A New Way to Quantify the Effect of Uncertainty*

Alexander W. Richter Nathaniel A. Throckmorton

First Draft: May 4, 2017

This Draft: May 14, 2018

ABSTRACT

This paper develops a new way to quantify the effects of aggregate uncertainty that accounts for exogenous and endogenous sources. First, we use Bayesian methods to estimate a nonlinear New Keynesian model with stochastic volatility and a zero lower bound constraint on the nominal interest rate. We discipline the model by matching data on uncertainty, in addition to common macro time series. Second, we use the Euler equation to decompose output into expected output and the expected variance and skewness of output. We then filter a time series for each term. Our method captures the effects of higher-order moments over horizons beyond 1 quarter by recursively decomposing expected output. Over a 1-quarter horizon, output uncertainty reduced output less than 0.01% every quarter, similar to volatility shocks in our model. Over horizons that remove the influence of expected output, output uncertainty on average reduced output 0.06% and the peak effect was 0.15% during the Great Recession, similar to structural VAR estimates. Other higher-order moments had much smaller effects on output.

Keywords: Endogenous Uncertainty; Stochastic Volatility; Nonlinear Estimation; Lower Bound
JEL Classifications: C11; D81; E32; E58

*Richter, Research Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX 75201 (alex.richter@dal.frb.org); Throckmorton, Department of Economics, William & Mary, P.O. Box 8795, Williamsburg, VA 23187 (nat@wm.edu). We especially want to thank Cosmin Ilut for discussing our paper at the NBER Workshop for DSGE Models and Mike Plante and Todd Walker for useful suggestions that improved the paper. We also thank Esteban Argudo, Pablo Cuba-Borda, Oliver de Groot, Ed Herbst, Ben Johannsen, Campbell Leith, Karel Mertens, Jim Nason, Frank Schorfheide, Andrea Tambalotti, Alex Wolman, Kei-Mu Yi and seminar participants at several conferences and institutions for helpful comments. We thank Eric Walter and Chris Stackpole for supporting the cluster at our institutions and Auburn University for providing access to its cluster. The views in the paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 INTRODUCTION

There is widespread agreement that uncertainty decreases economic activity. The debate rests on whether the effect is quantitatively significant, which is difficult to determine for two reasons. One, uncertainty is unobserved, so there is disagreement on what constitutes a good measure. Until recently, the literature has relied on proxies for uncertainty, such as realized or implied volatility, indexes based on keywords in print or online media, and survey-based forecast dispersion, which are often weakly correlated with each other and loosely connected with the definition of uncertainty.

Two, uncertainty is endogenous. Not only can uncertainty affect economic activity, as intuition suggests, what is happening in the economy can affect uncertainty. A few mechanisms emphasized in the literature include financial frictions and constraints that create an adverse feedback loop between net worth and asset prices [Brunnermeier and Sannikov (2014)], incomplete information that endogenously creates pessimism during recessions [Fajgelbaum et al. (2017); Saijo (2017); Van Nieuwerburgh and Veldkamp (2006)], and a zero lower bound (ZLB) constraint on the nominal interest rate that restricts a central bank's ability to stabilize the economy [Plante et al. (2018)].

The literature often uses exogenous volatility shocks to examine the effects of aggregate uncertainty. This paper develops a new approach that accounts for both exogenous and endogenous uncertainty sources. First, we use Bayesian methods to estimate a nonlinear New Keynesian model with stochastic volatility and an occasionally binding ZLB constraint. We discipline the model by matching data on uncertainty, in addition to common macro time series. This step allows us to decompose the sources of uncertainty and generate a data-driven policy function for any moment. Second, we use the Euler equation to decompose output into expected output and the expected variance and skewness of output. We then filter a time series for each term in the decomposition.

A major benefit of our method is that it captures the effects of higher-order moments over horizons beyond 1 quarter by recursively decomposing expected output. Over a 1-quarter horizon, output uncertainty reduced output less than 0.01% every quarter, similar to volatility shocks in our model. Over horizons that remove the influence of expected output, output uncertainty on average reduced output 0.06% and the peak effect was 0.15% during the Great Recession, similar to structural VAR estimates. Roughly one-third of the increase during the Great Recession was due to the ZLB. When we extend our model without capital so households can invest, the average effect of output uncertainty increases to -0.08% and the peak effect rises to -0.22%, but the differences are not statistically significant. Other higher-order moments had much smaller effects on output.

We conduct two exercises to uncover the drivers of our results. One, we use counterfactual simulations to decompose uncertainty into its endogenous and exogenous sources. Endogenous uncertainty—uncertainty that naturally arises due to first moment shocks—typically accounted for 95% of total uncertainty. However, nearly all of the changes in uncertainty were driven by the

volatility shocks. One exception is when the Fed was constrained. In 2009Q1, 8.5% of the increase in uncertainty was due to endogenous uncertainty while 38% was due to the endogenous amplification of volatility shocks. Two, we determine the importance of each parameter in our model for the results of our Euler equation decomposition using posterior predictive analysis. While price adjustment costs play an important role as others have emphasized, risk aversion and the monetary response to inflation had the largest impact on the effects of uncertainty among the deep parameters.

We conclude our analysis by calculating the welfare effects of uncertainty following the cost of business cycles literature. We find the welfare cost of volatility never exceeded 0.04% of consumption. We also compare impulse responses to a financial uncertainty shock in our nonlinear model to the same shock in a structural VAR using a recursive identification scheme, since that is the most common way to identify the effects of uncertainty. Using data simulated from the nonlinear model, the VAR generates quantitatively similar responses to our structural model. There is almost no response of output in data without the ZLB and a larger effect in data with a lengthy ZLB event. Similar differences in the responses occur in U.S. data with and without the ZLB period in the sample.

Although we use a familiar model as a starting point for understanding the effects of higher-order moments, our method is adaptable to a broad class of models. For example, it can be applied to models with limited information, irreversible investment, borrowing constraints, search frictions, heterogeneous agents, or other important sources of time-varying endogenous uncertainty. While those features may make the model too costly to estimate, approximate solutions are attainable either locally with perturbation methods or globally with projection methods. With a solution in hand, it is possible to calculate the expected variance or skewness surrounding any endogenous variable and link it to an empirical measure while filtering the data. Given a particular calibration, the filter can then generate time series for the terms in any Euler equation. Therefore, our method provides a way to compare the effects of uncertainty or other high-order moments across models.

The paper proceeds as follows. Section 2 places our work within the vast literature on uncertainty. Section 3 describes our model as well as the exogenous and endogenous sources of uncertainty. Section 4 outlines our solution and estimation procedures. Section 5 provides our estimation results, including the parameter estimates and the effects of uncertainty and skewness on output. Section 6 shows how our results change when we introduce capital. Section 7 draws comparisons between the impulse responses in our nonlinear model and a structural VAR. Section 8 concludes.

2 RELATED LITERATURE

Research that examines the effects of aggregate uncertainty has considered several different shocks. In a small open-economy real business cycle model, Fernández-Villaverde et al. (2011) examine volatility shocks to a country-specific interest rate spread. They find a 1 standard deviation shock lowers output 0.15%-0.2% in Argentina and Ecuador and 0.01%-0.02% in Brazil and Venezuela.

Most papers develop closed-economy New Keynesian models. Mumtaz and Zanetti (2013) focus on monetary policy volatility shocks in a model without capital. They find doubling the volatility reduces output growth by only 0.03%. Born and Pfeifer (2014) introduce variable capital utilization and investment adjustment costs. They show a simultaneous 2 standard deviation increase in uncertainty about government spending, monetary policy, and capital and labor taxes reduces output by only 0.065%. In contrast, Fernández-Villaverde et al. (2015) find a volatility shock to only capital taxes reduces output by 0.1% and the effects are much larger when the ZLB binds.

In a textbook model with recursive preferences, Basu and Bundick (2017) find a 1 standard deviation preference volatility shock—a proxy for demand uncertainty—reduces output by 0.2%. de Groot et al. (2018) show the way the shock enters their preferences creates an asymptote in the parameter space that amplifies the output response. Without the asymptote, preference volatility shocks have a small effect. Leduc and Liu (2016) include search frictions and habit formation and find a 1 standard deviation increase in technology volatility—a proxy for supply uncertainty—increases unemployment by 2.6%. In our paper, both supply and demand uncertainty varies exogenously due to stochastic volatility shocks to the risk premium and the growth rate of technology.

The volatility shocks in our model also have a small impact, but the effect of output uncertainty from our Euler equation decomposition is an order of magnitude larger. That result emphasizes the importance of accounting for the expected effects of uncertainty over horizons beyond 1 quarter. Another major benefit of our approach is that it directly links the measures of uncertainty in our model—second moments—to equivalent measures in the data with likelihood based methods, whereas previous work relied on first moments such as real activity, interest rates, and fiscal variables. We also quantify the effects of other higher-order moments, such as the skewness of output and the covariance between output and inflation, which have received less attention in the literature.

As an alternative to exogenous volatility shocks, several papers propose models that endogenously generate uncertainty. There are several mechanisms. One segment emphasizes the role of a financial sector under complete information, where the severity and duration of financial crises are stochastic. Most papers focus on crises that result from financial frictions and collateral constraints [Brunnermeier and Sannikov (2014); He and Krishnamurthy (2014); Mendoza (2010)], while a few papers incorporate the role of firm default [Arellano et al. (2016); Gourio (2014); Navarro (2014)].

