Household Leverage and the Recession

Appendix (not for publication)

Virgiliu Midrigan[*] Thomas Philippon[†]

May 2016

Contents

[*]New York University, virgiliu.midrigan@nyu.edu
[†]New York University, tphilipp@stern.nyu.edu

This appendix describes the data we have used, provides a more detailed discussion of the workings of our model, a more detailed discussion of how each parameter is identified in our indirect inference approach, as well as the details underlying the various robustness checks we have conducted.

A Data

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 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.

Figure A1 shows that the 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. The resulting series is shown in the right panel of Figure A1. 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.

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

Expected Duration of ZLB We downloaded this dataset from Bloomberg under the code MSM1KE. According to Bloomberg, this series uses Fed Funds futures contracts to measure the market's expected timing of rate hikes. The original series is monthly, so we have appropriately adjusted it to a quarterly frequency. Figure A2 shows this series, along with the model's forecasts of the duration of the ZLB when we impose a spread shock.

B Identification of Key Parameters

We next provide some intuition for what features of the data pin down individual parameters of our model by perturbing each of the six parameters and reporting how the model's fit changes. We report the results of these experiments in Table A1.

We first reduce α to 4, thus increasing the volatility of the taste shocks. Table A1 shows that all state-level aggregates are now much more volatile, with employment volatility increasing most (by about 80%), while house price volatility increasing least (about 40%). In contrast, the correlation between the model and data-produced series continues to be fairly high, in excess of 0.9. Thus, increasing the uncertainty about taste shocks makes state-level variables more sensitive to credit shocks. To understand this result, recall our discussion of the impulse responses to credit shocks in the main text. When uncertainty about taste shocks is high, the household finds it costly to reduce its level of assets in order to respond to a tightening of credit, and thus its consumption falls by more. Due to price and wage rigidities, more volatile state-level consumption fluctuations translate into more volatile employment and wage fluctuations as well.

We next increase γ to 0.975, thus increasing the duration of securities to 32 quarters. Table A1 shows that this perturbation again increases the volatility of state-level variables, more so for house prices (which are now 60% more volatile) than real variables which are now 20-30% more volatile. Intuitively, a greater γ implies that a current credit tightening will have effects on the household's ability to borrow far into the future, thus leading to a greater reduction in the collateral value of housing on impact. Importantly, increasing the duration of securities counterfactually reduces the correlation between house prices in the model and in the data. When γ is relatively high, house prices fall much more rapidly in response to a credit tightening than debt does, reducing the correlation between house prices and debt. Increasing the persistence of shocks, ρ, has a qualitatively similar impact as increasing γ and to conserve on space we do not report the results of these experiments here.

We next reduce the elasticity of substitution of labor varieties, ψ, to 3. We now see that wages are about 25% more volatile compared to the Benchmark model or the data, while employment is about 10% less volatile. Clearly, this parameter is identified based on the relative variability of wages and employment in the data.

Finally, we reduce the share of non-tradable goods in preferences to $\omega = 0.75$, thus decreasing the fraction of employment in the non-tradable goods sectors. We now find that employment, consumption and wages are all much less volatile (about 1/3 less volatile than in the data or the Benchmark) model. Intuitively, the larger the fraction of the workforce that produces tradable goods, the more insulated a state is from state-specific credit shocks.

C Workings of The Model

Here we discuss in greater detail the impulse responses to an economy-wide credit shock under the assumption that monetary policy follows a Taylor rule.

C.1 Benchmark Model

As we show in the paper, when prices are flexible, a credit tightening in the aggregate leads to a reduction in real interest rates and essentially no change in other real variables, such as employment. As is well known, a policy of strict inflation targeting would mimic the flexible price responses even in the presence of price and wage rigidities. Such a policy would ensure that the real interest rate in the sticky price economy mimics that in the flexible price economy, the latter often referred to as the *natural interest rate*, which we denote by r_t^n.

When prices are sticky and monetary policy does not fully offset the credit shock by a policy of strict inflation targeting, the real interest rate is too high compared to that under the flexible price allocations. If this is the case, consumption falls, as Figure A3 shows, as does output and employment (the two are equal up to a first-order approximation). The drop in consumption is larger, the greater the gap between the real interest rate and its natural level. Interestingly, even though we have assumed a unitary elasticity of intertemporal substitution, the decline in consumption in our model is not equal to the expected value of the sum of future interest rate gaps, $\mathbb{E}_t \sum_j \left(r^n_{t+j} - r_{t+j}\right)$, as would be the case in simpler versions of the New Keynesian models in which shocks to r^n_t are assumed exogenous. This point is illustrated by the upper-right panel of Figure A3 which shows that this sum falls by about twice more than consumption does.

The lower panels of Figure A3 provide some additional description of the mechanics behind the drop in consumption in the economy with price rigidities. We can decompose aggregate consumption as the product of two terms: the minimum amount $\underline{c}_t$ consumed by the $v = 1$ members of the household, times the mean-min consumption ratio. The minimum consumption level is given by

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

where recall that μ_t is the shadow value of wealth and π_{t+1} is the inflation rate. Similarly, the mean-min consumption ratio is given by

$$\text{mean-min consumption ratio} = \frac{\alpha}{\alpha - 1} \left(1 - \frac{1}{\alpha} \left(\frac{x_t}{\underline{c}_t}\right)^{1-\alpha}\right).$$

The lower panel of Figure A3 shows that with sticky prices μ_t increases in response to the credit shock, thus leading agents to reduce the minimum level of consumption. This effect is amplified by the tightening of liquidity constraints associated with the lower real interest rates which makes it optimal for households to reduce the amount transferred to individual members, $x_t/\underline{c}_t$, thus reducing the mean-min consumption ratio. This latter effect is more modest, however, as shown by the lower-right panel of Figure A3, and accounts for only about a quarter of the overall drop in consumption. Absent price rigidities, the shadow value of wealth would actually fall slightly, owing to the greater liquidity frictions, and the sharper decline in real interest rates would lead households to increase the minimum level of consumption by about 0.5%, as shown in the lower-center panel of Figure A3. This effect is essentially offset by the tightening of liquidity constraints, which reduce the mean-min consumption ratio, thus leading to a very small drop of consumption in the aggregate.

