Household Leverage and the Recession
Appendix

Not for Publication

Callum Jones Virgiliu Midrigan Thomas Philippon

Contents

1 Likelihood Function

We use Bayesian likelihood methods to estimate the parameters of the island economy and the shocks. We use a panel dataset across states together with aggregate data and the ZLB. We formulate the state-space of the model so as to separate our estimation into a regional component and aggregate component and make it computationally feasible.

We discuss the likelihood function of the state/regional component and then the likelihood function of the aggregate component.

1.1 Likelihood of the State Component

We use Bayesian methods. We first log-linearize the model. The log-linearized model has the state space representation:

$$x_t = \mathbf{J} + \mathbf{Q}x_{t-1} + \mathbf{G}\varepsilon_t \tag{1}$$

$$z_t = \mathbf{H}x_t. \tag{2}$$

The state vector is x_t. The error is distributed $\varepsilon_t \sim N(0, \Omega)$ where Ω is the covariance matrix of ε_t. We assume no observation error of the data z_t.

Denote by ϑ the vector of parameters to be estimated. Denote by $\mathcal{Z} = \{z_\tau\}_{\tau=1}^{T}$ the sequence of $N_z \times 1$ vectors of observable variables, combined over states. By Bayes law, the posterior $\mathcal{P}(\vartheta \mid \mathcal{Z})$ satisfies:

$$\mathcal{P}(\vartheta \mid \mathcal{Z}) \propto L(\mathcal{Z} \mid \vartheta) \times \mathcal{P}(\vartheta).$$

With Gaussian errors ε_t, the likelihood function $L(\mathcal{Z} \mid \vartheta)$ is computed using the sequence of structural matrices and the Kalman filter, described below:

$$\log L(\mathcal{Z} \mid \vartheta) = -\left(\frac{N_z T}{2}\right) \log 2\pi - \frac{1}{2}\sum_{t=1}^{T} \log\det \mathbf{S}_t - \frac{1}{2}\sum_{t=1}^{T} \tilde{y}_t^\top \left(\mathbf{S}_t\right)^{-1} \tilde{y}_t.$$

where $\tilde{y}_t$ is the vector of forecast errors and $\mathbf{S}_t$ is its associated covariance matrix.

1.2 Kalman Filter

The Kalman filter recursion is given by the following equations. The state of the system is $(\hat{x}_t, \mathbf{P}_{t-1})$. In the predict step, the structural matrices $\mathbf{J}$, $\mathbf{Q}$ and $\mathbf{G}$ are used to compute a forecast of the state $\hat{x}_{t|t-1}$ and the forecast covariance matrix $\mathbf{P}_{t|t-1}$ as:

$$\begin{aligned} \hat{x}_{t|t-1} &= \mathbf{J} + \mathbf{Q}\hat{x}_t \\ \mathbf{P}_{t|t-1} &= \mathbf{Q}\mathbf{P}_{t-1}\mathbf{Q}^\top + \mathbf{G}\Omega\mathbf{G}^\top. \end{aligned} \tag{3}$$

We update these forecasts with imperfect observations of the state vector. This update step involves computing forecast errors $\tilde{y}_t$ and its associated covariance matrix $\mathbf{S}_t$ as:

$$\begin{aligned} \tilde{y}_t &= z_t - \mathbf{H}\hat{x}_{t|t-1} \\ \mathbf{S}_t &= \mathbf{H}\mathbf{P}_{t|t-1}\mathbf{H}^\top. \end{aligned}$$

The Kalman gain matrix is given by:

$$\mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{H}^\top\mathbf{S}_t^{-1}.$$

With $\tilde{y}_t$, $\mathbf{S}_t$ and $\mathbf{K}_t$ in hand, the optimal filtered update of the state x_t is

$$\hat{x}_t = \hat{x}_{t|t-1} + \mathbf{K}_t\tilde{y}_t,$$

and for its associated covariance matrix:

$$\mathbf{P}_t = \left(I - \mathbf{K}_t\mathbf{H}\right)\mathbf{P}_{t|t-1}.$$

The Kalman filter is initialized with x_0 and $\mathbf{P}_0$ determined from their unconditional moments, and is computed until the final time period T of data. We can show that the stationary $\mathbf{P}_0$ has the expression:

$$\operatorname{vec}(\mathbf{P}_0) = \left(\mathbf{I} - \mathbf{Q} \otimes \mathbf{Q}\right)^{-1} \operatorname{vec}(\mathbf{G}\Omega\mathbf{G}^\top) \tag{4}$$

1.3 Kalman Smoother

With the estimates of the parameters on a sample up to time period T, the Kalman smoother gives an estimate of $x_{t|T}$, or an estimate of the state vector at each point in time given all available information. With $\hat{x}_{t|t-1}$, $\mathbf{P}_{t|t-1}$, $\mathbf{K}_t$ and $\mathbf{S}_t$ in hand from the Kalman filter, the vector $x_{t|T}$ is computed by:

$$x_{t|T} = \hat{x}_{t|t-1} + \mathbf{P}_{t|t-1} r_{t|T},$$

where the vector $r_{T+1|T} = 0$ and is updated with the recursion:

$$r_{t|T} = \mathbf{H}^\top \mathbf{S}_t^{-1} \left(z_t - \mathbf{H}\hat{x}_{t|t-1} \right) + \left(I - \mathbf{K}_t \mathbf{H} \right)^\top \mathbf{P}_{t|t-1}^\top r_{t+1|T}.$$

Finally, to get an estimate of the shocks to each state variable under this model's shock structure, denoted by e_t, we can compute:

$$e_t = \mathbf{G}\varepsilon_t = \mathbf{G} r_{t|T}.$$

1.4 Block structure

The regional component of the model has a block structure separated by state. For example, consider two states so that the log-linearized state-space representation is:

$$\begin{bmatrix} x_t^1 \\ x_t^2 \end{bmatrix} = \begin{bmatrix} \mathbf{J}^1 \\ \mathbf{J}^2 \end{bmatrix} + \begin{bmatrix} \mathbf{Q}^1 & 0 \\ 0 & \mathbf{Q}^2 \end{bmatrix} \begin{bmatrix} x_{t-1}^1 \\ x_{t-1}^2 \end{bmatrix} + \begin{bmatrix} \mathbf{G}^1 & 0 \\ 0 & \mathbf{G}^2 \end{bmatrix} \begin{bmatrix} \varepsilon_t^1 \\ \varepsilon_t^2 \end{bmatrix}$$

Under this block structure, the forecast covariance matrix $P_{t|t-1}$ also has a block structure. This is clear from the expressions (3) and (4).

The block structure is helpful for computational reasons. The log-likelihood becomes a weighted sum of state-by-state log-likelihood functions. To show this: because $\mathbf{P}_{t|t-1}$ has a block structure, so does $\mathbf{S}_t$. And because $\mathbf{S}_t$ has a block structure:

$$\log \det \mathbf{S}_t = \log \prod_j \det \mathbf{S}_t^j = \sum_j \log \det \mathbf{S}_t^j.$$

Also, because $\mathbf{S}_t$ has a block structure, so does its inverse, so that the last term in the log-likehood can also be separated by state. The log-likelihood is then:

$$\log L(\mathcal{Z} \mid \vartheta) = \sum_s \log L^s(\mathcal{Z}^s \mid \vartheta).$$

1.5 Weighting

We weight the contributions of the state likelihoods to account for differences in the size of states. Weights can, in principle, depend on the sample and the model's parameters. Agostinelli and Greco[1] discuss the asymptotic properties of the weighting function which are needed for the weighted likelihood to share the same asymptotic properties as the genuine likelihood function. Using population weights for state subsamples which are constant over time is a simple weighting function which satisfies these properties.

[1] Agostinelli, Claudio and Greco, Luca, *Weighted Likelihood in Bayesian Inference*, 46th Scientific Meeting of the Italian Statistical Society,2012.

2 Likelihood of the Aggregate Component

2.1 Solution with Zero Lower Bound

We write the model that approximates the ZLB in a way that agents have time-varying beliefs about the evolution of the model's structural parameters $\mathbf{A}_t$, $\mathbf{B}_t$, $\mathbf{C}_t$, $\mathbf{D}_t$ and $\mathbf{F}_t$, where:

$$\mathbf{A}_t x_t = \mathbf{C}_t + \mathbf{B}_t x_{t-1} + \mathbf{D}_t \mathbb{E}_t x_{t+1} + \mathbf{F}_t \epsilon_t.$$

For example, if the ZLB binds, the equation describing the Taylor rule becomes $i_t = 0$, changing the structural matrices $\mathbf{A}_t$, and so on. With time-varying structural matrices, the solution becomes:

$$x_t = \mathbf{J}_t + \mathbf{Q}_t x_{t-1} + \mathbf{G}_t \epsilon_t, \tag{5}$$

where $\mathbf{J}_t$, $\mathbf{Q}_t$ and $\mathbf{G}_t$ are conformable matrices which are functions of the evolution of beliefs about the time-varying structural matrices $\mathbf{A}_t$, $\mathbf{B}_t$, $\mathbf{C}_t$, $\mathbf{D}_t$ and $\mathbf{F}_t$ (Kulish and Pagan, 2017). These matrices satisfy the recursion:

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

where the final structures $\mathbf{Q}_T$ and $\mathbf{J}_T$ are known and computed from the time invariant structure above under the terminal period's structural parameters, ie the no-ZLB case.