Another segment examines the implications of incomplete information. Some of the papers feature learning with aggregate shocks [Fajgelbaum et al. (2017); Saijo (2017); Van Nieuwerburgh and Veldkamp (2006)], while others focus on firm-specific shocks [Ilut and Saijo (2016); Straub and Ulbricht (2015)]. In these models, an adverse shock under asymmetric learning lowers economic activity and makes it harder for households to learn about the economy, which amplifies the effects of first moment shocks. In our model, the effects of first and second moment shocks are amplified by the ZLB [Basu and Bundick (2017); Fernández-Villaverde et al. (2015); Nakata (2017); Plante

et al. (2018)]. We bridge the gap between the exogenous and endogenous uncertainty literatures by providing a flexible methodology that is easily applied to models with both types of uncertainty.

Our paper is also related to the cost of business cycles literature. Lucas (1987) examines the welfare cost of "instability" by calculating the fraction of consumption goods a household would give up each period to eliminate volatility. With constant relative risk aversion preferences, Lucas finds the welfare cost of the consumption volatility in post-World War II data ranges from 0.008% (log utility) to 0.17% (risk aversion, $\gamma = 20$). The conclusion is that the cost of instability is insignificant.[1] We build on this literature by calculating welfare at each point in our sample using an estimated model that matches both macro and uncertainty data. We find the welfare costs of first moment shocks are well within the range Lucas reported. Second moments shocks have an even smaller welfare effect, consistent with the values reported in Xu (2017). We view this important exercise as complementary to our Euler equation decomposition. However, a benefit of our decomposition is that it shows which higher-order moments are most important at each point in time.

3 NEW KEYNESIAN MODEL AND UNCERTAINTY MEASURES

We use a New Keynesian model similar to An and Schorfheide (2007), except it includes a ZLB constraint and stochastic volatility on technology growth and the risk premium on a nominal bond.

3.1 FIRMS

The production sector consists of a continuum of monopolistically competitive intermediate goods firms and a final goods firm. Intermediate firm $i \in [0, 1]$ produces a differentiated good, $y_t^f(i)$, according to $y_t^f(i) = z_t n_t(i)$, where $n(i)$ is the labor hired by firm i and $z_t = g_t z_{t-1}$ is technology, which is common across firms. Deviations from the balanced growth rate, $\bar{g}$, follow

$$g_t = (1 - \rho_g)\bar{g} + \rho_g g_{t-1} + \sigma_{g,t}\varepsilon_{g,t},\ 0 \leq \rho_g < 1,\ \varepsilon_g \sim \mathbb{N}(0, 1), \tag{1}$$

$$\sigma_{g,t} = \bar{\sigma}_g(\sigma_{g,t-1}/\bar{\sigma}_g)^{\rho_{\sigma_g}} \exp(\sigma_{\sigma_g}\varepsilon_{\sigma_g,t}),\ 0 \leq \rho_{\sigma_g} < 1,\ \varepsilon_{\sigma_g} \sim \mathbb{N}(0, 1), \tag{2}$$

where the standard deviation of the technology shock, σ_g, follows an independent log-normal process (σ_g and ε_g are uncorrelated) to add a source of time-varying supply uncertainty to the model.

The final goods firm purchases $y_t^f(i)$ units from each intermediate firm to produce the final good, $y_t^f \equiv [\int_0^1 y_t^f(i)^{(\theta-1)/\theta} di]^{\theta/(\theta-1)}$, where $\theta > 1$ is the elasticity of substitution. It then maximizes dividends to determine its demand function for intermediate good i, $y_t^f(i) = (p_t(i)/p_t)^{-\theta} y_t^f$, where $p_t = [\int_0^1 p_t(i)^{1-\theta} di]^{1/(1-\theta)}$ is the price level. Following Rotemberg (1982), each intermediate firm pays a price adjustment cost, $adj_t^f(i) \equiv \varphi_f(p_t(i)/(\bar{\pi} p_{t-1}(i)) - 1)^2 y_t^f/2$, where $\varphi_f > 0$ scales the cost and $\bar{\pi}$ is the gross inflation rate along the balanced growth path. Therefore, firm i chooses $n_t(i)$ and $p_t(i)$ to maximize the expected discounted present value of future dividends,

[1] Several papers examine these estimates in different settings [Lester et al. (2014); Otrok (2001); Tallarini (2000)].

$E_t \sum_{k=t}^{\infty} q_{t,k} d_k(i)$, subject to its production function and the demand for its product, where $q_{t,t} \equiv 1$, $q_{t,t+1} \equiv \beta(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma}$ is the pricing kernel between periods t and $t+1$, $q_{t,k} \equiv \prod_{j=t+1}^{k>t} q_{j-1,j}$, $d_t(i) = p_t(i) y_t^f(i)/p_t - w_t n_t(i) - adj_t^f(i)$, and a tilde denotes a detrended variable (i.e., $\tilde{x} = x/z$). In symmetric equilibrium, all firms make identical decisions, so the optimality conditions imply

$$\tilde{y}_t^f = n_t, \tag{3}$$

$$mc_t = \tilde{w}_t, \tag{4}$$

$$\varphi_f(\pi_t^{gap} - 1)\pi_t^{gap} = 1 - \theta + \theta mc_t + \beta \varphi_f E_t[(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma}(\pi_{t+1}^{gap} - 1)\pi_{t+1}^{gap}(\tilde{y}_{t+1}^f/\tilde{y}_t^f)], \tag{5}$$

where $\pi_t^{gap} \equiv \pi_t/\bar{\pi}$ is the inflation gap. In the special case where prices are perfectly flexible (i.e., $\varphi_f = 0$), $\tilde{w}_t = (\theta - 1)/\theta$, which equals the inverse of the gross markup of price over marginal cost.

3.2 HOUSEHOLDS The representative household chooses $\{c_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \beta^t [((c_t/z_t)^{1-\gamma} - 1)/(1-\gamma) - \chi n_t^{1+\eta}/(1+\eta)]$, where γ is the coefficient of relative risk aversion, $\chi > 0$ is a preference parameter that determines the steady state labor supply, $1/\eta$ is the Frisch elasticity of labor supply, c is consumption, n is labor hours, b is the real value of a privately-issued 1-period nominal bond that is in zero net supply, and E_0 is the mathematical expectation operator conditional on information in period 0. Following An and Schorfheide (2007), households receive utility from consumption relative to the level of technology, which is a proxy for the habit stock. That assumption allows us to use additively separable preferences and parameterize the degree of risk aversion while maintaining a balanced growth path. The household's choices are constrained by $c_t + b_t/(i_t s_t) = w_t n_t + b_{t-1}/\pi_t + d_t$, where π is the gross inflation rate, w is the real wage rate, i is the gross nominal interest rate set by the central bank, and d is a real dividend received from owning the intermediate goods firms. Following Smets and Wouters (2007) and Gust et al. (2017), s is a shock to the risk premium on the nominal bond and it evolves according to

$$s_t = (1 - \rho_s) + \rho_s s_{t-1} + \sigma_{s,t} \varepsilon_{s,t},\ 0 \leq \rho_s < 1,\ \varepsilon_s \sim \mathbb{N}(0,1), \tag{6}$$

$$\sigma_{s,t} = \bar{\sigma}_s (\sigma_{s,t-1}/\bar{\sigma}_s)^{\rho_{\sigma_s}} \exp(\sigma_{\sigma_s} \varepsilon_{\sigma_s,t}),\ 0 \leq \rho_{\sigma_s} < 1,\ \varepsilon_{\sigma_s} \sim \mathbb{N}(0,1), \tag{7}$$

where the standard deviation of the risk premium shock, σ_s, follows an independent log-normal process (σ_s and ε_s are uncorrelated) to introduce time-varying demand uncertainty into the model.

The first order conditions to the household's constrained optimization problem imply

$$\tilde{w}_t = \chi n_t^{\eta} \tilde{c}_t^{\gamma}, \tag{8}$$

$$1 = \beta E_t[(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma}(s_t i_t/(\bar{\pi} \pi_{t+1}^{gap} g_{t+1}))]. \tag{9}$$

Equation (9) is the Euler equation we will use to show the effects of various higher-order moments.

3.3 MONETARY POLICY

The central bank sets the gross nominal interest rate according to

$$i_t = \max\{1, i_t^n\}, \tag{10}$$

$$i_t^n = (i_{t-1}^n)^{\rho_i}(\bar{\imath}(\pi_t^{gap})^{\phi_\pi}(g_t\tilde{y}_t/(\bar{g}\tilde{y}_{t-1}))^{\phi_y})^{1-\rho_i}\exp(\sigma_i\varepsilon_{i,t}),\ 0 \leq \rho_i < 1,\ \varepsilon_i \sim \mathbb{N}(0,1), \tag{11}$$

where y is output (the amount of final goods, y^f, minus the resources lost due to price adjustment costs, adj^f), i^n is the gross notional interest rate, $\bar{\imath}$ and $\bar{\pi}$ are the steady-state or target values of the inflation and nominal interest rates, and ϕ_π and ϕ_y determine the central bank's responses to deviations of inflation from the target rate and deviations of output growth from the balanced growth rate.