To conclude, consumption (and thus output and employment, which are equal to consumption up a first-order approximation) falls in our model mostly because of a sharp increase in the shadow value of wealth μ_t, rather than due to a tightening of liquidity constraints per se. In turn, the shadow value of wealth μ_t satisfies

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

an Δ_t is implicitly a function of the amount of credit available. To see this, recall that Δ is implicitly defined by

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

where $\frac{q_t b_{t+1}}{y_t}$ is the debt-to-consumption ratio in the economy.

The workings of our model are thus fairly similar to those of the simple New Keynesian model, in which $\mu_t = 1/c_t$. The key difference is that in our setup the natural real rate is pinned down endogenously, by the availability of credit as well as the strength of the precautionary-savings motive.

C.2 Economy with Fixed Prices

In Figure A4 we contrast the responses to an aggregate credit shock under the baseline parameterization reported in the paper with those in an economy with fixed prices ($\lambda_p = \lambda_w = 1$.) In this experiment we have chosen the size of the credit shock to ensure a maximal drop in debt of 0.25 of steady state income, roughly the drop in household debt observed in the U.S. during the 2007-2010 period. As earlier, the fact that the credit limit m_t follows a persistent autoregressive process as well as the fact that the credit limit applies only to new debt issues, implies that debt falls gradually and then recovers slowly as well. This is shown by the solid line in the upper-left panel of Figure A4 which scales the value of household debt $q_t b_{t+1}$ by the steady state income and reports deviations from the steady state. The solid line in the lower-left panel of Figure A4 shows the response of output (and, up to a first-order approximation, consumption and employment), expressed as % deviations from the steady state. These real variables fall by about 1.5% in response to the 25% decline in household credit, and gradually mean-revert.

The right panels of Figure A4 show the response of the nominal interest rate and inflation. The nominal interest rate falls gradually, by about 2% (annualized) at its trough and follows a path similar to that of the debt series. Inflation, in contrast, falls immediately, by about 1.5% (annualized), almost as much as output. Our parameter estimates thus imply a fairly steep slope of the Phillips curve. Moreover, they imply that interest rates fall by about 2%, too little compared to the drop in interest rates observed in the U.S. data, a point to which we return below.

Consider next the impulse responses in an economy with fixed prices and wages. Notice that the response of output is now much larger (about 3.5% at its trough which occurs about 12 quarters after the shock) and much more persistent: output returns to only about 2.8% below its steady state level even 30 quarters after the shock and only gradually recovers. Thus, with permanently fixed prices the model predicts what essentially amounts to *secular stagnation* – a long-term decline in both output and interest rates below their long-run trends.

Figure A5 provides some intuition for why the economy with greater price stickiness predicts a larger and more persistent response of output to household credit shocks. The figure reports the response of the real interest rate in our Benchmark model, the economy with fixed prices, as well as the natural rate of interest. Recall that if monetary policy were to mimic the latter, inflation would not react to a credit shock in our model, while output would fall by very little. Figure A5 shows that although the real interest rate falls by more initially in the economy with fixed prices, it recovers much more quickly than the real interest rate in our Benchmark model and is thus higher than the natural rate of interest starting from about four years (16 quarters) after the shock. This is a direct consequence of the particular parameterization of the Taylor rule that we have used, one that puts a relatively high weight on inflation and relatively little on output, thus implying a less responsive Taylor rule in the fixed price economy. Since output in the model depends on the entire expected path of real interest rates, the anticipation of persistently high real interest rates (relative to the natural rate) in the fixed price economy reduces output in a persistent fashion.

Figure A6 illustrates the model's responses to a persistent shock to the spread between the interest rate at which agents save/borrow and the federal funds rate. As discussed in the paper, we now assume that

$$i_t = f_t + \xi_t,$$

where the wedge itself follows an autoregressive process:

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

and modify the Taylor rule by assuming that the federal funds rate reacts to changes in such spreads and fully offsets them whenever possible:

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

The key point of this figure is that spread and credit shocks strongly reinforce each other. Without credit shocks, monetary policy offsets the increase in the spread by reducing its target rate f_t one-for-one (shown in the upper-right panel of the figure, for a spread shock of 2.5%), resulting in no change on output and inflation. In the presence of the household credit shock (chosen, as earlier, to generate a 25% drop in debt relative to steady state income), monetary policy can no longer fully respond as it is constrained by the zero lower bound for about 10 quarters. Consequently, output falls much more than in our previous experiment, by about 5%, instead of 1.5% in the absence of the increase in the spread.

D Robustness Checks

In this section we describe in some more detail some of the robustness checks we have conducted. We considered five additional experiments. First, we increased the degree of idiosyncratic uncertainty. Second, we studied an economy with more general CRRA preferences with a lower elasticity of intertemporal substitution. Third, we consider an alternative calibration of the model, using state-level wage data adjusted for worker composition from Beraja, Hurst and Ospina (2015). Fourth, we study a version of our model with a lower steady-state equilibrium interest rate. Finally, we introduce shocks to markups so that the model better matches the dynamics of inflation during the Great Recession.

D.1 Economy with Large Idiosyncratic Uncertainty

Figure A7 reports the results of our main experiment in which we feed the model a path for shocks that allow it to perfectly replicate the household debt-to-income data in the U.S., but in which we reduce α to 3.25. Recall that a lower α increases the uncertainty about idiosyncratic taste shocks and thus strengthens the precautionary savings motive.

As the upper-left panel of Figure A7 indicates, the same path for household debt is now associated with much larger swings in the natural interest rate, from 0 in 2001 to about 5% in 2008 and 0 again in 2013. All aggregate variables are consequently much more volatile. The nominal interest rate increases to 10% at the peak of the boom and reaches the zero lower bound by 2012. Consequently, employment is now much more volatile, almost as volatile as in the data. Inflation is extremely volatile as well, owing to the relatively steep slope of the Phillips curve implied by our estimates using state-level data.

Although this version of the model does a better job at reproducing the dynamics of employment in the U.S. data, it has a number of counterfactual implications. First, the model implies much wilder swings in both the nominal and real interest rate compared to that observed in the data. Part of this discrepancy is due to the large volatility of inflation, but the natural rate of interest itself fluctuates a lot as well. Second, the model now does a poor job at accounting for the state-level comovement between employment and household debt. Figure A8 illustrates this point, by contrasting the model's predictions for employment and comparing it to the data for a subset of states in the U.S. The model counterfactually predicts a much greater volatility of employment in response to state-level HH debt shocks than observed in the data.