Given a sequence of ZLB durations, the state-space of the model is:

$$\begin{aligned}
x_t &= \mathbf{J}_t + \mathbf{Q}_t x_{t-1} + \mathbf{G}_t \varepsilon_t \\
z_t &= \mathbf{H}_t x_t.
\end{aligned}$$

The observation equation is time-varying because the nominal interest rate becomes unobserved when it is at its bound.

Denote by ϑ the vector of parameters to be estimated and by $\mathbf{T}$ the vector of ZLB durations that are observed each period. Denote by $\mathcal{Z} = \{z_\tau\}_{\tau=1}^{T}$ the sequence of vectors of observable variables. With Gaussian errors, the likelihood function $L(\mathcal{Z}, \mathbf{T} \mid \vartheta)$ for the aggregate component is computed using the sequence of structural matrices associated with the sequence of ZLB durations, and the Kalman filter:

$$\log L(\mathcal{Z}, \mathbf{T} \mid \vartheta) = -\left(\frac{N_z T}{2}\right) \log 2\pi - \frac{1}{2} \sum_{t=1}^{T} \log \det \mathbf{H}_t \mathbf{S}_t \mathbf{H}_t^\top - \frac{1}{2} \sum_{t=1}^{T} \tilde{y}_t^\top \left(\mathbf{H}_t \mathbf{S}_t \mathbf{H}_t^\top\right)^{-1} \tilde{y}_t.$$

2.2 Kalman filter

The state of the system is $(\hat{x}_t, \mathbf{P}_{t-1})$. In the predict step, the structural matrices $\mathbf{J}_t$, $\mathbf{Q}_t$ and $\mathbf{G}_t$ are used to compute a forecast of the state $\hat{x}_{t|t-1}$ and the forecast covariance matrix $\mathbf{P}_{t|t-1}$ as:

$$\begin{aligned}
\hat{x}_{t|t-1} &= \mathbf{J}_t + \mathbf{Q}_t \hat{x}_t \\
\mathbf{P}_{t|t-1} &= \mathbf{Q}_t \mathbf{P}_{t-1} \mathbf{Q}_{t|t-1}^\top + \mathbf{G}_t \Omega \mathbf{G}_t^\top.
\end{aligned}$$

This formulation differs from the time-invariant Kalman filter used at the state-level because in the forecast stage the matrices $\mathbf{J}_t$, $\mathbf{Q}_t$ and $\mathbf{G}_t$ can vary over time. We update these forecasts with imperfect observations of the state vector. This update step involves computing forecast errors $\tilde{y}_t$ and its associated covariance matrix $\mathbf{S}_t$ as:

$$\begin{aligned}\tilde{y}_t &= z_t - \mathbf{H}_t \hat{x}_{t|t-1} \\ \mathbf{S}_t &= \mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^\top.\end{aligned}$$

The Kalman gain matrix is given by:

$$\mathbf{K}_t = \mathbf{P}_{t|t-1} \mathbf{H}_t^\top \mathbf{S}_t^{-1}.$$

With $\tilde{y}_t$, $\mathbf{S}_t$ and $\mathbf{K}_t$ in hand, the optimal filtered update of the state x_t is

$$\hat{x}_t = \hat{x}_{t|t-1} + \mathbf{K}_t \tilde{y}_t,$$

and for its associated covariance matrix:

$$\mathbf{P}_t = \left(I - \mathbf{K}_t \mathbf{H}_t\right) \mathbf{P}_{t|t-1}.$$

The Kalman filter is initialized with x_0 and $\mathbf{P}_0$ determined from their unconditional moments, and is computed until the final time period T of data.

2.3 Kalman smoother

With the estimates of the parameters and durations in hand at time period T, the Kalman smoother gives an estimate of $x_{t|T}$, or an estimate of the state vector at each point in time given all available information (Hamilton, 1994). With $\hat{x}_{t|t-1}$, $\mathbf{P}_{t|t-1}$, $\mathbf{K}_t$ and $\mathbf{S}_t$ in hand from the Kalman filter, the vector $x_{t|T}$ is computed by:

$$x_{t|T} = \hat{x}_{t|t-1} + \mathbf{P}_{t|t-1} r_{t|T},$$

where the vector $r_{T+1|T} = 0$ and is updated with the recursion:

$$r_{t|T} = \mathbf{H}_t^\top \mathbf{S}_t^{-1} \left(z_t - \mathbf{H}_t \hat{x}_{t|t-1}\right) + \left(I - \mathbf{K}_t \mathbf{H}_t\right)^\top \mathbf{P}_{t|t-1}^\top r_{t+1|T}.$$

Finally, to get an estimate of the shocks to each state variable under this model's shock structure, denoted by e_t, we compute:

$$e_t = \mathbf{G}_t \varepsilon_t = \mathbf{G}_t r_{t|T}.$$

3 Posterior Sampler

This section describes the sampler used to obtain the posterior distribution of interest. We compute the likelihood function at the state-level and the aggregate level, together with the prior. The posterior of our full model $\mathcal{P}(\vartheta \mid \mathbf{T}, \mathcal{Z})$ satisfies:

$$\mathcal{P}(\vartheta \mid \mathbf{T}, \mathcal{Z}) \propto L(\mathcal{Z}, \mathbf{T} \mid \vartheta) \times \mathcal{P}(\vartheta).$$

Our priors are diffuse and are the same, or slightly wider, than those used in Smets and Wouters (2007).

We use a Markov Chain Monte Carlo procedure to sample from the posterior. It has a single block, corresponding to the parameters ϑ.[2] The sampler at step j is initialized with the last accepted draw of the structural parameters ϑ_{j-1}.

First start by selecting which parameters to propose new values. For those parameters, draw a new proposal ϑ_j from a proposal density centered at ϑ_{j-1} and with thick tails to ensure sufficient coverage of the parameter space and an acceptance rate of roughly 20% to 25%. The proposal ϑ_j is accepted with probability $\frac{\mathcal{P}(\vartheta_j|\mathbf{T},\mathcal{Z})}{\mathcal{P}(\vartheta_{j-1}|\mathbf{T},\mathcal{Z})}$. If ϑ_j is accepted, then set $\vartheta_{j-1} = \vartheta_j$.

4 Data

4.1 State-level

We use state-level data on employment, consumption spending, wages, debt-to-income, and house prices. The observed state data is annual. The model is quarterly. So we used a mixed-frequency estimation procedure. The data is only used to compute forecast errors in the first quarter of the year.

To construct the data, we first take each state's series relative to its 2001 value, compute the devation of each state's observation from the state mean, regress that series on time dummies, weighted by the state's relative population, and work with the residuals. As discussed in the text, this formulation allows us to separate the state-based and aggregate-components of the state-space model. The resulting series for each state from 2001 to 2012 are plotted in Figures 1 to 5.

Here, we provide more details on each series:

- Consumption: We use state-level data on Total Personal Consumption Expenditures by State from the BEA, net of healthcare and housing. The data is available for download at the BEA website.

- Employment: We use state-level data on Total Employment net of employment in the construction sector from the BEA annual table SA4. In our empirical analysis we scale this measure of employment by each state's population.

- Population: We use state-level data on Population from the BEA annual table SA1-3.

- Labor Compensation: We use state-level data on Compensation of Employees by Industry from the BEA annual table SA6N.

- Wages: We divide total labor compensation by the number of employed individuals using the two series described above.

- Income: We use state-level data on Personal Income from the BEA annual table SA4.

- Household Debt: We use data from the FRBNY Consumer Credit Panel Q4 State statistics by year. Our measures of debt include auto loans, credit card debt, mortgage debt and student loans. This database also provides information on the number of

[2] It is worth noting that, as in Kulish, Morley and Robinson (2017), one can estimate in addition to the structural parameters ϑ, the expected zero lower bound durations can be estimated together with the structural parameters, in which case an additional block is needed in the posterior sampler.

individuals with credit scores in each state, which we use to express the debt data in per-capita terms. We then construct a debt-to-income series by dividing this measure of per-capita debt by per-capita income using the data described above on income and population from the BEA.

- House Prices: We used data on the Not Seasonally Adjusted House Price Index available on the FHFA website.

4.2 Aggregate level

Inflation, employment, output, household debt, house prices, wages, Fed Funds rates, ZLB durations from NY Federal Reserve Survey Data. Construction of the data series largely follows Smets and Wouters (2007). The codes for each raw data series are as follows:

- Gross Domestic Product: Implicit Price Deflator (GDPDEF). Index, 2009=100.
- Personal Consumption Expenditures (PCEC). Current, $.
- Nonfarm Business Sector: Real Compensation Per Hour (COMPRNFB). Index, 2009=100.
- Civilian Employment (CE16OV), Thousands of Persons.
- Total population (CNP16OV), Thousands of Persons.
- Nonfarm Business Sector: Average Weekly Hours (PRS85006023). Index, 2009=100.
- Civilian Noninstitutional Population. (CNP16OV). Thousands of Persons.
- Household Debt.
- House Prices.

To map these data series to the model variables, we do the following transformations.