3.4 COMPETITIVE EQUILIBRIUM

The aggregate resource constraint is given by

$$\tilde{c}_t = \tilde{y}_t, \tag{12}$$

$$\tilde{y}_t = (1 - \varphi_f(\pi_t^{gap} - 1)^2/2)\tilde{y}_t^f. \tag{13}$$

To make the model stationary, we redefined all of the variables that grow along the balanced growth path in terms of technology (i.e., $\tilde{x}_t \equiv x_t/z_t$). A competitive equilibrium consists of infinite sequences of quantities, $\{\tilde{c}_t, \tilde{y}_t, \tilde{y}_t^f, n_t\}_{t=0}^{\infty}$, prices, $\{\tilde{w}_t, mc_t, i_t, i_t^n, \pi_t^{gap}\}_{t=0}^{\infty}$, and exogenous variables, $\{s_t, g_t, \sigma_{g,t}, \sigma_{s,t}\}_{t=0}^{\infty}$, that satisfy the detrended equilibrium system, (1)-(13), given the initial conditions, $\{c_{-1}, i_{-1}^n, g_0, s_0, \varepsilon_{i,0}, \sigma_{g,0}, \sigma_{s,0}\}$, and the sequences of shocks, $\{\varepsilon_{g,t}, \varepsilon_{s,t}, \varepsilon_{i,t}, \varepsilon_{\sigma_g,t}, \varepsilon_{\sigma_s,t}\}_{t=1}^{\infty}$.

3.5 MEASURES OF UNCERTAINTY

The stochastic volatility processes, (2) and (7), create exogenous sources of time-varying supply and demand uncertainty. Uncertainty is measured by the expected standard deviation of future technology growth and the future risk premium, which equal

$$U_{g,t} \equiv \sqrt{E_t[(g_{t+1} - E_t g_{t+1})^2]} = \sqrt{E_t[\sigma_{g,t+1}^2]},$$

$$U_{s,t} \equiv \sqrt{E_t[(s_{t+1} - E_t s_{t+1})^2]} = \sqrt{E_t[\sigma_{s,t+1}^2]}.$$

We classify these types of uncertainty as exogenous because they fluctuate due to temporary changes in the standard deviation of each shock. For example, if the volatility of technology growth temporarily increases, then supply uncertainty also increases and lowers economic activity.

Uncertainty also arises endogenously in any nonlinear macro model. Following Plante et al. (2018), the endogenous uncertainty surrounding trended output growth, $y_t^g \equiv g_t\tilde{y}_t/\tilde{y}_{t-1}$, is given by

$$U_{y^g,t} \equiv \sqrt{E_t[(y_{t+1}^g - E_t[y_{t+1}^g])^2]}, \tag{14}$$

which is the same way we measure exogenous uncertainty, except it is calculated with an endoge-

nous variable. Both measures of uncertainty remove the predictable component of the forecasted variable instead of only a constant trend, so they distinguish between uncertainty and conditional volatility. However, the endogenous uncertainty measure not only fluctuates due to exogenous volatility shocks, but also due to events that happen in the economy. For example, when the notional interest rate is negative, the economy is more sensitive to first moment shocks that adversely affect the economy, which increases the endogenous uncertainty about output growth. The ZLB constraint also creates uncertainty by amplifying the effect of the two exogenous volatility shocks.

4 NUMERICAL METHODS AND DECOMPOSITION

4.1 SOLUTION METHOD We solve the nonlinear model with the policy function iteration algorithm described in Richter et al. (2014), which is based on the theoretical work on monotone operators in Coleman (1991). The presence of stochastic volatility complicates the solution method because the realizations of g and s depend on the realizations of the stochastic volatility processes.

We discretize the state space and then approximate the stochastic volatility processes, (2) and (7), and first moment shocks, ε_g, ε_s, and ε_i, using the N-state Markov chain described in Rouwenhorst (1995). The Rouwenhorst method is attractive because it only requires us to interpolate along the dimensions of the endogenous state variables, which makes the solution more accurate and faster than quadrature methods. For each combination of the first and second moment shocks, we calculate the future realizations of technology and the risk premium according to (1) and (6). To obtain initial conjectures for the nonlinear policy functions, we solve the log-linear analogue of our nonlinear model with Sims's (2002) gensys algorithm. Then we minimize the Euler equation errors on every node in the discretized state space and compute the maximum distance between the updated policy functions and the initial conjectures. Finally, we replace the initial conjectures with the updated policy functions and iterate until the maximum distance is below the tolerance level.

The algorithm produces policy functions for consumption and inflation. To estimate the model, we also create a policy function for output growth uncertainty, (14), by interpolating the implied policy function for output and integrating. See Appendix C for a description of the solution method.

4.2 ESTIMATION PROCEDURE We estimate the nonlinear model with quarterly data on per capita real GDP, $RGDP/CNP$, the GDP implicit price deflator, DEF, the federal funds rate, FFR, the macro uncertainty series in Jurado et al. (2015), UM, and the financial uncertainty series in Ludvigson et al. (2017), UF, from 1986Q1 to 2016Q2. The vector of observables is given by

$$\hat{\mathbf{x}}_t^{data} \equiv [\Delta \log(RGDP_t/CNP_t),\ \Delta \log(DEF_t),\ \log(1 + FFR_t/100)/4,\ z(UM_t),\ z(UF_t)],$$

where Δ denotes a difference and $z(\cdot)$ is a standardized variable. Appendix A provides our sources.

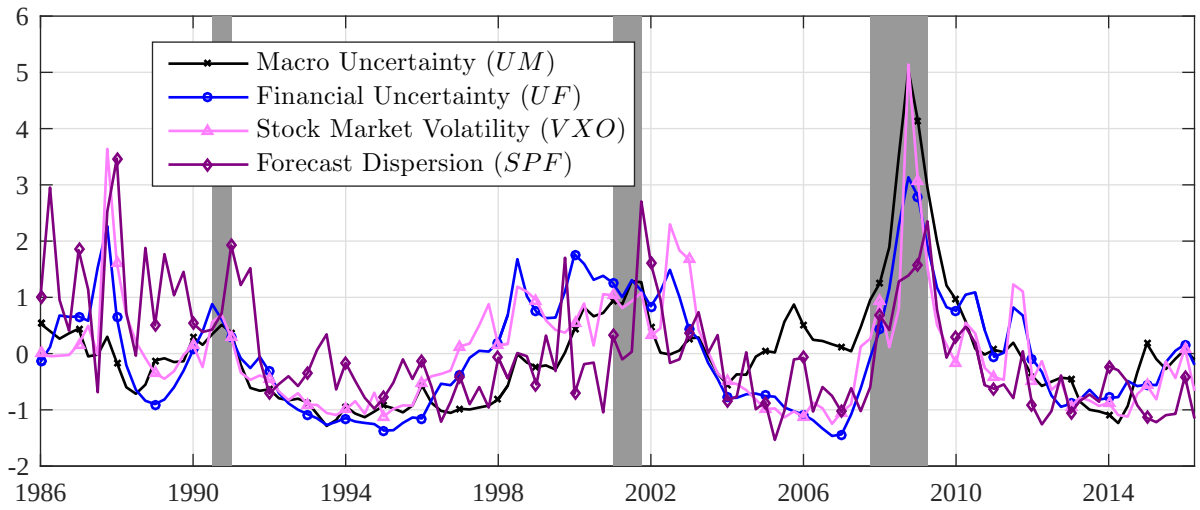

Figure 1: Measures of uncertainty in the data.

Figure 1 plots the standardized 1-quarter ahead UM and UF series, which inform the parameters in our model and ensure it produces the same fluctuations in uncertainty as the data. The series are based on a factor augmented VAR that accounts for 132 macroeconomic and 147 financial variables. Repeated simulations of the FAVAR are used to obtain estimates of uncertainty for each macro (financial) variable and then averaged to obtain the UM (UF) time series. The benefit of these series is that they are calculated the same way as (14), so they distinguish between uncertainty and conditional volatility and reflect the uncertainty surrounding a rich set of variables.

For comparison, we also plot two other popular measures of uncertainty: the Chicago Board Options Exchange S&P 100 Volatility Index (VXO) and the dispersion in forecasts of real GDP growth 1-quarter ahead from the Survey of Professional Forecasters (SPF). The different uncertainty measures generally move together, but they also show significant independent variation. For example, sharp increases in the VXO, SPF, and UF series occur with some regularity, but they are far less frequent in the UM series. After the start of the Great Recession, the correlations between the uncertainty measures all exceeded 0.7, but they are near 0.4 prior to that date. The one exception is the correlation between UF and the VXO, which was above 0.8 in both subperiods.

We calibrate four parameters (table 1). The subjective discount factor, β, is set to 0.9987. The preference parameter, χ, is set so the labor supply along the balanced growth path equals $1/3$ of the available time. The elasticity of substitution between goods, θ, is set to 6, which matches the estimate in Christiano et al. (2005) and corresponds to a 20% average markup of price over marginal cost. The Frisch labor supply elasticity, $1/\eta$, is set to 3, to match the estimate in Peterman (2016).

We use Bayesian methods to estimate the remaining parameters in our model. For each draw from the parameter distribution, we solve the nonlinear model and approximate the likelihood using a particle filter. We determine whether to accept a draw with a random walk Metropolis-

Balanced Growth Discount Factor	$\bar{\beta}$	0.9987	Real GDP Growth Rate ME SD	σ_{me,y^g}	0.00268
Frisch Elasticity of Labor Supply	$1/\eta$	3	Inflation Rate ME SD	$\sigma_{me,\pi}$	0.00109
Elasticity of Substitution	θ	6	Federal Funds Rate ME SD	$\sigma_{me,i}$	0.00094
Balanced Growth Labor Supply	$\bar{n}$	0.33	Macro Uncertainty ME SD	$\sigma_{me,um}$	0.44721
Number of Particles	N_p	40,000	Financial Uncertainty ME SD	$\sigma_{me,uf}$	0.44721

Table 1: Calibrated parameters for the nonlinear model and particle filter.