D.2 Lower Elasticity of Intertemporal Substitution

We have assumed logarithmic preferences in our Benchmark model, thus restricting the elasticity of intertemporal substitution to unity. We next relax this assumption and study an economy with more general CRRA preferences. In particular, we assume that preferences of consumers on island s are now of the form

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

where σ is the inverse elasticity of intertemporal substitution. Notice that we also scale the idiosyncratic shocks v_{it} by σ, but this is an innocuous normalization given that we reestimate the volatility of the idiosyncratic shocks below using state-level data. The minimum consumption level (for the members with $v = 1$) is now given by

$$\underline{c}_t(s) = \left(\frac{1}{\beta \mathbb{E}_t \mu_{t+1}(s) R_{t+1}} \frac{1}{p_t(s)} \right)^{\frac{1}{\sigma}}, \tag{3}$$

The optimal amount of transfers now satisfies

$$\frac{\sigma}{\alpha - \sigma} \left(\frac{x_t(s)}{\underline{c}_t(s)} \right)^{-\alpha} = \left(\beta \mathbb{E}_t \frac{\mu_{t+1}(s)}{\mu_t(s)} R_{t+1} \right)^{-1} - 1, \tag{4}$$

while aggregate consumption is, as earlier,

$$\frac{c_t(s)}{\underline{c}_t(s)} = \frac{\alpha}{\alpha - 1} \left(1 - \frac{1}{\alpha} \left(\frac{x_t(s)}{\underline{c}_t(s)} \right)^{1-\alpha} \right). \tag{5}$$

We set $\sigma = 3$, corresponding to an elasticity of intertemporal substitution of 1/3 and have re-estimated all the parameters of the model using the state-level auxiliary panel regressions. The column labeled 'Low EIS' in Table A2 reports how the model fits the data and the new coefficients that achieve this fit. As Panel A of this table shows, the overall model R^2, which summarizes the fraction of variation in the data accounted for by the model is 0.94 now, almost as high as in the Benchmark model. Once again, the model accounts well for the variation in state-level consumption, employment, wages and house prices quite. Panel B shows that the model now requires a smaller dispersion in idiosyncratic taste shocks ($\alpha = 6.5$ compared to 5.5 in the Benchmark economy with logarithmic preferences), and a relatively smaller dispersion in the relative volatility of housing taste shocks. Two additional changes are worth noting. First, we now estimate a much greater elasticity of substitution between labor varieties ($\psi = 13.5$ compared to 5.4 in the Benchmark model). Intuitively, a lower EIS also implies much stronger wealth effects on labor supply and thus volatility in wages, which must be offset by increasing the degree of real rigidity in the model. Second, notice that the annualized discount rate needed to match a 2% real interest rate is now higher and equal to 5.3%. The gap between the discount rate and the interest rate is thus more than twice higher than in the Benchmark model (3.3% vs. 1.5%). Thus, the economy with lower EIS is characterized by a stronger precautionary savings motive, even though it matches the state-level data as well as the economy with logarithmic preferences.

Figure A9 reports the model's implications for the aggregate time-series when we confront it with a sequence of shocks that allow it to reproduce the household debt to income series in the U.S. For simplicity, we no longer report the debt-to-income series since this is identical to that in Figure A7.

Notice in the upper-left panel of Figure A9 that the natural interest rate is now about twice more volatile than in the Benchmark economy: it increases from about 1% to 4% during the boom and falls to about 0.5% in 2013. The larger movements in the natural interest rate are a consequence of the stronger precautionary savings motive implied by the new estimates of the model.

The upper-right panel of Figure A9 shows that the nominal interest rate does not reach the ZLB in this economy: the nominal rate falls to only about 2% by 2013. Intuitively, even though the interest rate needs to fall by more now during the recession, the level of the nominal interest rates at their peak is now higher (almost 7% in 2007), owing to the large rise in the natural interest rate during the boom, thus giving monetary policy more room to maneuver during the recession. As earlier, the failure of household credit shocks to trigger the ZLB on their own implies that such shocks are offset fairly well, so fluctuations in employment are fairly modest compared to those observed in the data.

We have also experimented with raising σ even further. As we do so, natural interest rates indeed become more and more volatile, and consequently so does employment. However, the model implies counterfactually large increases in nominal interest rates during the boom. Household debt shocks alone can only trigger the zero lower bound if the volatility of the natural rate of interest is high enough for nominal interest rates to increase above 10% by 2007, which is clearly counterfactual.

D.3 Beraja, Hurst and Ospina (2015) Wage Data

Recall that in our Benchmark experiments we have used a measure of wage data from the BEA, by dividing labor compensation by state employment. In a recent paper Beraja, Hurst and Ospina (2015) construct an alternative wage panel using the Current Population Survey,[1] which accounts for variation in the composition of the labor force (age, race, educational attainment, race). This new data exhibits much more time-series variation in the cross-section of states compared to the BEA data.

The column labeled 'BHO Wages' in Table A2 reports results from our re-estimation of the model using this new wage panel. Notice in parentheses that in the data the elasticity of wages to current HH debt is equal to 0.18 (compared to 0.09 for the BEA data), and the elasticity of wages to lagged HH debt is equal to -0.09 (compared to -0.04 in the BEA data). The new wage data is thus twice more volatile than the BEA series (or the unadjusted CPS data, which shows a very similar pattern as that of the BEA data). The re-estimated model matches this pattern well, by reducing the extent of real wage rigidity (ψ now falls to 2.33 compared to 5.38 in the Benchmark model).

Figure A10 presents the aggregate implications of this version of the model. Notice that all macroeconomic time-series, including inflation, are virtually identical to those in our Benchmark calibration. Although the Phillips curve is somewhat steeper now, movements in employment are now less volatile and consequently inflation falls by nearly as much as in our Benchmark model.

A starker way to see the role of price rigidities in shaping the model's responses to

[1]We are grateful to Erik Hurst for sharing the data with us.

household credit shocks is to assume that prices and wages are permanently fixed. Figure A11 shows that employment now fluctuates considerably more, from about 2% above the steady-state at the peak to -1% at the bottom, thus a 3% drop or about 40% of that observed in the data. Of course, the model's implications for prices and wages would now be grossly at odds with the state-level evidence documented by Beraja, Hurst and Ospina (2015).