1. Construct the series LNSindex, which is an index of CNP16OV where 1992Q3=1. We adjust the CNP16OV series to account for breaks in the series each January, due to revisions from updated Census reports, which can be substantial. To do this, we impute an estimate of each January's monthly change in population and construct an estimate of the revised change in population from the actual change to the constructed imputed change. We then distribute that revised change in population across the preceding 12 months.
2. Construct the series CE16OVIndex, which is an index of CE16OV where 1992Q3=1.
3. Inflation = Π_t = ln(GDPDEF / GDPDEF(-1)) * 100.
4. Consumption = C_t = ln((PCEC / GDPDEF) / LNSindex) * 100. Then compute the percentage change in consumption as an observable, $\ln C_t - \ln C_{t-1}$.

5. Real wages = w_t = ln(COMPRNFB) * 100.[3] Then compute the percentage change in real wages as an observable, $\ln w_t - \ln w_{t-1}$.

6. Labor supply = ln(PRS85006023 * CE16OVIndex) / LNSIndex. Then take as an observable n_t as the value of labor supply relative to its 1948Q1 to 2015Q1 mean.

Household debt. U.S. household debt to income ratio exhibits a trend, starting from about 0.5 in 1975 to about 1 in the last decade. Since we do not allow for trends in our model, we de-trend the data by subtracting a linear trend. We smooth this series to eliminate high frequency noise, by projecting it on a cubic spline of order 15 — the smoothed series is reported with dotted lines in the figure.

Fed Funds rate: The interest rate is the Federal Funds Rate, taken from the Federal Reserve Economic Database.

ZLB Durations: We follow the approach of Kulish, Morley and Ronbinson (2017) and use the ZLB durations extracted from the New York Federal Reserve Survey of Primary Dealers, conducted eight times a year, from 2011Q1 onwards.[4] We take the mode of the distribution implied by these surveys. Before 2011, we use responses from the Blue Chip Financial Forecasts survey.

5 Priors and Posteriors

5.1 Priors

For the persistence and standard deviation of the AR(1) shocks, we use the same priors as Smets and Wouters (2007) use in their aggregate estimation. The persistence parameters are centered around 1/2 and are diffuse. We use wide priors on the standard deviations of the shocks.

For the wage and price Calvo stickiness parameters, we use a slightly more diffuse prior than Smets and Wouters (2007) but centered around the same mode of $\lambda_p = \lambda_w = 1/2$. This is because we use a more recent sample and wider priors are consistent with a flattening of the Phillips Curve.

For the degree of idiosyncratic uncertainty α, we use a wide prior, centered around a value of 2.5. As the first panel of Figure 7 shows, this prior is wide enough to allow the data to find strong or very weak effects of credit shocks, as we discuss in the paper. We have tried estimations with uniform priors over α and find very similar results.

5.2 Posteriors

Kernel density estimates for the posterior distributions are plotted in red in Figure 7. They display single peaks, and are largely different from the priors, suggesting they are well identified.

[3] Note, this series differs marginally from the series Smets and Wouters (2007) use for real wages, which is defined as the Bureau of Labor Statistics' series PRS85006103 - Nonfarm Business, All Persons, Hourly Compensation Duration, divided by the GDP deflator GDPDEF. COMPRNFB was convenient because it was readily available on the Federal Reserve Economic Database.

[4] See the website here. For example in 2011, the survey conducted on January 18, one of the questions asked was: "Of the possible outcomes below, please indicate the percent chance you attach to the timing of the first federal funds target rate increase." (Question 2b). Responses were given in terms of a probability distribution across future quarters.

The convergence of the posterior distributions for each parameter is analyzed in Figure 8. Here we plot the Gelman-Rubin convergence diagnostic, along the length of the chain on the x-axes, up to the maximum length 150,000. Intuitively, the diagnostic compares the estimated between-chain and within-chain variances. As originally suggested by Brooks and Gelman (1997), values of the diagnostic below 1.2 suggest convergence. We find values well below 1.2 for our baseline posterior estimates with a chain of 150,000.

We discard half of the chain as a burn-in.

6 Theoretical Variance Decompositions

The theoretical forecast error variance decompositions at the 2Q, 4Q, and 12Q horizons are shown in Tables 2 to 4. Credit shocks account for a sizeable fraction of the differential changes in employment and consumption at the state-level. The aggregate component of credit shocks has almost no role in accounting for aggregate consumption or employment.

Credit shocks play little role in accounting for high frequency changes in the Fed Funds rate, but do become more important at longer horizons.

7 Robustness Estimations

We report two further estimations for robustness purposes, in addition to those reported in the text.

Aggregate Data We estimate the model using aggregate data only to identify the structural parameters α, λ_p, and λ_w. Regional data is only used to estimate the regional shock processes. Aggregate data is also used to identify the aggregate shock processes. The results show that α is lower at 2.5 than our baseline estimate of 2.7. We also find even greater degrees of price stickiness relative to our baseline, with λ_p centered around 0.97 compared to 0.96 in our baseline, and λ_w centered aorund 0.93, compared to 0.88 in our baseline.

The first panel of Figure 9 compares what the model predicts when only credit shocks are used to compute the state-by-state series for differential changes in employment. Because our estimate of the degree of idiosyncratic uncertainty is higher, as is our estimates of the stickiness of prices and wages, we find that credit shocks generate an even stronger correlation between our counterfactual and the data, with an estimated slope of 0.7 in the boom, compared to 0.6 in our baseline, and a slope of 0.9 in the bust, compared to 0.7 in our baseline.

The aggregate consequences under this estimation are shown in the first panel of Figure 10. Credit shocks have a larger effect, owing to the lower estimate of α and higher degree of price and wage stickiness.

Estimated Taylor Rule We also estimated the Taylor Rule parameters, using the same priors as in Smets and Wouters (2007). We find very similar estimates to Justiniano and Primiceri (2008). As a result, as discussed in the text, the regional and aggregate results are largely the same as found in our baseline.

Table 1: Estimated Parameters, Robustness

	Prior				A. Aggregate Data			B. Taylor Rule		
Parameter	Dist	Median	10%	90%	Mode	10%	90%	Mode	10%	90%
A. Structural Parameters										
α	N	2.6	1.5	3.8	2.53	2.38	2.78	2.86	2.61	3.12
λ_p	B	0.5	0.2	0.8	0.97	0.95	0.98	0.95	0.94	0.96
λ_w	B	0.5	0.2	0.8	0.93	0.89	0.94	0.86	0.83	0.89
α_r	B	0.8	0.6	0.9	-	-	-	0.73	0.67	0.78
α_p	N	1.5	1.1	2.0	-	-	-	1.72	1.38	2.06
α_x	N	0.1	0.0	0.2	-	-	-	0.29	0.24	0.36
α_y	N	0.1	0.0	0.2	-	-	-	0.05	0.04	0.09
B. Regional Shock Processes										
ρ_z	B	0.5	0.2	0.8	0.91	0.66	0.94	0.89	0.26	0.91
ρ_m	B	0.5	0.2	0.8	0.73	0.47	0.82	0.60	0.37	0.82
ρ_h	B	0.5	0.2	0.8	0.86	0.73	0.91	0.86	0.79	0.92
ρ_n	B	0.5	0.2	0.8	0.82	0.55	0.87	0.64	0.46	0.77
ρ_b	B	0.5	0.2	0.8	0.93	0.81	0.95	0.92	0.84	0.95
$100 \times \sigma_z$	IG	0.6	0.3	1.9	0.60	0.48	1.29	0.61	0.55	2.62
σ_m	IG	0.6	0.3	1.9	0.70	0.54	1.33	0.67	0.55	1.91
σ_h	IG	0.6	0.3	1.9	0.51	0.34	1.45	0.67	0.36	1.12
σ_n	IG	0.6	0.3	1.9	2.39	1.68	7.13	2.94	1.09	3.58
$1000 \times \sigma_b$	IG	0.6	0.3	1.9	0.91	0.69	3.05	0.85	0.68	2.37
C. Aggregate Shock Processes										
ρ_z	B	0.5	0.2	0.8	0.97	0.94	0.98	0.97	0.94	0.98
ρ_m	B	0.5	0.2	0.8	0.97	0.95	0.98	0.96	0.94	0.97
ρ_h	B	0.5	0.2	0.8	0.94	0.92	0.95	0.95	0.92	0.96
ρ_n	B	0.5	0.2	0.8	0.05	0.04	0.18	0.12	0.06	0.21
ρ_b	B	0.5	0.2	0.8	0.88	0.86	0.90	0.89	0.88	0.92
ρ_p	B	0.5	0.2	0.8	0.54	0.43	0.69	0.51	0.48	0.70
$100 \times \sigma_z$	IG	0.6	0.3	1.9	0.62	0.55	0.69	0.59	0.55	0.70
σ_m	IG	0.6	0.3	1.9	0.18	0.14	0.23	0.20	0.16	0.28
σ_h	IG	0.6	0.3	1.9	0.14	0.11	0.20	0.15	0.11	0.21
$\frac{1}{100} \times \sigma_n$	IG	0.6	0.3	1.9	0.59	0.53	0.68	0.21	0.15	0.39
$1000 \times \sigma_b$	IG	0.6	0.3	1.9	1.58	1.36	1.89	1.19	0.88	1.49
$1000 \times \sigma_p$	IG	0.6	0.3	1.9	0.91	0.64	1.13	0.87	0.62	1.08
$100 \times \sigma_r$	IG	0.6	0.3	1.9	1.19	1.06	1.39	0.95	0.82	1.11