Hastings algorithm. The filter uses $40{,}000$ particles and systematic resampling with replacement following Kitagawa (1996). To help the model better match outliers during the Great Recession, we adapt the particle filter described in Fernández-Villaverde and Rubio-Ramírez (2007) to include the information contained in the current observation according to Algorithm 12 in Herbst and Schorfheide (2016). See Appendix D for a more complete description of our estimation procedure.

A major difference from other filters is that the particle filter requires measurement error (ME) to avoid degeneracy—a situation when all but a few particle weights are near zero, so the equation linking the observables to equivalent variables in the model is given by $\hat{\mathbf{x}}_t^{data} = \hat{\mathbf{x}}_t^{model} + \xi_t$, where

$$\hat{\mathbf{x}}_t^{model} = [\log(y_t^g),\ \log(\pi_t),\ \log(i_t),\ z(U_{y^g,t}),\ z(U_{s,t})],$$

$\xi \sim \mathbb{N}(0, \Sigma)$ is a vector of MEs and $\Sigma = \text{diag}([\sigma^2_{me,y^g}, \sigma^2_{me,\pi}, \sigma^2_{me,i}, \sigma^2_{me,um}, \sigma^2_{me,uf}])$. Following Herbst and Schorfheide (2016), we set the ME variances to 20% of their variance in the data, except the ME variance for the policy rate is set to 2% because the federal funds rate is less noisy and it affects the level of uncertainty predicted by the model near the ZLB. We link output uncertainty to the macro uncertainty index and risk premium uncertainty to the financial uncertainty index because Ludvigson et al. (2017) find financial uncertainty is an exogenous impulse that causes recessions, whereas macro uncertainty endogenously responds to other shocks that affect the business cycle. In our model, output uncertainty is endogenous, whereas risk premium uncertainty is exogenous.

The entire algorithm is programmed in Fortran using Open MPI and executed on a cluster with 512 cores. We parallelize the nonlinear solution by distributing the nodes in the state space across the available cores. To increase the accuracy of the filter, we calculate the model likelihood on each core and then evaluate whether to accept a candidate draw based on the median likelihood. This important step reduces the variance of the model likelihood across multiple runs of the particle filter.

Our estimation procedure has three stages. First, we conduct a mode search to create an initial variance-covariance matrix for the parameters. The covariance matrix is based on the parameters corresponding to the 90th percentile of the likelihoods from $5{,}000$ draws. Second, we perform an initial run of the Metropolis Hastings algorithm with $25{,}000$ draws from the posterior distribution. We burn off the first $5{,}000$ draws and use the remaining draws to update the variance-covariance matrix from the mode search. Third, we conduct a final run of the Metropolis Hastings algorithm.

We obtain $100{,}000$ draws from the posterior distribution and then thin by 100 to limit the effects of serial correction in the parameter draws, so our posterior distribution has a sample of $1{,}000$ draws.

4.3 EULER EQUATION DECOMPOSITION Our goal is to determine how changes in uncertainty affect output, taking into account all first and second moment shocks as well as endogenous dynamics. One way to quantify the effect of uncertainty in a given period is by decomposing output with the Euler equation, (9). A third-order approximation around the balanced growth path implies

$$\begin{aligned}\hat{y}_t \approx{}& E_t\hat{y}_{t+1} - \tfrac{1}{\gamma}\hat{r}_t - \operatorname{cov}_t(\hat{\pi}_{t+1}, \hat{y}_{t+1}) - \operatorname{cov}_t(\hat{g}_{t+1}, \hat{y}_{t+1}) - \tfrac{1}{\gamma}\operatorname{cov}_t(\hat{\pi}_{t+1}, \hat{g}_{t+1}) \\ &- \tfrac{1}{2\gamma}(\operatorname{var}_t \hat{g}_{t+1} + \operatorname{var}_t \hat{\pi}_{t+1} + \gamma^2 \operatorname{var}_t \hat{y}_{t+1}) + \tfrac{1}{6\gamma}(\operatorname{skew}_t \hat{g}_{t+1} + \operatorname{skew}_t \hat{\pi}_{t+1} + \gamma^3 \operatorname{skew}_t \hat{y}_{t+1}),\end{aligned} \tag{15}$$

where var_t, skew_t, and cov_t denote the variance, third moment, and covariance of a variable conditional on information at time t, $\hat{r}_t \equiv \hat{\imath}_t + \hat{s}_t - E_t\hat{\pi}_{t+1} - E_t\hat{g}_{t+1}$ is the *ex-ante* real interest rate, and a hat denotes log deviation from the balanced growth path.[2] Appendix B provides the derivation.

We omit higher-order covariance terms as well as fourth-order and higher terms because they had almost no effect on output in our sample. The variance, skewness, and covariance terms quantify the effect of the uncertainty, upside and downside risk, and the pairwise linear relationships between output, inflation, and technology growth. Higher risk aversion makes households less willing to intertemporally substitute consumption goods, which makes them less sensitive to the real interest rate and more sensitive to the variance and skewness of output. Most of our analysis will focus on the variance of output. That term will have the same effect on output regardless of which Euler equation is used for the decomposition because the pricing kernel always enters in the same way.

The decomposition shows how the different types of uncertainty and skewness affect current output over a 1-quarter horizon. If we recursively substitute for expected future output, we obtain

$$\begin{aligned}\hat{y}_t \approx{}& E_t\hat{y}_{t+q} - \tfrac{1}{\gamma}E_t\textstyle\sum_{j=1}^{q}\hat{r}_{t+j-1} \\ &- \textstyle\sum_{j=1}^{q}(\operatorname{cov}_t(\hat{\pi}_{t+j}, \hat{y}_{t+j}) + \operatorname{cov}_t(\hat{g}_{t+j}, \hat{y}_{t+j}) + \tfrac{1}{\gamma}\operatorname{cov}_t(\hat{\pi}_{t+j}, \hat{g}_{t+j})) \\ &- \tfrac{1}{2\gamma}\textstyle\sum_{j=1}^{q}(\operatorname{var}_t \hat{g}_{t+j} + \operatorname{var}_t \hat{\pi}_{t+j} + \gamma^2 \operatorname{var}_t \hat{y}_{t+j}) \\ &+ \tfrac{1}{6\gamma}\textstyle\sum_{j=1}^{q}(\operatorname{skew}_t \hat{g}_{t+j} + \operatorname{skew}_t \hat{\pi}_{t+j} + \gamma^3 \operatorname{skew}_t \hat{y}_{t+j}),\end{aligned} \tag{16}$$

where $q \geq 1$ is the forecast horizon. The sum of each variance term over q quarters captures the effect of a given type of uncertainty, conditional on expected output in quarter q. When q becomes

[2] Decompositions of equilibrium conditions have been used to study other topics. Basu and Bundick (2015) derive a similar decomposition to ours in an endowment economy model to provide intuition for how the Fed can offset the effects of uncertainty at and away from the ZLB, but they do not quantify the terms. Parker and Preston (2005) use the Euler equation to decompose consumption growth into a forecast error, the real interest rate, a measure of preferences, and a precautionary saving channel. Chung and Leeper (2007), Hall and Sargent (2011), Berndt et al. (2012), and Mason and Jayadev (2014) all use the government budget constraint to determine the key drivers of government debt.

sufficiently large, the conditional expectation drops out of the decomposition, so we are able to determine the unconditional effects of each higher-order moment. Expected output can hide the effect of higher-order moments in future quarters. By decomposing expected future output, we can show how the uncertainty, skewness, and covariance terms affect output over horizons beyond 1 quarter.[3]

Given a draw from the posterior distribution, we quantify the effect of each term on output in three steps. First, we create policy functions for the $10q+1$ variables in the decomposition by integrating across $10{,}000$ q-quarter simulations initialized at each node in the state space. Although the variables are represented in deviations from the balanced growth path, the policy functions inherit the nonlinearities from the solution. Second, we create time series for the variables in the decomposition at each horizon by interpolating the policy functions at the median filtered states and shocks in each time period. Third, we weight each variable by its coefficient in the decomposition.

5 ESTIMATED EFFECTS OF UNCERTAINTY

We first show the posterior parameter distributions, impulse responses, and sources of output uncertainty. Then we show the results of our Euler equation decomposition and analyze which parameters are most important. The section concludes by showing the welfare cost of business cycles.

5.1 PRIOR AND POSTERIOR DISTRIBUTIONS

The first four columns of table 2 display the estimated parameters and information about the priors. The prior for the coefficient of relative risk aversion is taken from An and Schorfheide (2007). The priors for the steady state growth rate and the target inflation rate are set to the average per capita GDP growth rate and the average inflation rate over our sample period. The priors for the monetary policy parameters, which follow Guerrón-Quintana and Nason (2013), are chosen so the distributions cover the values in Taylor (1993) as well as stronger responses that could explain data during the ZLB period. The priors for the persistence parameters are diffuse, but all of the means, except for the growth rate, are set to 0.6 since a modest degree of persistence is needed to explain the data. The priors for the standard deviations are also diffuse but less diffuse than in An and Schorfheide (2007) and Smets and Wouters (2007), since our nonlinear model generates more volatility than analogous unconstrained linear models.