D.4 Lower Steady State Real Interest Rate

We have also re-estimated the model by targeting a lower steady state real interest rate of 0.5% (compared to 2% in our Benchmark model). The last column of Table A2 reports the moments and parameters for this version of the model. Note that we now need much more volatile house preference shocks (σ_η is 18 compared to 7.5 earlier) and a lower discount rate ($-\log(\beta)$ is now equal to 1.98 compared to 3.49 in the Benchmark model). The model's fit is just as good as that of our Benchmark model.

Figure A12 shows that the model's implications for real variables are virtually unchanged. Although both the natural and nominal interest rates shift down by 1.5% now, the nominal interest rate continues to be above zero and so the model's real implications are unchanged.

D.5 Matching Inflation Dynamics in the Great Recession

Recall that when we feed our Benchmark model both household credit shocks as well as shocks to spreads, we find that employment falls considerably, but so does inflation. Because of the relatively steep slope of the Phillips curve implied by our state-level estimates, inflation falls to about -0.5%, much lower than the 2% observed in the data. Here we argue that our model's implications for employment are not critically affected when we study a variation of the model that matches the inflation data.

To match the inflation dynamics in the data, we choose a path for shocks to the elasticity of substitution between varieties of goods, so that the firm's desired markup changes. These shocks add an additional term to the inflation Phillips curve of our model. When confronting the model with this additional series of shocks, we re-calibrate the persistence of the shock to spreads to ensure the same expected duration at the ZLB of 5 quarters as in our Benchmark model.

Figure A13 shows that although the model matches now the inflation data post 2007 much better, its implications for employment are virtually unchanged. This reflects the fact that with a lower drop in inflation we need a greater persistence of the spread shocks to match the expected duration of the ZLB in the data. Conditional on matching this duration, the marginal effect of household credit shocks on employment is therefore unchanged.

E Solution Method

We briefly discuss the solution method we used to study the economy in the presence of the ZLB. We then show that given our piecewise-linear approximation method, the marginal effect of a HH credit shock is uniquely pinned down by agents' expectations of the duration of the ZLB.

The algorithm we used was developed by Jones (2015), itself a variation of the OccBin toolbox of Guerrieri and Iacoviello (2015) and Kulish and Pagan (2016).[2] We follow the New Keynesian literature by log-linearizing all the equilibrium conditions in the absence of a zero lower bound. We thus have

$$\mathbf{A}\mathbf{x}_t = \mathbf{C} + \mathbf{B}\mathbf{x}_{t-1} + \mathbf{D}\mathbb{E}_t\mathbf{x}_{t+1} + \mathbf{F}\epsilon_t, \tag{6}$$

where $\mathbf{x}_t$ is a vector collecting the endogenous variables and ϵ_t are the aggregate shocks.[3] One of the equations in this system is the interest rate rule that, for now, ignores the zero lower bound. We use standard methods to characterize the solution of this system, which gives

$$\mathbf{x}_t = \mathbf{J} + \mathbf{Q}\mathbf{x}_{t-1} + \mathbf{G}\epsilon_t. \tag{7}$$

In periods in which the ZLB binds, we replace the interest rate rule with $i_t = 0$, thus modifying the system of equations to

$$\mathbf{A}^*\mathbf{x}_t = \mathbf{C}^* + \mathbf{B}^*\mathbf{x}_{t-1} + \mathbf{D}^*\mathbb{E}_t\mathbf{x}_{t+1} + \mathbf{F}^*\epsilon_t. \tag{8}$$

The algorithm we use is based on a piece-wise linear solution of the equilibrium conditions in these two regimes, under the assumption that agents observe the credit shock ϵ_t in each period, but believe that no other shocks are possible in the future. The algorithm thus computes the perfect-foresight transition dynamics for any given point $(\mathbf{x}_{t-1}, \epsilon_t)$.

To make things concrete, suppose that $(\mathbf{x}_{t-1}, \epsilon_t)$ is such that agents conjecture that the interest rate i_t will be equal to zero from some date $t' \geq t$ until T. Note that we allow for the

[2]We are extremely grateful to Callum Jones for sharing his codes and helping us with the computations described in this section.

[3]Because we have linearized the model, island-level shocks wash out, by assumption, in the aggregate.

possibility that a given shock today causes agents to forecast a binding ZLB in the *future,* owing for example to the hump-shaped nature of the response of HH debt to shocks.

Since the ZLB stops binding at T, the law of motion for aggregate variables in that period is given by $\mathbf{x}_T = \mathbf{J} + \mathbf{Q}\mathbf{x}_{T-1}$. In period $T-1$, the ZLB binds, so we use (8) to calculate date $T-1$ variables. We thus have

$$\mathbf{x}_{T-1} = \left[\mathbf{A}^* - \mathbf{D}^*\mathbf{Q}\right]^{-1}\left[\mathbf{C}^* + \mathbf{D}^*\mathbf{J} + \mathbf{B}^*\mathbf{x}_{T-2}\right] = \mathbf{J}_{T-1} + \mathbf{Q}_{T-1}\mathbf{x}_{T-2}.$$

We continue iterating backwards, until period t, using (8) in those periods in which the ZLB is conjectured to bind, and (6) in all other periods. This backward recursion gives a path for all variables of the model, $\mathbf{x}_{t+j}$. If this path is such that the original conjecture agents have made regarding the timing of when the ZLB is in effect is correct, we stop the algorithm. Otherwise, we update the guess for when the ZLB binds and continue iterating accordingly. Once the algorithm has converged, the solution of the model is

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

The critical thing to notice is that we compute this solution by iterating backward starting from the date T at which the ZLB is expected to stop binding. Thus, the marginal impulse response of a given shock ϵ_t, encoded in $\mathbf{Q}_t$ and $\mathbf{G}_t$, is only a function of the expected ZLB duration.

References

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

[2] Jones, Callum. 2015. “Unanticipated Shocks and Forward Guidance At the Zero Lower Bound.” NYU Working Paper.