Table 2: Unconditional Variance Decomposition, 2 Quarter Horizon

Variable \ Shock	Collateral	Housing	Productivity	Leisure	Discount	Policy	Markup
A. State-level							
Spending	38.5	0.4	0.3	2.4	58.4		
Employment	23.3	0.2	38.2	2.4	35.9		
Wages	0.3	0.0	0.0	99.6	0.1		
Debt-to-income	99.8	0.2	0.0	0.0	0.0		
House prices	5.0	94.2	0.0	0.2	0.6		
B. Aggregate-level							
Consumption	0.1	0.0	0.1	0.2	45.0	49.4	5.3
Employment	0.0	0.0	32.5	0.1	30.4	33.4	3.6
Wages	0.0	0.0	0.0	93.5	0.0	0.0	6.5
Debt-to-income	51.5	0.3	0.5	1.0	45.6	0.8	0.2
House prices	0.1	85.4	0.0	0.0	7.4	6.5	0.7
Fed Funds rate	0.0	0.0	0.0	0.0	90.9	8.2	0.7
Inflation	1.1	0.0	1.9	4.0	0.1	0.0	92.8

Table 3: Unconditional Variance Decomposition, 4 Quarter Horizon

Variable \ Shock	Collateral	Housing	Productivity	Leisure	Discount	Policy	Markup
A. State-level							
Spending	43.0	0.4	0.3	2.8	53.4		
Employment	26.8	0.2	35.0	3.7	34.2		
Wages	0.4	0.0	0.0	99.5	0.1		
Debt-to-income	99.7	0.3	0.0	0.0	0.0		
House prices	5.3	93.8	0.0	0.2	0.7		
B. Aggregate-level							
Consumption	0.0	0.0	0.3	0.6	44.8	46.6	7.7
Employment	0.0	0.0	31.4	0.4	30.8	32.1	5.3
Wages	0.0	0.0	0.1	90.1	0.0	0.0	9.7
Debt-to-income	47.1	0.2	0.7	1.4	49.4	1.0	0.3
House prices	0.0	86.5	0.0	0.0	5.9	6.5	1.0
Fed Funds rate	0.3	0.0	0.0	0.1	89.3	9.2	1.0
Inflation	2.1	0.0	3.1	6.6	0.1	0.1	88.0

Table 4: Unconditional Variance Decomposition, 12 Quarter Horizon

Variable \ Shock	Collateral	Housing	Productivity	Leisure	Discount	Policy	Markup
A. State-level							
Spending	53.6	0.6	0.5	4.7	40.6		
Employment	33.9	0.4	25.8	12.4	27.6		
Wages	1.0	0.0	0.0	98.9	0.1		
Debt-to-income	99.4	0.6	0.0	0.0	0.0		
House prices	6.6	91.7	0.1	0.5	1.2		
B. Aggregate-level							
Consumption	0.1	0.0	1.5	3.3	45.4	41.3	8.4
Employment	0.1	0.0	29.9	2.3	32.3	29.4	6.0
Wages	0.0	0.0	0.7	85.0	0.2	0.2	13.9
Debt-to-income	40.7	0.4	1.4	2.9	52.7	1.3	0.6
House prices	0.0	88.7	0.2	0.4	3.4	6.2	1.2
Fed Funds rate	5.9	0.1	0.3	0.6	80.2	11.6	1.3
Inflation	5.8	0.0	5.5	11.6	0.3	0.2	76.6