The last four columns display the posterior means, standard deviations, and 90% credible sets for the estimated parameters. Low frequency movements in the macro and financial uncertainty time series coupled with sharp increases in both series during the Great Recession generate highly persistent stochastic volatility processes with large shock standard deviations. For example, a two standard deviation supply uncertainty shock causes a 25.1% increase in the volatility of technology growth with a half-life of about 15.9 quarters. The monetary policy parameters imply a high degree

[3] After iterating, we obtain $E_t[\mathrm{cov}_{t+j}(x_{t+j+1}, y_{t+j+1})] = \mathrm{cov}_t(x_{t+j}, y_{t+j}) - \mathrm{cov}_t(E_{t+j}[x_{t+j+1}], E_{t+j}[y_{t+j+1}])$ by the law of total covariance. In our derivation, we ignore the second term because its effects are quantitatively small.

	Prior			Posterior			
Parameter	Dist.	Mean	SD	Mean	SD	5%	95%
Risk Aversion (γ)	Gamm	2.0000	0.5000	3.00551	0.44806	2.35252	3.81243
Price Adjustment Cost (φ_f)	Norm	100.0000	20.0000	141.00914	19.95554	110.36190	175.77686
Inflation Response (ϕ_π)	Norm	2.0000	0.2500	2.54332	0.19854	2.21212	2.85598
Output Response (ϕ_y)	Norm	0.5000	0.2000	1.04678	0.15152	0.79593	1.29649
Average Growth ($\bar{g}$)	Norm	1.0040	0.0010	1.00439	0.00058	1.00337	1.00534
Average Inflation ($\bar{\pi}$)	Norm	1.0055	0.0010	1.00649	0.00041	1.00579	1.00718
Int. Rate Persistence (ρ_i)	Beta	0.6000	0.2000	0.84086	0.01902	0.80740	0.87024
Growth Persistence (ρ_g)	Beta	0.4000	0.2000	0.51433	0.12352	0.29503	0.70706
Risk Persistence (ρ_s)	Beta	0.6000	0.2000	0.91050	0.01084	0.89163	0.92723
Growth SV Persistence (ρ_{σ_g})	Beta	0.6000	0.2000	0.95721	0.01890	0.92614	0.98109
Risk SV Persistence (ρ_{σ_s})	Beta	0.6000	0.2000	0.93308	0.01617	0.90404	0.95725
Int. Rate Shock SD (σ_i)	IGam	0.0025	0.0025	0.00127	0.00017	0.00102	0.00157
Growth Shock SD ($\bar{\sigma}_g$)	IGam	0.0075	0.0075	0.00371	0.00054	0.00288	0.00463
Risk Shock SD ($\bar{\sigma}_s$)	IGam	0.0025	0.0025	0.00139	0.00022	0.00107	0.00177
Growth SV Shock SD (σ_{σ_g})	IGam	0.1000	0.0250	0.11216	0.02350	0.07647	0.15372
Risk SV Shock SD (σ_{σ_s})	IGam	0.1000	0.0250	0.11855	0.02218	0.08428	0.15666

Table 2: Prior and posterior distributions of the estimated parameters. The last two columns show the 5th and 95th percentiles of each marginal posterior distribution. The model is estimated with quarterly data from 1986Q1 to 2016Q2.

of interest rate smoothing and strong responses to real GDP growth and inflation, which are necessary for the model to explain the long ZLB period. The mean estimates of the annualized technology growth and inflation rates are 1.77% and 2.62%, which are slightly higher than the values in the data since they are unconditional and must compensate for the expectation of the ZLB period. The mean coefficient of relative risk aversion is consistent with An and Schorfheide (2007). The price adjustment cost parameter implies a slope of the Phillips curve of about 0.035, which is in line with other estimates in the literature. Overall, the priors and posterior means are consistent with Gust et al. (2017), who estimate a similar model with an occasionally binding ZLB constraint but without stochastic volatility. Appendix G provides additional estimation results, including the kernel densities of the parameters, median filtered observables and shocks, and unconditional moments.

5.2 IMPULSE RESPONSES We begin our analysis by showing impulse responses to first and second moment shocks to illustrate the underlying dynamics in the model. Figure 2 plots the responses to a 2 standard deviation positive risk premium, risk premium volatility, growth, and growth volatility shock. The parameters are set to their posterior means and the simulations are initialized at two different states. Our benchmark simulation is initialized at the stochastic steady state and reflective of any state of the economy where there is little expectation of hitting the ZLB.

We compare the baseline impulse responses to the responses when the notional rate is negative by initializing the simulation at the filtered state vector corresponding to 2009Q2. The effect of mean reversion is removed from the responses by plotting the percentage point difference (percent

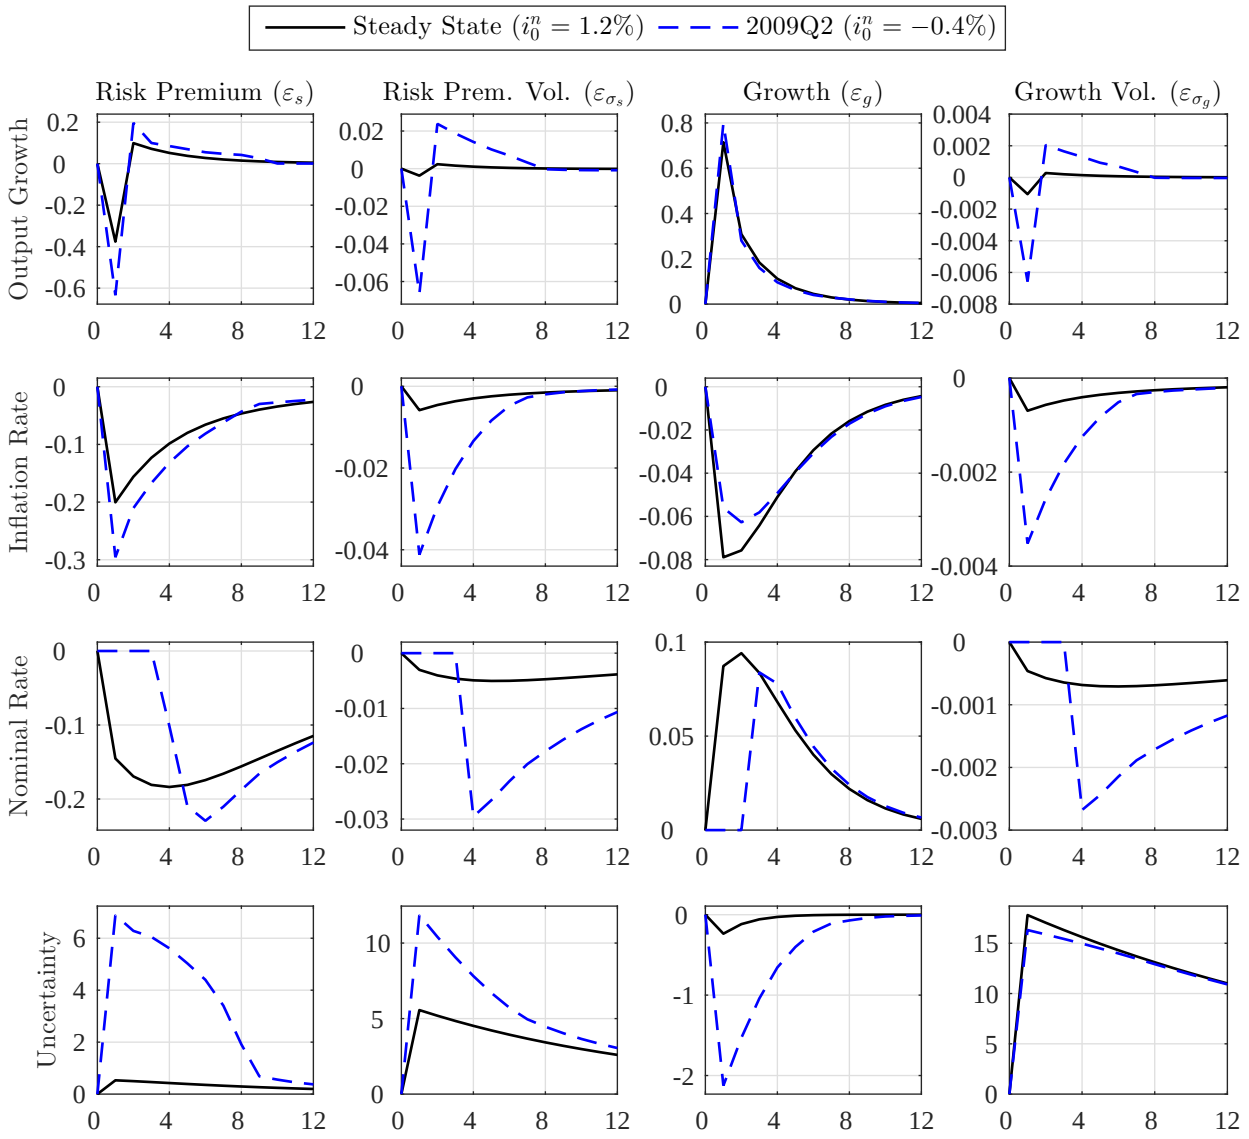

Figure 2: Impulse responses to a 2 standard deviation positive shock at and away from the ZLB. The steady-state simulation (solid line) is initialized at the stochastic steady state. The other simulation (dashed line) is initialized at the filtered state corresponding to 2009Q2 so the ZLB binds. The vertical axes are in percentage point deviations from the baseline simulation, except uncertainty is a percent change. The horizontal axes denote the time period in quarters.

change for output uncertainty) from a counterfactual simulation without a shock in the first quarter. The risk premium and growth volatilities are initialized at their stochastic steady states in both simulations, so the level shocks are not amplified by exogenous changes in volatility over time and the impact effects of the volatility shocks are not distorted by the log-normal volatility processes.