[3] Kulish, M. and A. Pagan, 2016. “Estimation and Solution of Models with Expectations and Structural Changes”, forthcoming, *Journal of Applied Econometrics.*

A1: Identification of Key Parameters

	Data	Benchmark	$\alpha = 4$	$\gamma = 0.975$	$\psi = 3$	$\omega = 0.75$
log employment on current debt	0.18	0.18	0.32	0.22	0.16	0.12
log employment on lagged debt	-0.15	-0.15	-0.28	-0.20	-0.14	-0.10
log consumption on current debt	0.30	0.30	0.55	0.37	0.30	0.25
log consumption on lagged debt	-0.20	-0.21	-0.42	-0.27	-0.21	-0.19
log wages on current debt	0.09	0.11	0.20	0.14	0.14	0.08
log wages on lagged debt	-0.04	-0.05	-0.13	-0.06	-0.08	-0.05
log house prices on current debt	1.94	1.84	2.97	3.17	1.84	1.81
log house prices on lagged debt	-1.40	-1.51	-2.46	-2.83	-1.51	-1.50
relative std dev. employment (model/data)	-	1.00	1.80	1.27	0.89	0.68
relative std dev. consumption (model/data)	-	0.91	1.56	1.12	0.91	0.70
relative std dev. wages (model/data)	-	1.05	1.59	1.49	1.26	0.63
relative std dev. house prices (model/data)	-	0.88	1.43	1.59	0.88	0.87
correlation employment (model vs. data)	-	0.98	0.99	1.00	1.00	0.99
correlation consumption (model vs. data)	-	0.99	0.93	0.98	0.98	0.93
correlation wages (model vs. data)	-	0.99	0.94	1.00	0.98	0.94
correlation house prices (model vs. data)	-	0.93	0.91	0.79	0.93	0.91

A2: Robustness Checks: Parameterization

	Data	Benchmark	Low EIS	BHO Wages	Low r
			A. Coefficients in Auxiliary Regressions		
log employment on current debt	0.18	0.18	0.17	0.18	0.17
log employment on lagged debt	-0.15	-0.15	-0.14	-0.15	-0.14
log consumption on current debt	0.30	0.30	0.29	0.33	0.28
log consumption on lagged debt	-0.20	-0.21	-0.21	-0.23	-0.20
log wages on current debt	0.09	0.11	0.10	0.18 (0.17)	0.10
log wages on lagged debt	-0.04	-0.05	-0.05	-0.10 (-0.09)	-0.05
log house prices on current debt	1.94	1.84	1.89	1.87	1.88
log house prices on lagged debt	-1.40	-1.51	-1.50	-1.40	-1.48
R^2_{Model}		0.94	0.96	0.99	0.96
			B. Parameters		
Tail Pareto taste shocks	α	5.50	6.50	5.50	5.51
Persistence coupon payments	γ	0.95	0.95	0.94	0.94
AR(1) shocks	ρ	0.76	0.81	0.85	0.79
Volatility housing shocks	σ_η	7.49	3.54	8.50	18.00
Elasticity labor varieties	ψ	5.38	13.48	2.33	4.82
Weight on non-traded goods	ω	0.87	0.86	0.94	0.88
Annual discount rate, %	$-\ln(\beta)$	3.49	5.32	3.49	1.98

Figure A1: Household Debt in the U.S.