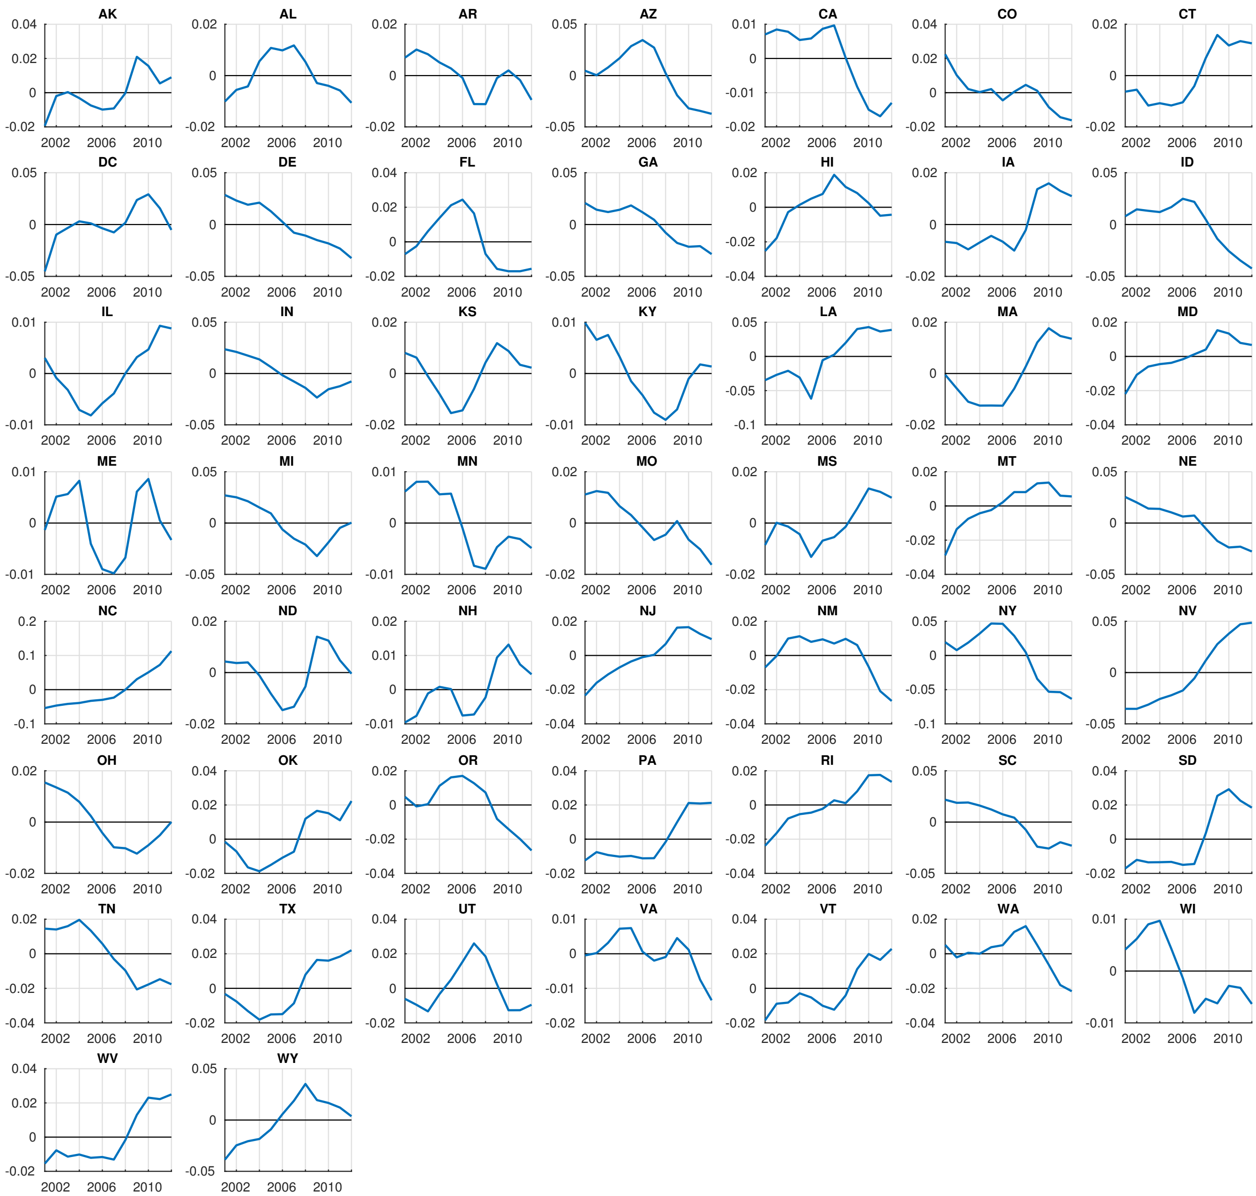

Figure 1: State Data: Employment

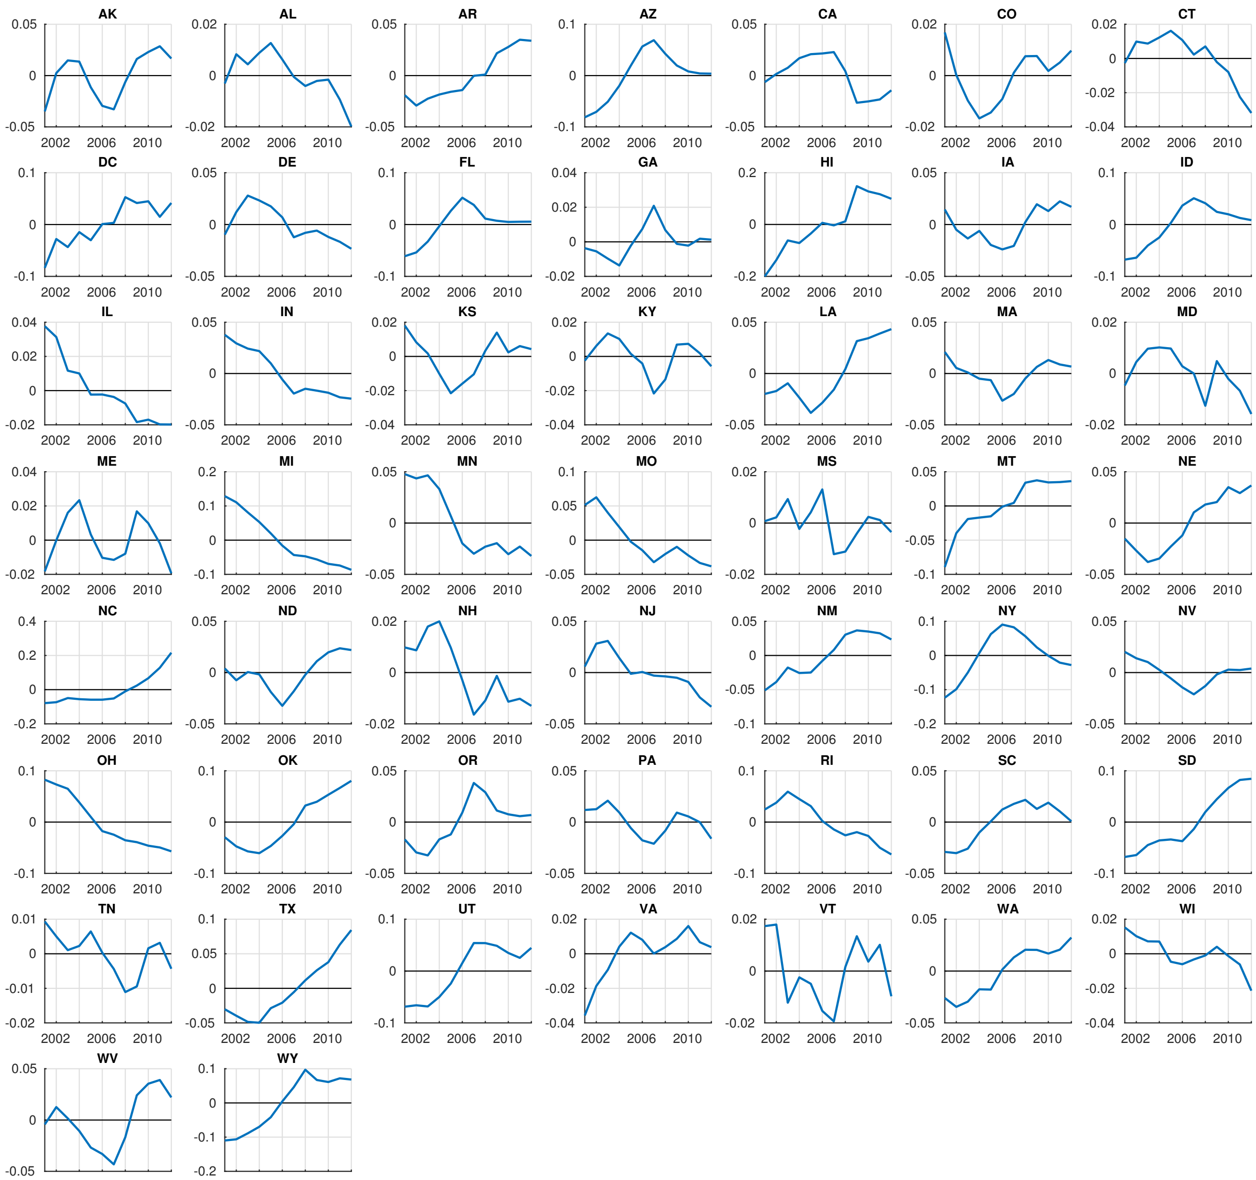

Figure 2: State Data: Household Spending

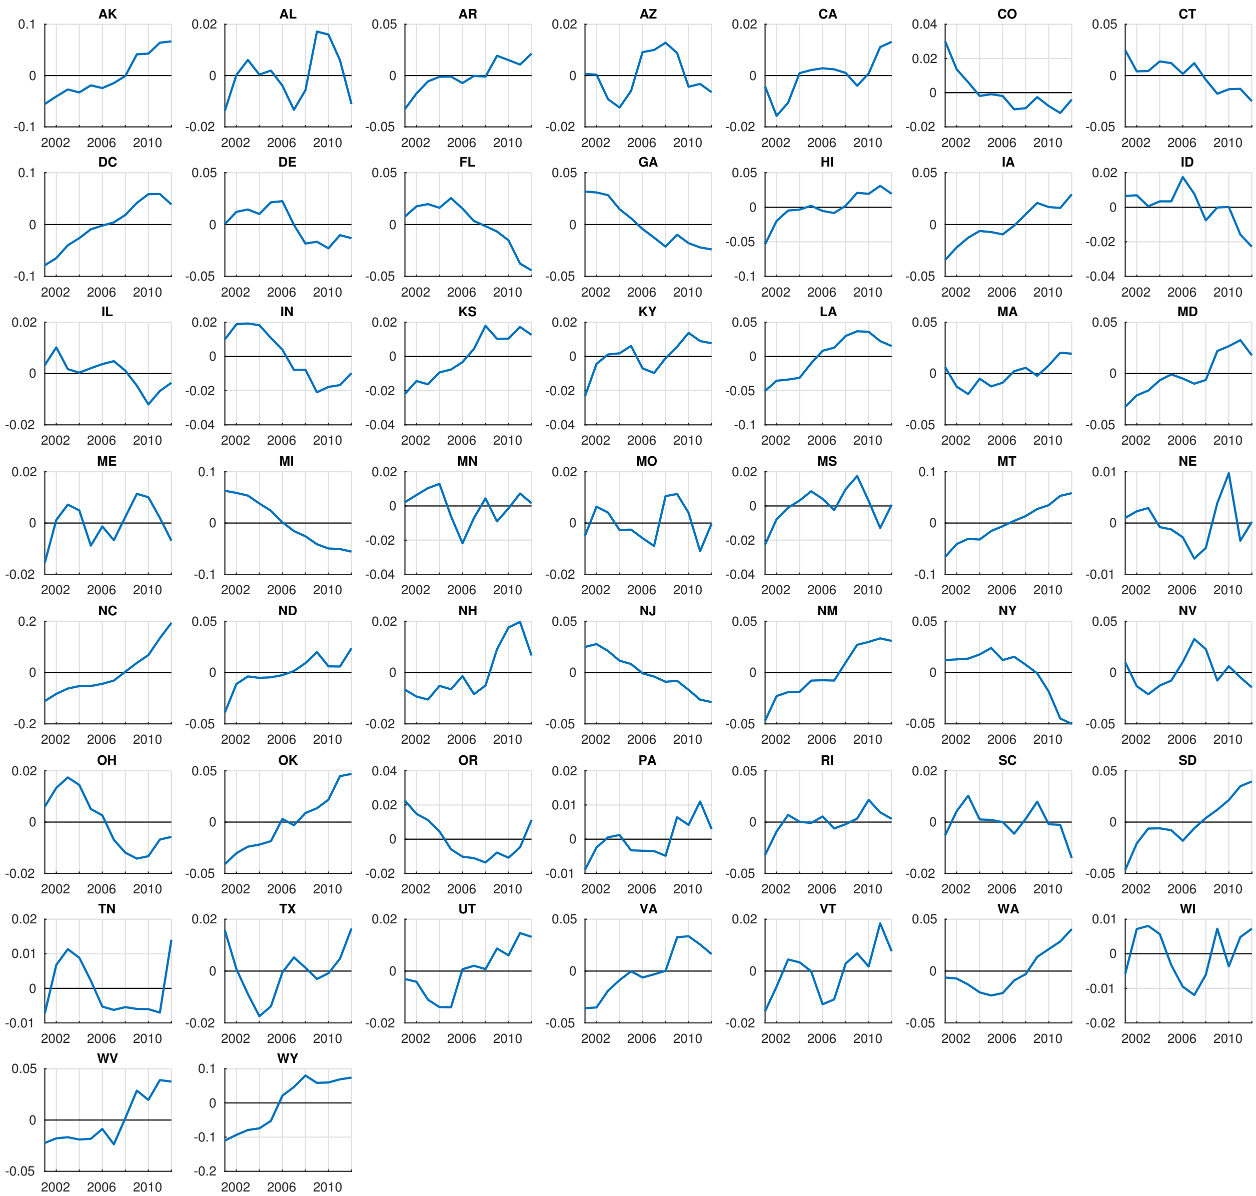

Figure 3: State Data: Wages

