great Recession and have abstracted from other forces, such as constraints on financial intermediaries, that played an important role during the acute phase of the crisis. A limitation of our approach is that we cannot estimate the *interactions* between household credit risk and the net worth of intermediaries. Though beyond the scope of this paper, explicitly modeling such interactions is an important area for future research.

References

Adelman, Saul W., Mark L. Cross, and David G. Shrider, "Why Do Homeowners Make Mortgage Curtailment Payments?," *Journal of Housing Research*, 2010, *19* (2), 195–212.

Amromin, Gene, Jennifer Huang, and Clemens Sialm, "The Tradeoff Between Mortgage Prepayments and Tax-Deferred Retirement Savings," *Journal of Public Economics*, 2007, *91* (10), 2014–2040.

Arellano, Cristina and Ananth Ramanarayanan, "Default and the Maturity Structure in Sovereign Bonds," *Journal of Political Economy*, 2012, *120*, 187–232.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 2019, *87* (6), 1789–1833.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *The Review of Economic Studies*, 2018, *85* (3), 1502–1542.

Bernanke, Ben S and Mark Gertler, "Agency Costs, Net Worth, and Business Fluctuations," *The American Economic Review*, 1989, *79*, 14–31.

_ , _ , **and Simon Gilchrist**, "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, 1999, *1*, 1341–1393.

Blanco, Andres, "Optimal Inflation Target in an Economy with Menu Costs and Zero Lower Bound," *American Economic Journal: Macroeconomics*, 2019.

Boar, Corina, Denis Gorea, and Virgiliu Midrigan, "Liquidity Constraints in the US Housing Market," 2020. NBER WP 23345.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, "Understanding Booms and Busts in Housing Markets," *Journal of Political Economy*, 2016, *124* (4), 1088–1147.

Challe, Edouard and Xavier Ragot, "Precautionary Saving Over the Business Cycle," *The Economic Journal*, 2016, *126* (590), 135–164.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Coibion, Olivier and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 197–232.

Cooley, Thomas F and Gary D Hansen, "The Inflation Tax in a Real Business Cycle Model," *The American Economic Review*, 1989, pp. 733–748.

Eggertsson, Gauti and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *1*, 139–234.

Eggertsson, Gauti B and Paul Krugman, "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach," *The Quarterly Journal of Economics*, 2012, *127* (3), 1469–1513.

Faria-e-Castro, Miguel, "Fiscal Multipliers and Financial Crises," 2018.

Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 2017, *125* (1), 140–223.

Fernald, John G, Robert E Hall, James H Stock, and Mark W Watson, "The Disappointing Recovery of Output After 2009," *Brookings Papers on Economic Activity*, 2017, *48* (1), 1–81.

Garriga, Carlos, Rodolfo Manuelli, and Adrian Peralta-Alva, "A Macroeconomic Model of Price Swings in the Housing Market," *American Economic Review*, 2019, *109* (6), 2036–72.

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

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

Gourinchas, Pierre-Olivier, Thomas Philippon, and Dimitri Vayanos, "The Analytics of the Greek Crisis," in "NBER Macroeconomics Annual 2016, Volume 31" NBER Chapters, National Bureau of Economic Research, Inc, 2016.

Guerrieri, Luca and Matteo Iacoviello, "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, 2015, *70*, 22–38.

_ **and** _ , "Collateral Constraints and Macroeconomic Asymmetries," *Journal of Monetary Economics*, 2017, *90*, 28–49.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap*," *The Quarterly Journal of Economics*, 03 2017, *132* (3), 1427–1467.

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

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

Jones, Callum, "Unanticipated Shocks and Forward Guidance at the ZLB," 2017.

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

_ , _ , **and** _ , "Credit Supply and the Housing Boom," *Journal of Political Economy*, 2019, *127* (3), 1317–1350.

_ , **Giorgio Primiceri, and Andrea Tambalotti**, "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 2011, *14*, 101–121.

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

_ , **Kurt Mitman, and Giovanni L Violante**, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Kehoe, Patrick, Virgiliu Midrigan, and Elena Pastorino, "Debt Constraints and the Labor Wedge," *American Economic Review Papers & Proceedings*, 2016, *106* (5), 548–53.

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov, "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 2011, *43* (2-3), 255–296.

_ **and John Moore**, "Credit Cycles," *Journal of Political Economy*, 1997, *105* (2), 211–248.

_ **and** _ , "Liquidity, Business Cycles, and Monetary Policy," *Journal of Political Economy*, 2019, *127* (6), 2926 – 2966.

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

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

Lepetyuk, Vadym, Lilia Maliar, and Serguei Maliar, "Should Central Banks Worry About Nonlinearities of their Large-scale Macroeconomic Models?," 2017. Bank of Canada WP.

Lucas, Robert E., "Liquidity and Interest Rates," *Journal of Economic Theory*, 1990, *50* (2), 237–264.

_ **and Nancy Stokey**, "Liquidity Crises," *Manuscript, University of Chicago*, 2011.

Lustig, Hanno N and Stijn G Van Nieuwerburgh, "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *The Journal of Finance*, 2005, *60* (3), 1167–1219.

Martin, Philippe and Thomas Philippon, "Inspecting the Mechanism: Leverage and the Great Recession in the Eurozone," *American Economic Review*, 2017, *107* (7), 1904–1937.

Mendoza, Enrique G, "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 2010, *100* (5), 1941–66.

Mian, Atif and Amir Sufi, "House prices, home equity-based borrowing, and the US household leverage crisis," *American Economic Review*, 2011, *101* (5), 2132–56.

_ **and** _ , "What explains the 2007–2009 drop in employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

_ **, Kamalesh Rao, and Amir Sufi**, "Household balance sheets, consumption, and the economic slump," *The Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

Nakamura, Emi and Jon Steinsson, "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 2014, *104* (3), 753–92.

Negro, Marco Del, Marc P Giannoni, and Frank Schorfheide, "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 168–96.

Rocheteau, Guillaume, Pierre-Olivier Weill, and Tsz-Nga Wong, "A tractable model of monetary exchange with ex post heterogeneity," *Theoretical Economics*, 2018, *13* (3), 1369–1423.

Shi, Shouyong, "Liquidity, Assets and Business Cycles," *Journal of Monetary Economics*, 2015, *70*, 116 – 132.

Simonovska, Ina and Michael E Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586–606.

Werning, Ivan, "Managing a Liquidity Trap: Monetary and Fiscal Policy," 2011. NBER WP 17344.