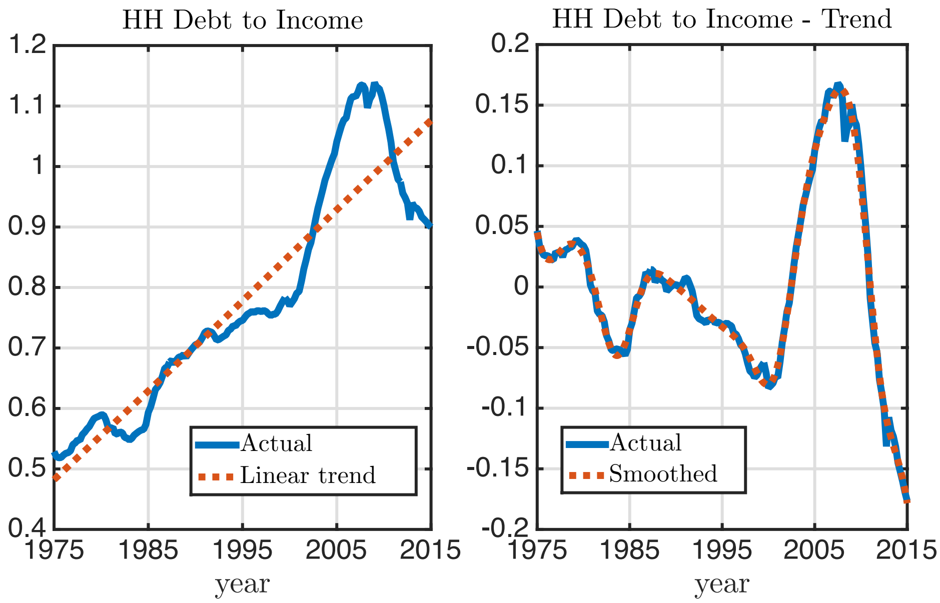

Figure A2: Expected Duration of ZLB

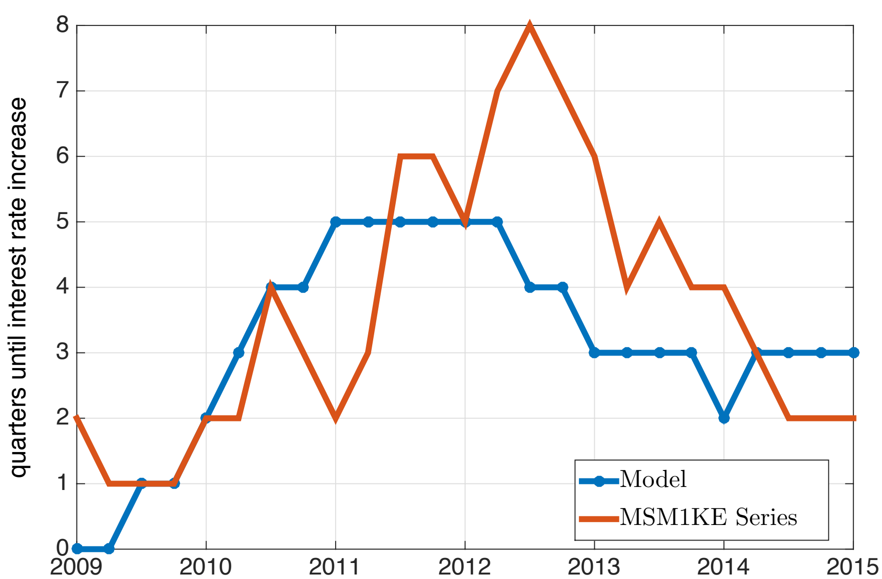

Figure A3: Mechanics of the Model

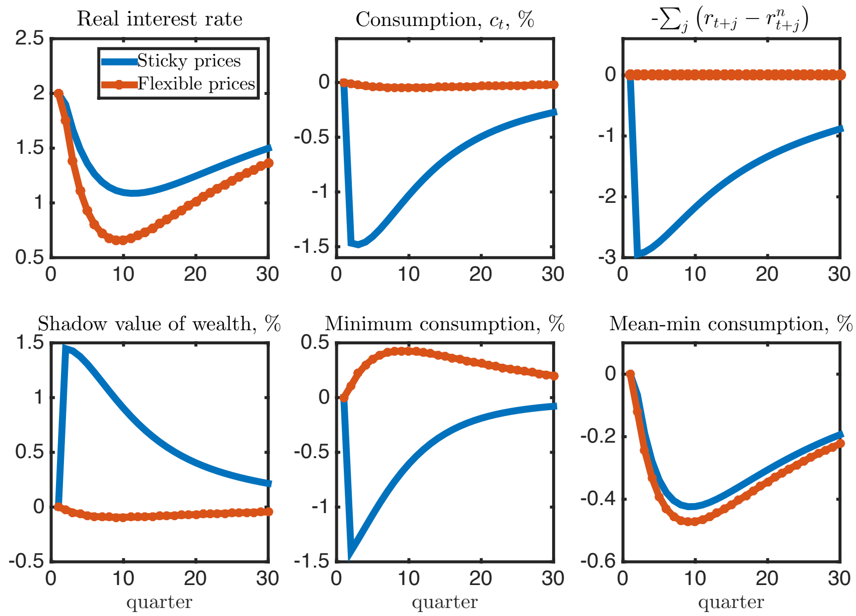

Figure A4: Comparison with Economy with Fixed Prices

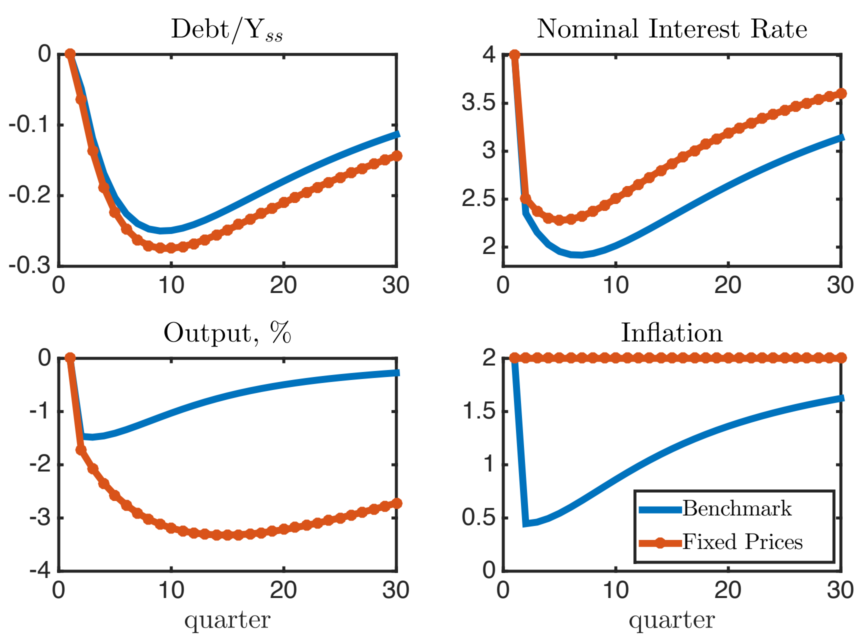

Figure A5: Interest Rates after an Aggregate Credit Shock

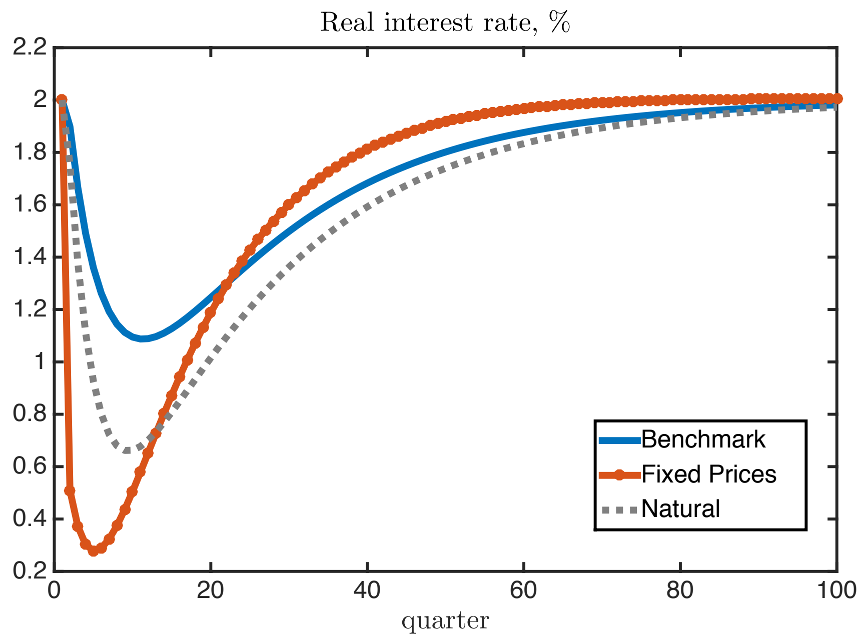