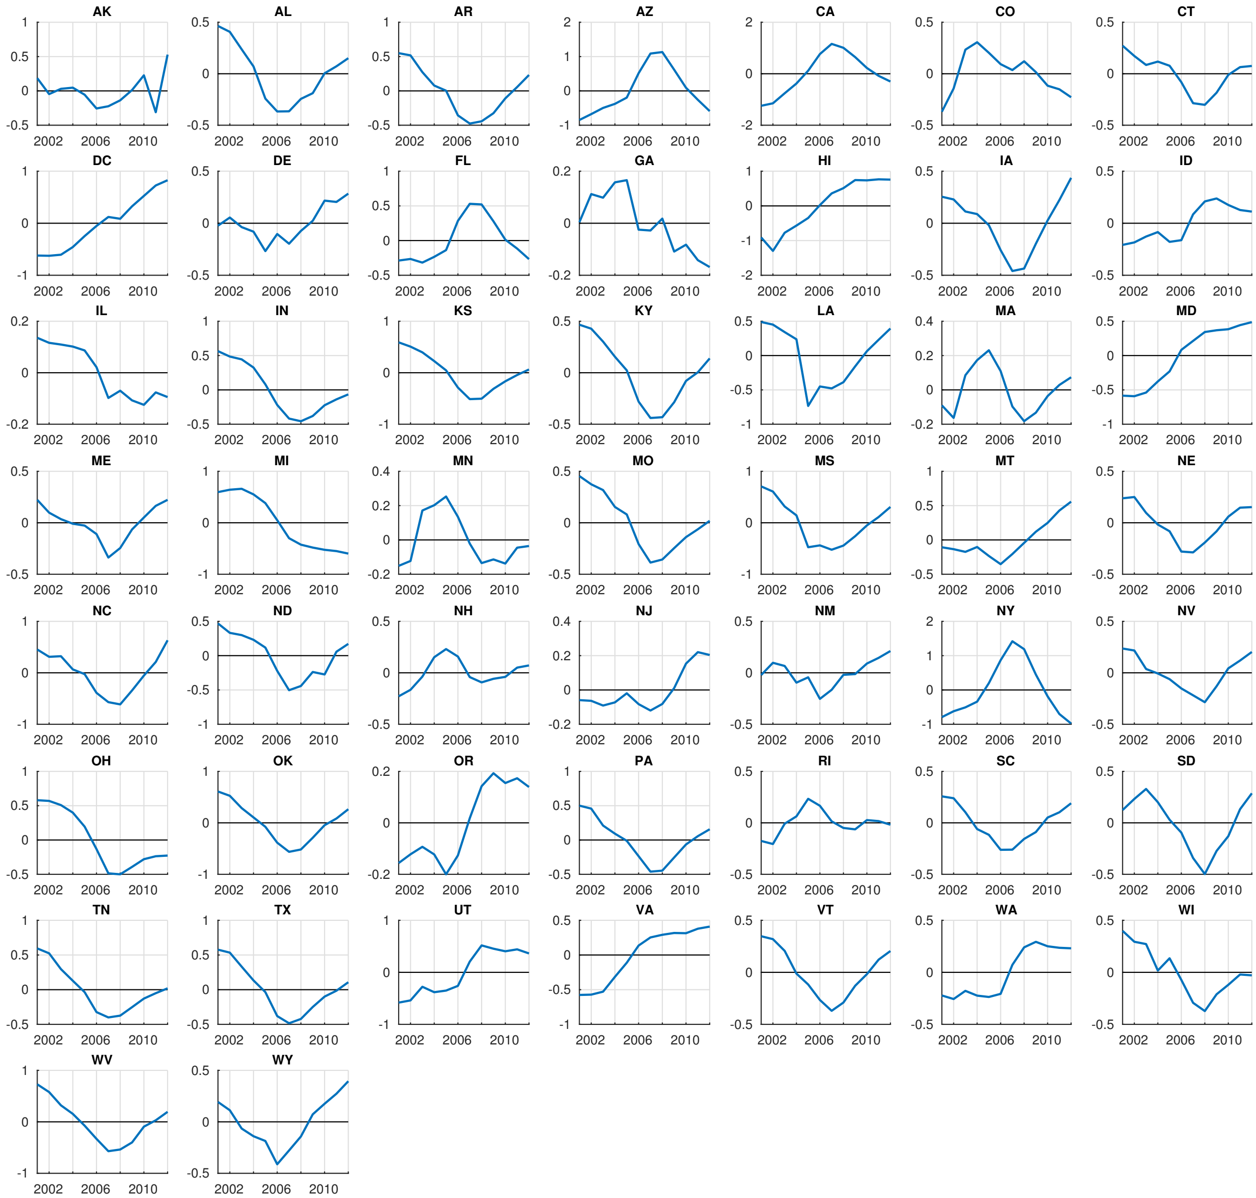

Figure 4: State Data: Household Debt

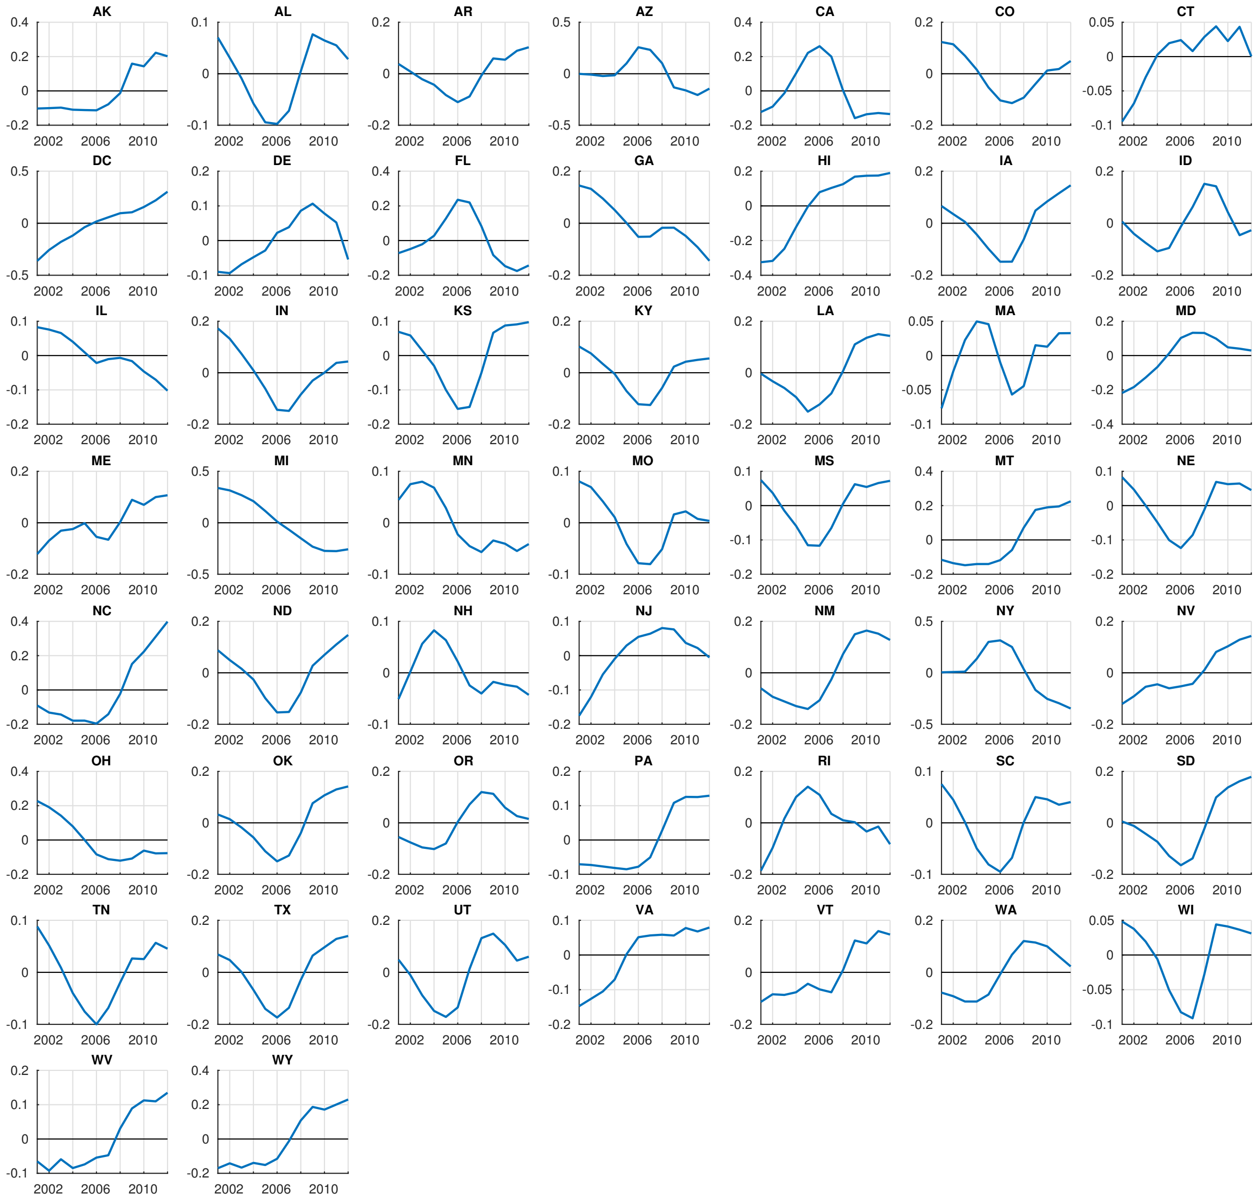

Figure 5: State Data: House Prices

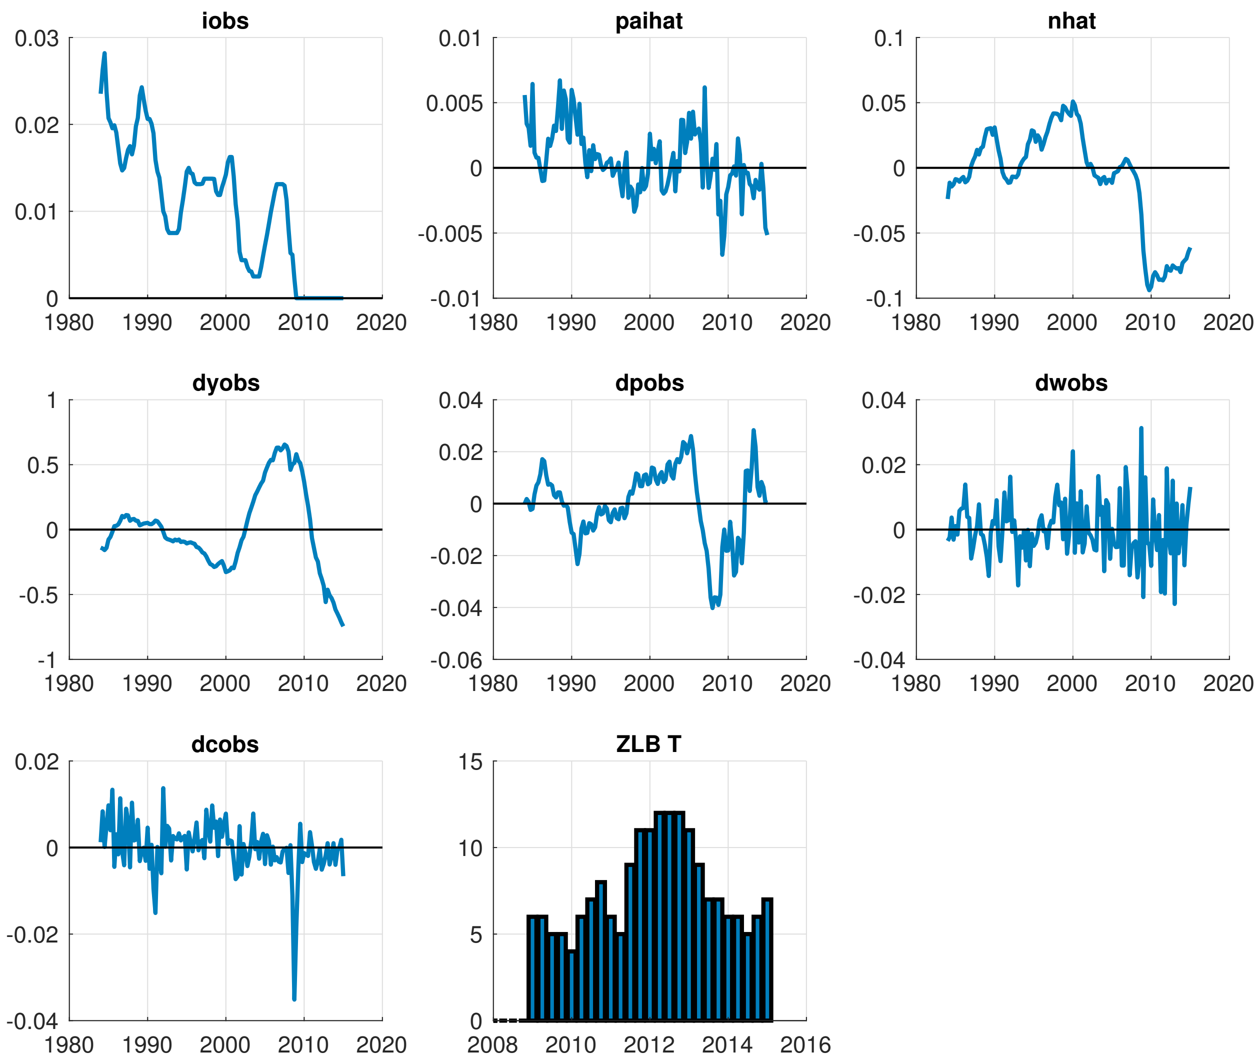

Figure 6: Aggregate Data