A higher risk premium (first column) in either initial state causes households to postpone consumption, which lowers output growth and inflation on impact. When the Fed is not constrained by the ZLB, it responds to the shock by reducing its policy rate. The impact on output uncertainty is small since the Fed is able to stabilize the economy. In 2009Q2, the higher risk premium leads to

an expected ZLB duration of 2 quarters on impact. The Fed cannot respond by lowering its policy rate, which causes a larger decline in output growth. The result is a larger increase in output uncertainty since households expect a wider range of future realization of output growth. In other words, the model endogenously generates uncertainty when the ZLB binds due to a risk premium shock.

Similar to the level shock, a positive shock to the volatility of the risk premium (second column) lowers output growth and inflation. In steady state, the Fed adjusts its policy rate to stabilize the economy, so the effect of the volatility shock is small even though output uncertainty rises far more than it does in response to the level shock. When the ZLB binds, however, the increase in uncertainty nearly doubles, which magnifies the effect on output growth and inflation. Hence, the model also endogenously creates uncertainty by amplifying the effects of second moment shocks.

Level and volatility shocks to technology growth have qualitatively and quantitatively different effects than risk premium shocks. A positive shock to technology growth (third column) increases output growth and decreases inflation like a typical supply shock, so the Fed faces a tradeoff between stabilizing inflation and output growth unlike with a risk premium shock. In steady state, the policy rate immediately increases since the response to the output gap dominates the response to the inflation gap. The ZLB initially binds in 2009Q2, but the increase in the notional rate causes a quick exit from the ZLB after 1 quarter. The delayed increase in the policy rate causes a slightly larger boost in output growth and a smaller decline in inflation. In contrast with the risk premium shock, a positive technology growth shock causes output uncertainty to decline because it reduces the probability that the ZLB binds next period. However, the responses are smaller in magnitude.

Growth volatility shocks cause bigger changes in uncertainty than level shocks. Similar to a risk premium volatility shock, a positive growth volatility shock (fourth column) reduces output growth and inflation, which leads to a lower nominal interest rate. However, the responses differ in a few ways. One, growth volatility directly affects output volatility. Therefore, uncertainty increases more than it does in response to a risk premium volatility shock. Two, the response of output uncertainty is similar in both initial states. Three, the increase in output uncertainty away from the ZLB is much larger than the increase from a risk premium volatility shock. Therefore, growth volatility shocks play a larger role in explaining the fluctuations in uncertainty when the ZLB does not bind.

5.3 SOURCES OF UNCERTAINTY

The impulse responses show uncertainty is time-varying due to exogenous volatility shocks or first moment shocks that interact with the state of the economy. Figure 3a decomposes output growth uncertainty into its exogenous and endogenous sources using counterfactual simulations conditional on the posterior mean parameters of our model. To isolate the contribution of technology growth uncertainty, we turn off the risk premium volatility shocks. Similarly, we zero out the technology growth volatility shocks to identify the amount of risk premium uncertainty. We then turn off both volatility shocks to determine the amount of endogenous

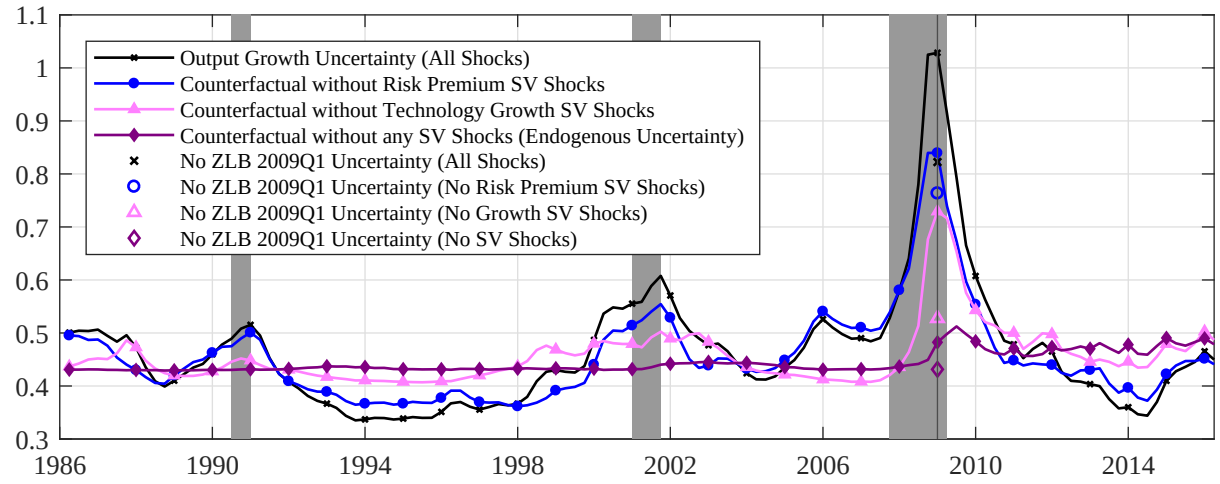

(a) Decomposition of the exogenous and endogenous sources of uncertainty.

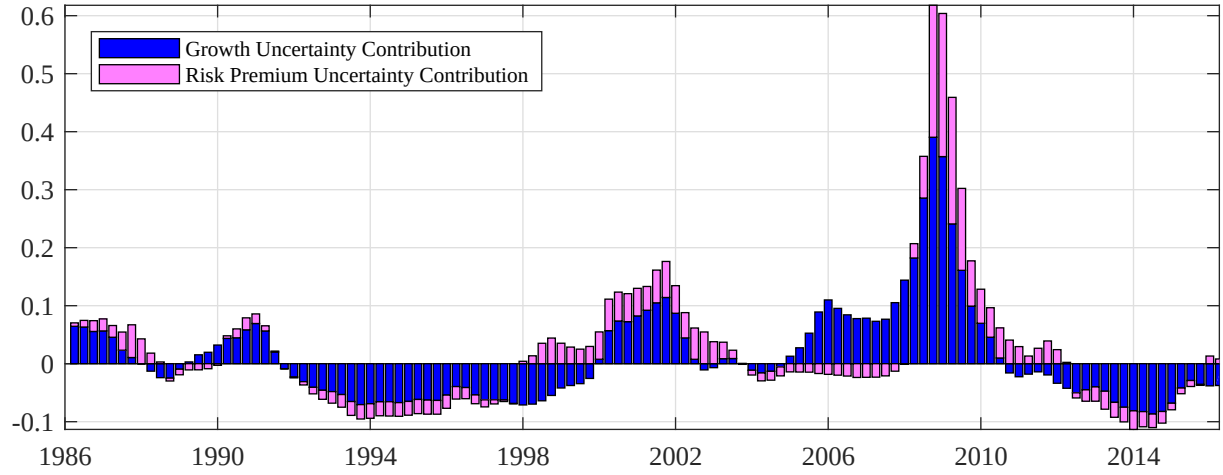

(b) Relative contribution of the exogenous sources of uncertainty.

Figure 3: Sources of output growth uncertainty in our baseline model.

uncertainty. We also show the endogenous amplification of the exogenous volatility shocks when the Fed was most constrained using the solution to the unconstrained nonlinear model.

On average about 95% of output growth uncertainty is due to the uncertainty that occurs without second moment shocks, which we refer to as endogenous uncertainty. However, most of the *changes* in uncertainty are driven by the exogenous volatility shocks. Growth volatility shocks are the key driver in most periods, but risk premium volatility shocks play an important role in certain parts of our sample. Typically, endogenous uncertainty is fairly constant, but it increases when the policy rate is near or at its ZLB, which occurs in the mid 2000s and from 2009 to the end of the sample. The sharp increase in uncertainty in 2009, however, primarily occurred due to the endogenous amplification of the exogenous volatility shocks, rather than through first moment shocks. The markers in 2009Q1 show the counterfactual increase in uncertainty that would have occurred if the

Fed was not constrained. Those results indicate that about 8.5% ($(0.48-0.43)/(1.03-0.43)$) of the increase in uncertainty in 2009Q1 was due to endogenous uncertainty and about 38% ($(1.03-0.82)/(1.03-0.48)$) was due to the endogenous amplification of second moment shocks.

Despite some nonlinear interactions between the exogenous volatility shocks and the ZLB, we are able to approximate the relative contribution of each volatility shock over time, similar to a variance decomposition in a linear model. The dark bars in figure 3b represent the technology growth counterfactual relative to the endogenous uncertainty counterfactual (circles minus diamonds) and the light bars represent the risk premium counterfactual relative to the endogenous uncertainty counterfactual (triangles minus diamonds), which is approximately equal to output growth uncertainty relative to the endogenous uncertainty counterfactual (solid minus diamonds).

The results reiterate that technology growth uncertainty is typically the biggest contributor to changes in output growth uncertainty, but the two sources of exogenous uncertainty typically move together. There are two notable exceptions. One, the model predicts that risk premium uncertainty precedes the 2001 recession. Two, technology growth uncertainty increases before the rise in risk premium uncertainty during the Great Recession, but the effects of risk premium uncertainty linger while the impact of technology growth uncertainty is negligible for a few years after the Great Recession. During the Great Recession, technology growth and risk premium volatility shocks have nearly equal roles. By the end of the sample, output growth uncertainty declined to its lowest point.