Figure A6: Impulse Response to a Shock to Spreads

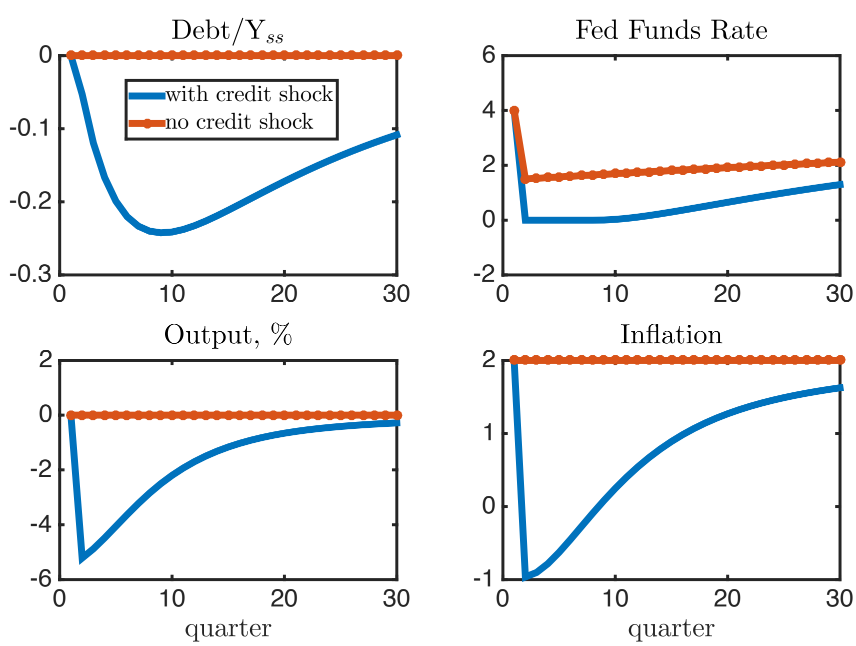

Figure A7: Economy with Large Idiosyncratic Uncertainty

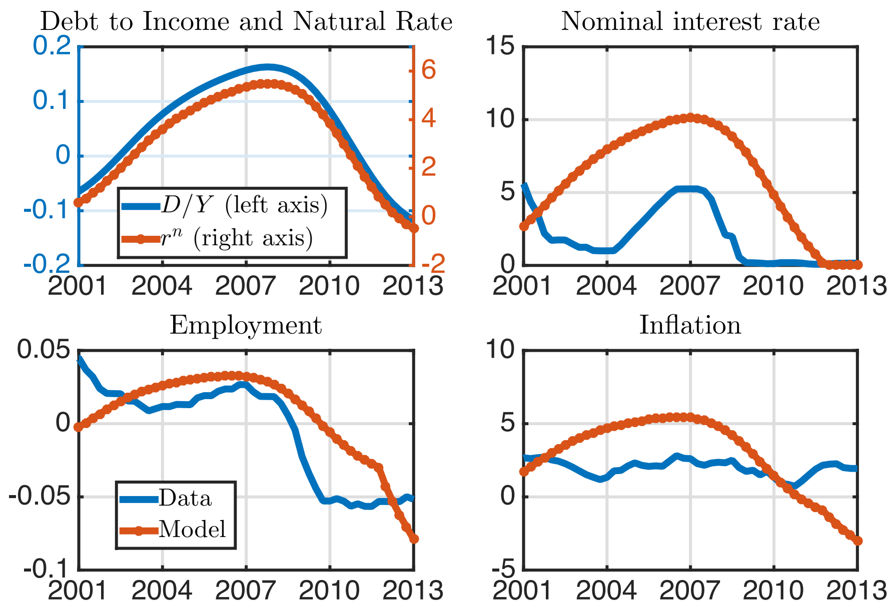

Figure A8: State Employment. Large Idiosyncratic Uncertainty

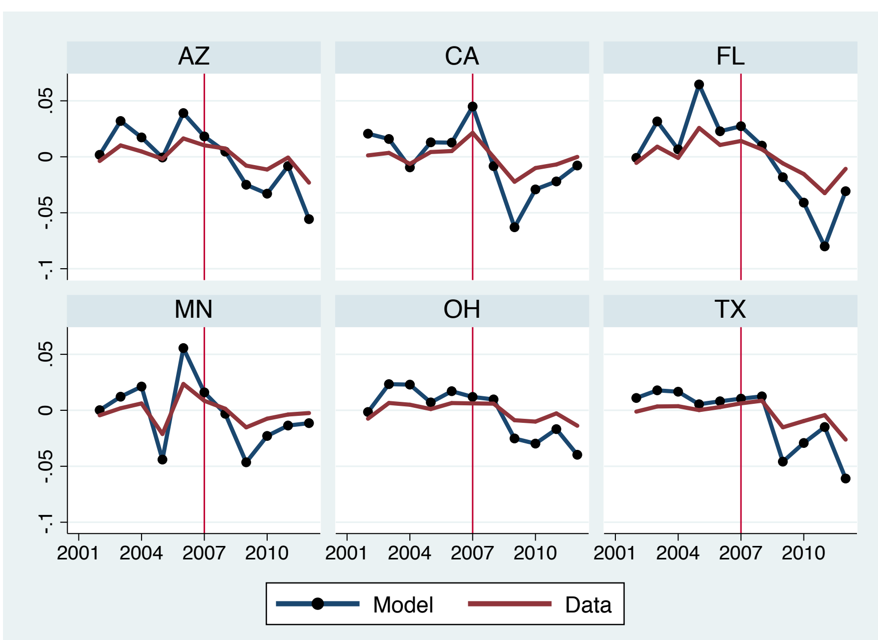

Figure A9: Aggregate Implications. Low EIS Economy