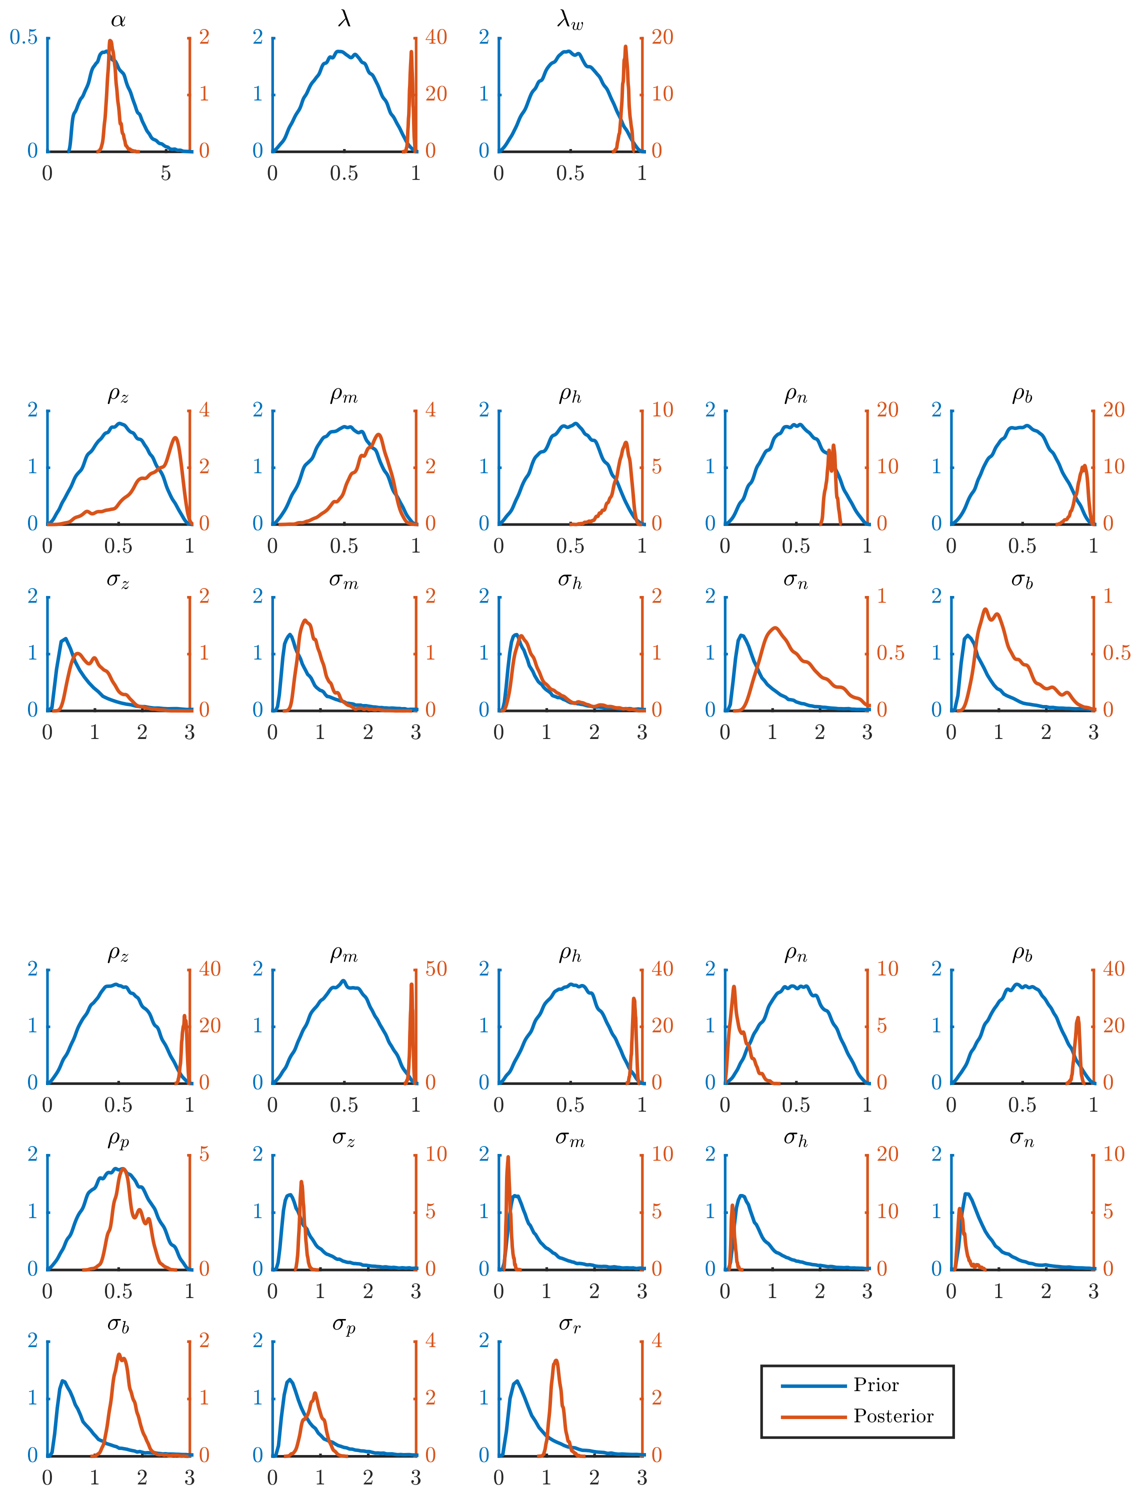

Figure 7: Prior and Posterior Distributions in Baseline Estimation

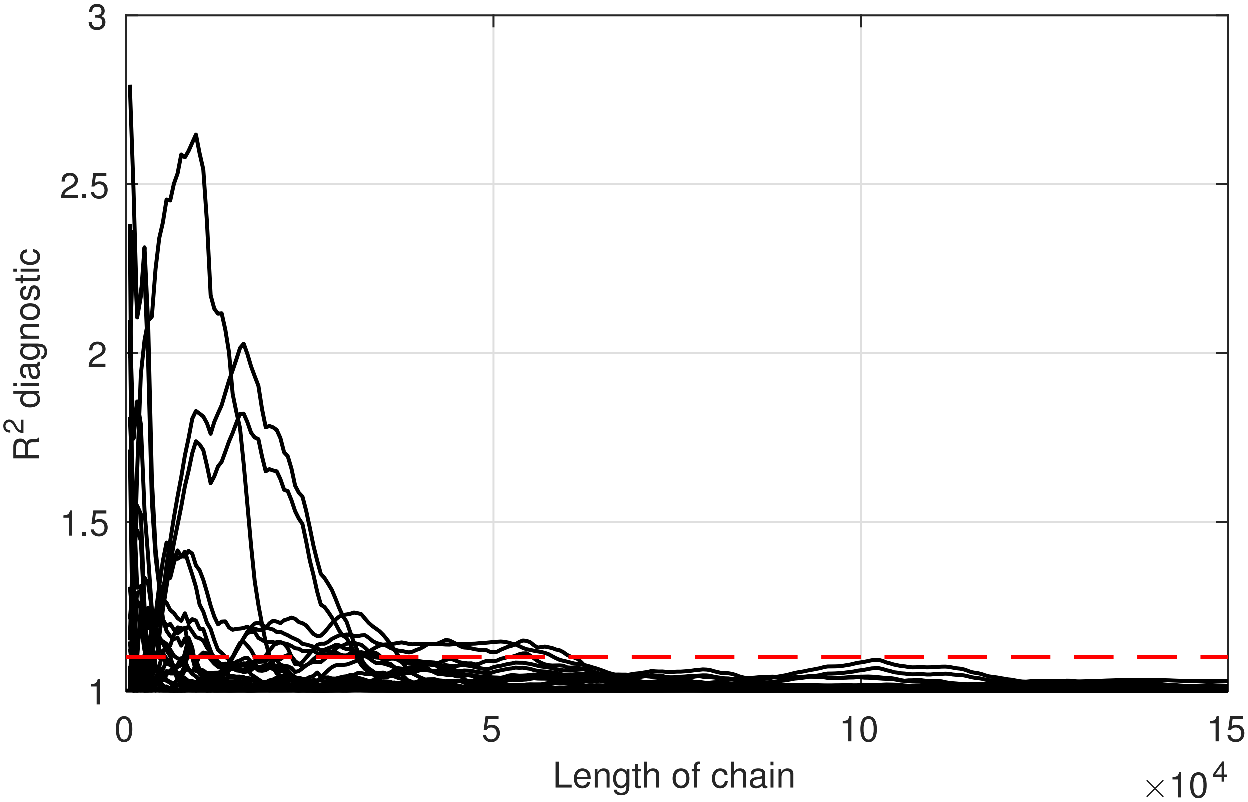

Figure 8: Convergence of Parameter Posterior Distributions Across MCMC Chains