5.4 EULER EQUATION DECOMPOSITION The rest of this section focuses on the effects of uncertainty and other higher-order moments. Figure 4 shows a filtered time series of each term in the Euler equation decomposition in (16) over different forecast horizons. The values on the vertical axes are the effects on current output in percentage point deviations from the balanced growth path.

The top row shows the decomposition over a 1-quarter horizon. The changes in output are almost entirely driven by expectations about output next quarter. The real interest rate had a smaller role, typically reducing output by about 0.1%. The peak effect was -0.37% during the Great Recession, but that effect quickly declined as the economy rebounded. The higher-order terms show output uncertainty had time-varying adverse effects on current output. The largest effect occurred during the Great Recession, since the ZLB constraint made the economy more sensitive to adverse shocks. However, that effect was short-lived because the notional rate was negative only until 2011.

The effects of uncertainty were small throughout our sample. Even during the Great Recession, the peak increase in uncertainty reduced output by less than 0.01%. Output skewness and both inflation uncertainty and inflation skewness also had very small effects. Interestingly, the effects of uncertainty over a 1-quarter horizon were similar to the impact effects of the exogenous volatility shocks shown in figure 2. However, the results understate the effects of uncertainty because they hide the impact that future real interest rates and higher-order moments have on expected output.

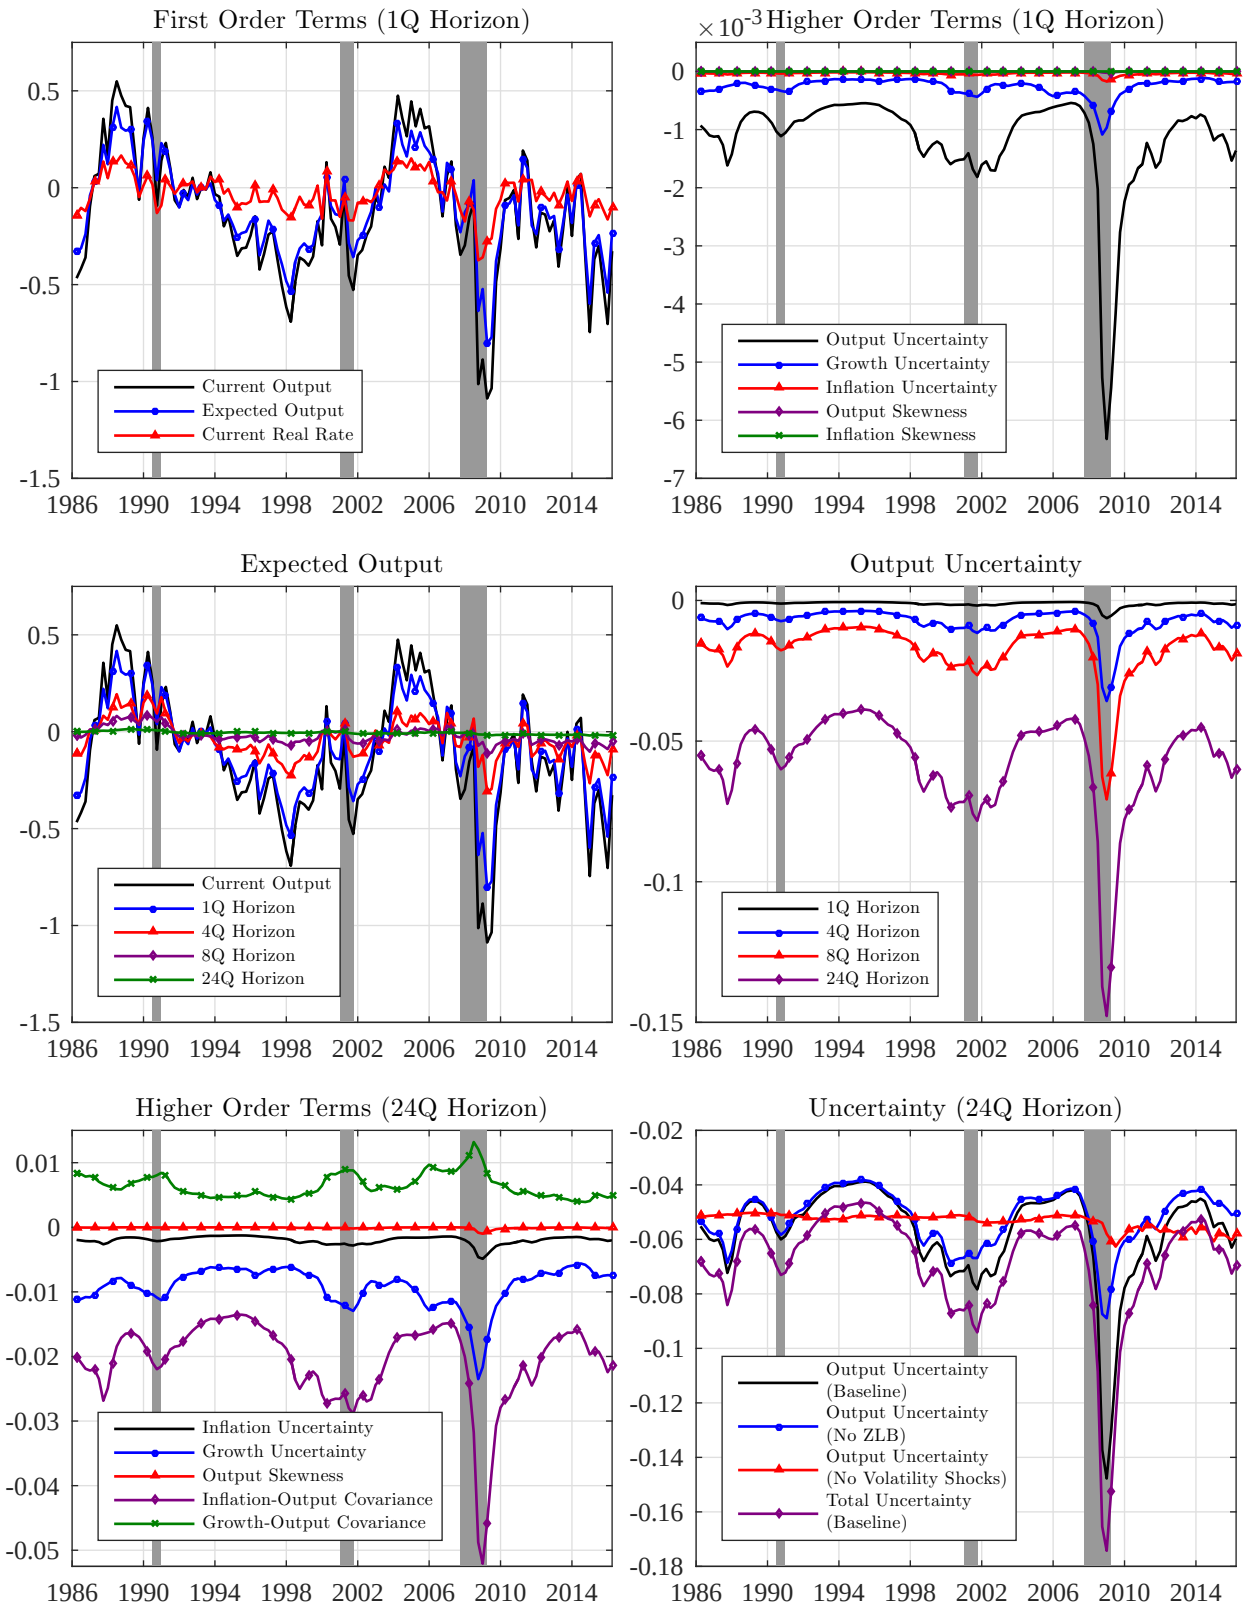

Figure 4: Filtered decomposition of the effects on current output. The shaded regions denote NBER recessions. The values on the vertical axes are the contributions to the percentage point deviation of detrended output from steady state.

The middle left panel shows how expected output affected current output over horizons up to 24 quarters. In most periods, the differences between current and expected output were much larger over horizons beyond 1-quarter, which indicates that other factors, such as the real interest rate and uncertainty, explained a larger fraction of the changes in output. We focus on a 24-quarter horizon because it is long enough that expected output barely matters for current output. For example, in 2009Q2 (the last quarter of the Great Recession) expected output in 2009Q3 explained 74.3% of the decline in current output, whereas expected output in 2015Q3 explained only 1.6% of the decline. Over those same horizons, the contribution of output uncertainty increased from 2.9% to 11.9%.

The middle right panel shows the effect of output uncertainty over the horizons shown in the left panel, but the values on the vertical axis are cumulative effects. Although the effect of output uncertainty is small when it is conditional on expected output over a 1-quarter horizon, it is more significant over longer horizons that decompose the influence of expected future output. Over a 24-quarter horizon, output uncertainty on average decreased current output by about 0.06% and the largest effect was about 0.15% in 2009Q1, which accounted for 16.6% of the decline in that quarter.

The other higher order moments are shown in the bottom left panel. During the Great Recession, the peak effects of technology growth uncertainty, inflation uncertainty, and output skewness over a 24-quarter horizon were -0.023%, -0.005%, and -0.001%, respectively, and the average effects were much smaller. We do not show the effect of inflation skewness because it is always near zero. It is not surprising that inflation uncertainty and skewness had small effects because the Fed aggressively targeted inflation throughout our sample. However, we expected a larger effect of output skewness, especially during the Great Recession. The ZLB creates downside risk since it prevents the Fed from responding to adverse shocks through conventional channels. Evidently, those effects are small when controlling for other moments. The covariance between inflation and output on average lowered output by 0.02%, the second largest effect behind output uncertainty.