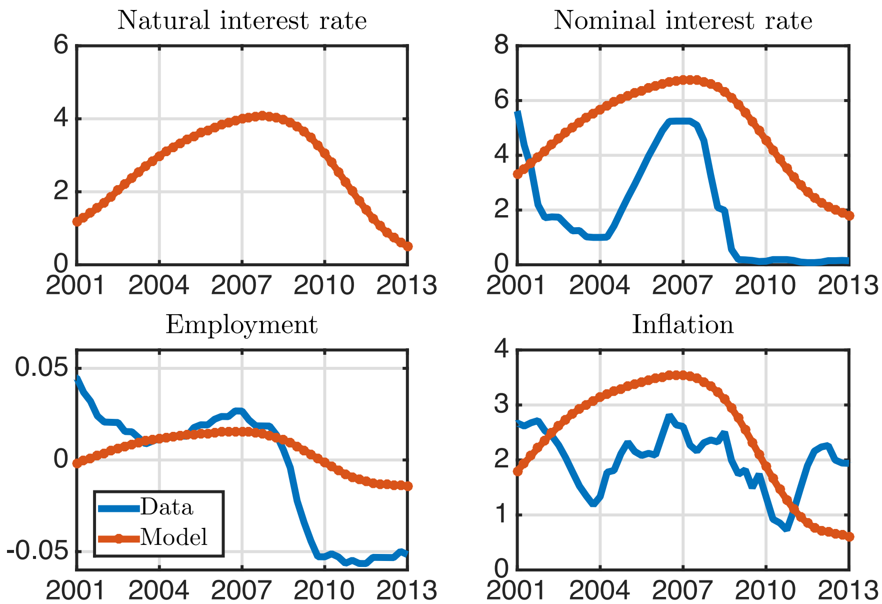

Figure A10: Aggregate Implications. Match BOH Wage Data

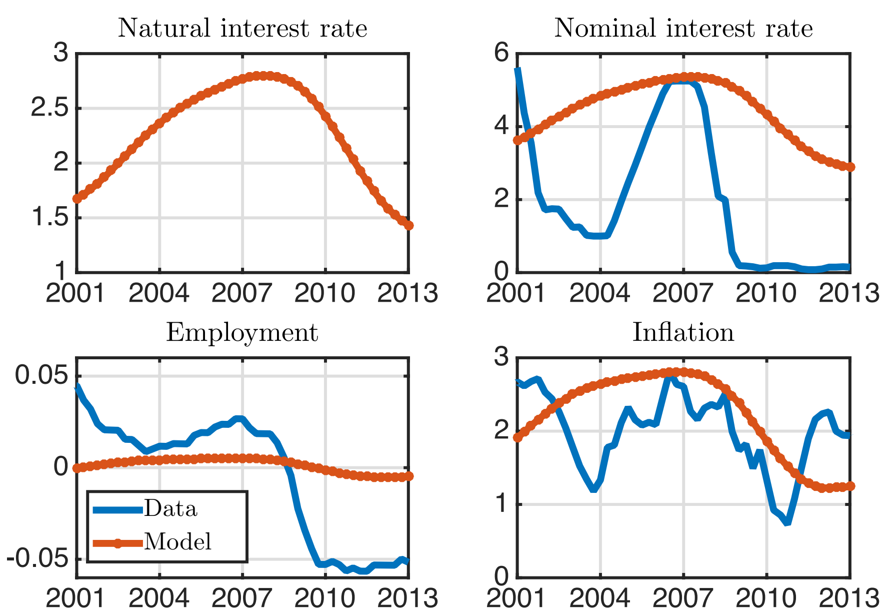

Figure A11: Aggregate Implications. Economy with Fixed Prices

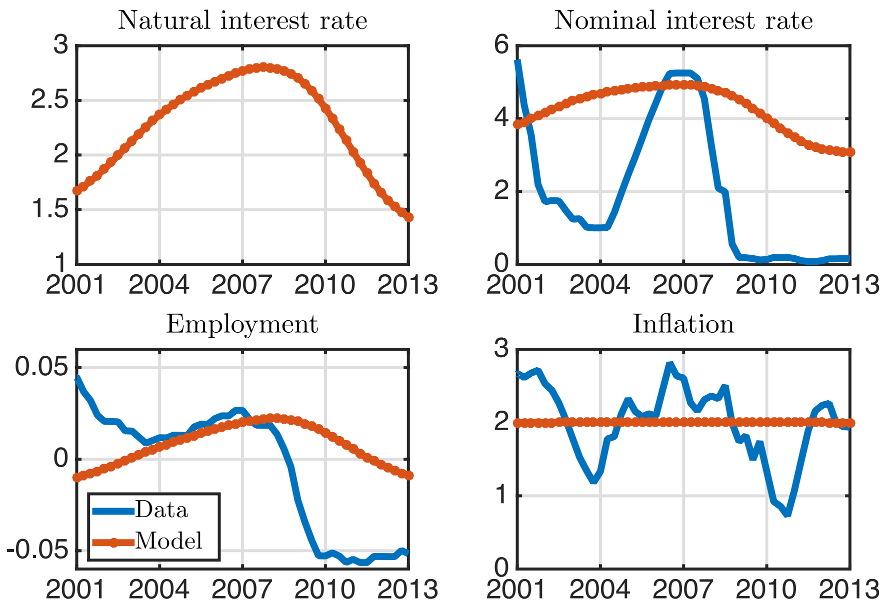

Figure A12: Aggregate Implications. Low r Economy

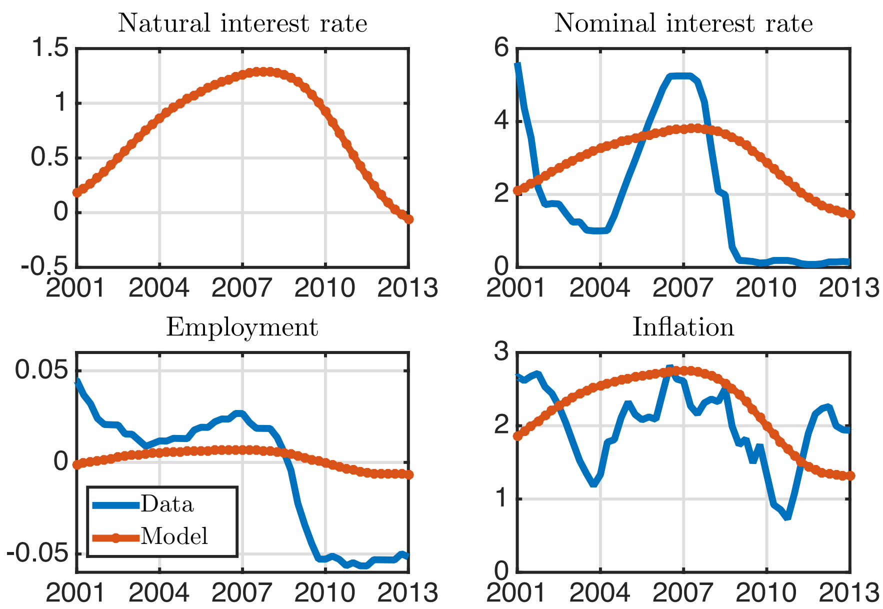

Figure A13: Aggregate Implications. Matching Inflation Dynamics