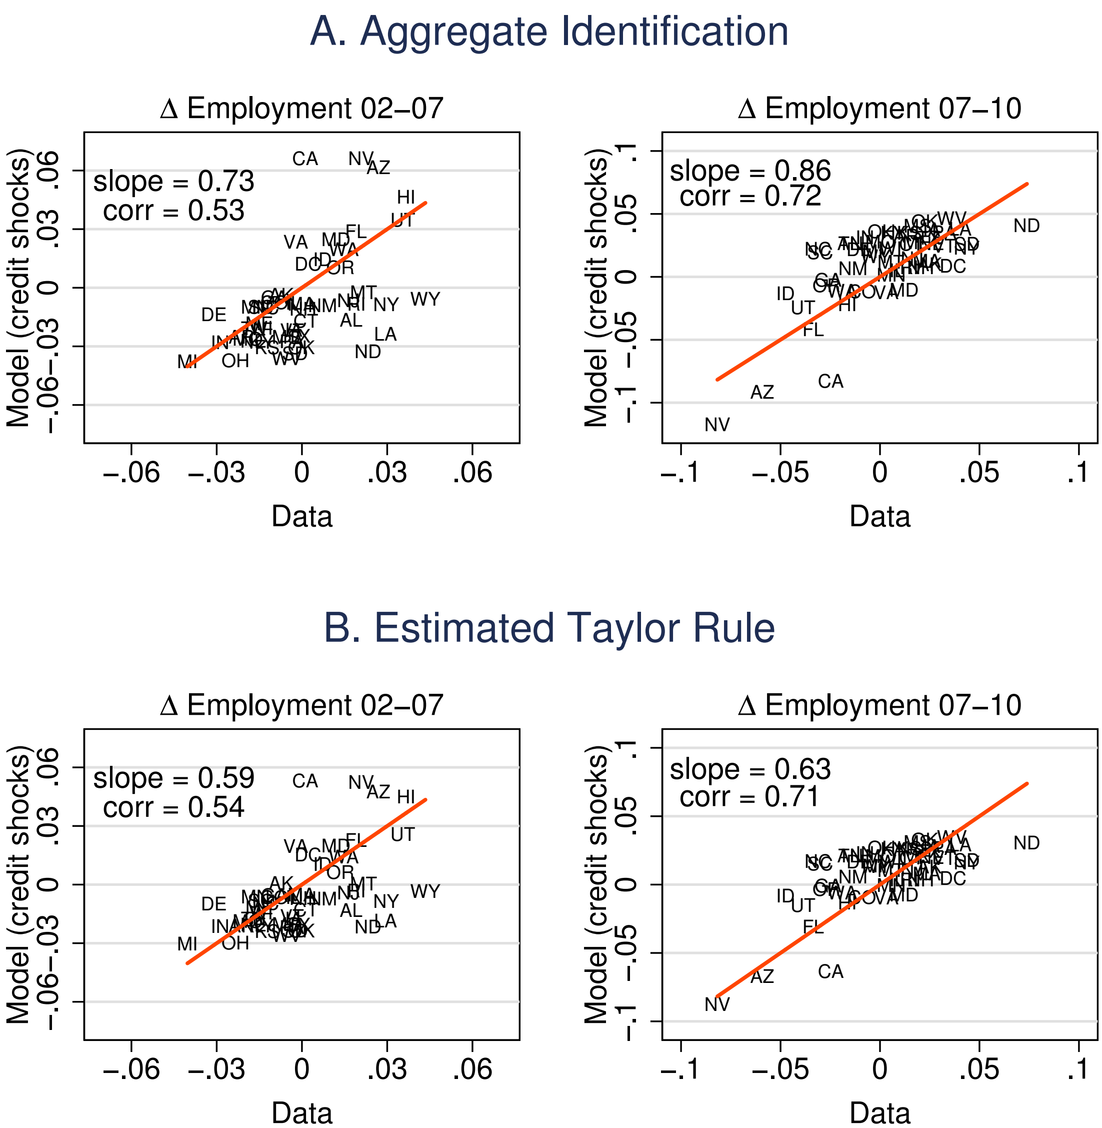

Figure 9: Effects of Credit Shocks at the Regional Level: Robustness

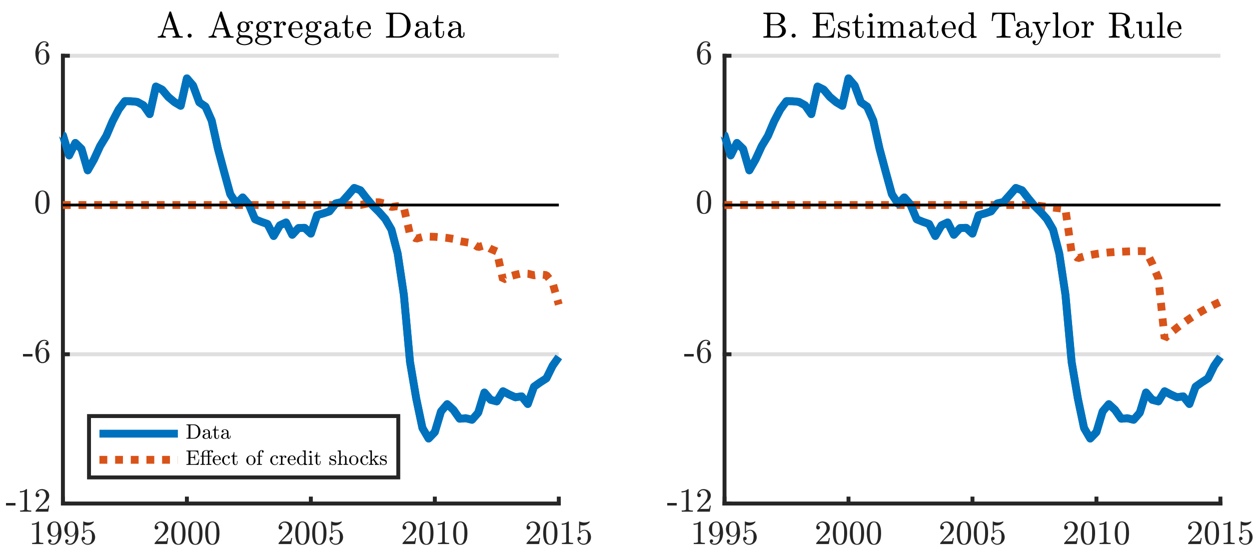

Figure 10: Effects of Credit Shocks at the Aggregate Level: Robustness