The bottom right panel shows the effect of output uncertainty over a 24-quarter horizon along with two of the counterfactuals shown in figure 3a. First, we plot the effect of output uncertainty after removing the influence of the ZLB using the solution to the unconstrained nonlinear model. The differences from the baseline path show how much the ZLB increased the adverse effects of uncertainty. In most quarters, the differences are small because there is a low probability of going to and staying at the ZLB. Larger differences between the two paths occurred from 2008Q4 to 2009Q4, when the notional interest rate was well below zero. For example, in 2009Q1 output uncertainty reduced output by about 0.06 percentage points more than if the Fed was not constrained.

Second, we restore the ZLB constraint but zero out both exogenous volatility shocks. Since first moment shocks are the main source of most of the uncertainty in the economy, they are also the primary source of the adverse effects of output uncertainty. At its peak, endogenous uncertainty only increased the adverse effects of output uncertainty by about 0.01 percentage points, whereas

exogenous volatility shocks played a much larger role during the last two recessions. For example, the volatility shocks without the ZLB contributed about 0.04 percentage points to the decline in output growth in 2009Q1 and their amplification contributed an additional 0.06 percentage points.

The bottom right panel also plots the total effect of uncertainty—the sum of output, technology growth, and inflation uncertainty—over a 24-quarter horizon. On average, the three sources of uncertainty lowered output by 0.07 percentage points with a peak decline of 0.17 percentage points.

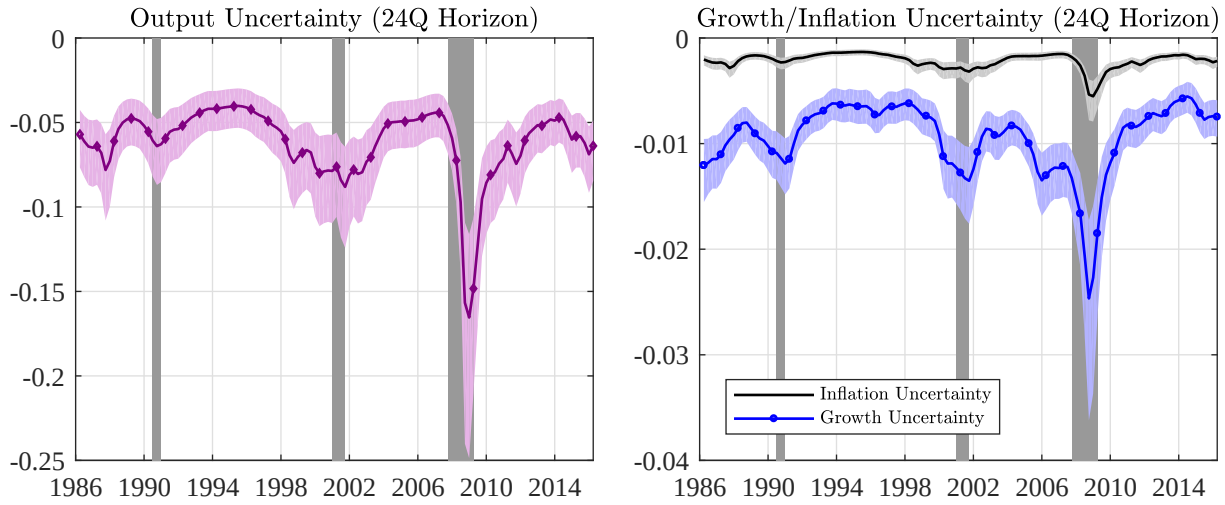

Figure 5: 68% credible sets of the filtered effects on current output. The vertical bars denote NBER recessions. The values on the vertical axes are the contributions to the percentage point deviation of detrended output from steady state.

The results in figure 4 are based on the posterior mean. However, we can generate time series for every draw from the posterior distribution. Figure 5 shows the 1 standard deviation (16%-84%) credible sets for the effects of each type uncertainty over a 24-quarter horizon. In a typical quarter, the effect of output uncertainty ranges from -0.01% to $+0.015\%$ of the median effect. The effects, however, are more asymmetric during recessions. For example, during the peak of the Great Recession there was a 68% chance output uncertainty decreased current output by at least 0.12% and it could have decreased it by as much as 0.25%. The effects of technology growth and inflation uncertainty are always much smaller than output uncertainty, even in the tail of the parameter distribution. In all three cases, the credible sets are much tighter than the range of estimates in the literature.

5.5 KEY PARAMETERS We determine the relative importance of each parameter for our Euler equation decomposition by conducting posterior predictive analysis with the draws from the posterior distribution, $\{\hat{\theta}_i\}_{i=1}^{1000}$. We focus on the effect of output uncertainty over a 24-quarter horizon,

$$h(\theta, t) \equiv -(\gamma/2) \sum_{j=1}^{24} \text{var}_t(\hat{y}_{t+j}|\theta, \bar{\varepsilon}_t, \bar{z}_t),$$

where $\bar{\varepsilon}_t$ and $\bar{z}_t$ are the median filtered states and shocks conditional on the posterior mean parameters. Fixing the states and shocks isolates the role of each parameter. We calculate $h(\hat{\theta}_i, t)$ for all i

to generate the credible sets in figure 5, which capture the effects of output uncertainty given each posterior draw. Define $\bar{\theta}_{i,\ell}$ as the ith posterior draw conditional on the posterior mean of parameter ℓ. As a counterfactual, we first calculate $h(\bar{\theta}_{i,\ell}, t)$ using the procedure described at the end of section 4. We then calculate the root-mean square-deviation (RMSD) from the counterfactual given by

$$RMSD(\ell, t) = \sqrt{\frac{1}{1000}\sum_{i=1}^{1000}(h(\hat{\theta}_i, t) - h(\bar{\theta}_{i,\ell}, t))^2}.$$

Figure 6 plots time series of the RMSD for the nine most consequential parameters in the model. A high RMSD for parameter ℓ means the effect of output uncertainty on current output is sensitive to that parameter. The risk premium persistence (ρ_s) and shock standard deviation (σ_s) have the largest average RMSDs. Of the deep parameters, the coefficient of relative risk aversion (γ) and the monetary response to inflation (ϕ_π) are the most important parameters. There is also considerable variation in the importance of the parameters across time. For example, during recessions the RMSD of each parameter increases, but the process parameters (right panel) become relatively more important than the deep parameters. Outside recessions, the deep parameters (left panel) are relatively more important, though the average RMSD of each parameter is much lower.

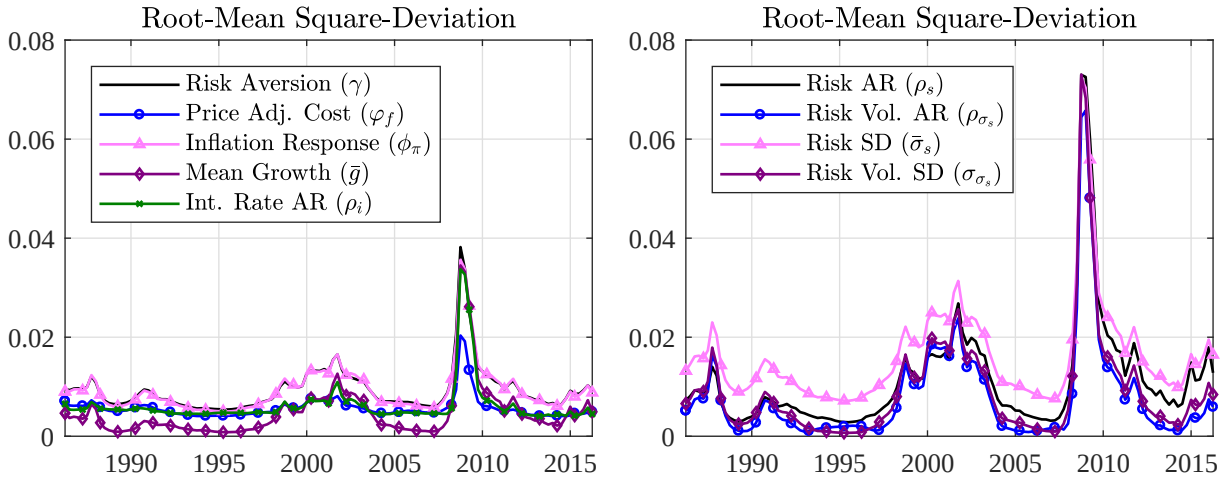

Figure 6: Time series of the root-mean square-deviation of the effect of output uncertainty over a 24-quarter horizon (i.e., the effect on current output at the posterior draw minus the effect after fixing a parameter at its posterior mean).

The RMSD statistic summarizes the importance of a given parameter in every quarter of our sample, but it does not show whether that parameter increases or decreases the effect of uncertainty. By conditioning on a particular quarter, we can determine the sign. Figure 7 shows scatter plots of the deviation, $\Delta_{i,\ell,t} \equiv h(\hat{\theta}_i, t) - h(\bar{\theta}_{i,\ell}, t)$, for all posterior draws, conditional on parameter ℓ and $t = 2008Q4$. In other words, it shows the changes in the effect of output uncertainty that occur when a given parameter deviates from its posterior mean. A positive (negative) value of $\Delta_{i,\ell,t}$ means output uncertainty has a smaller (larger) adverse effect on current output for a given draw.

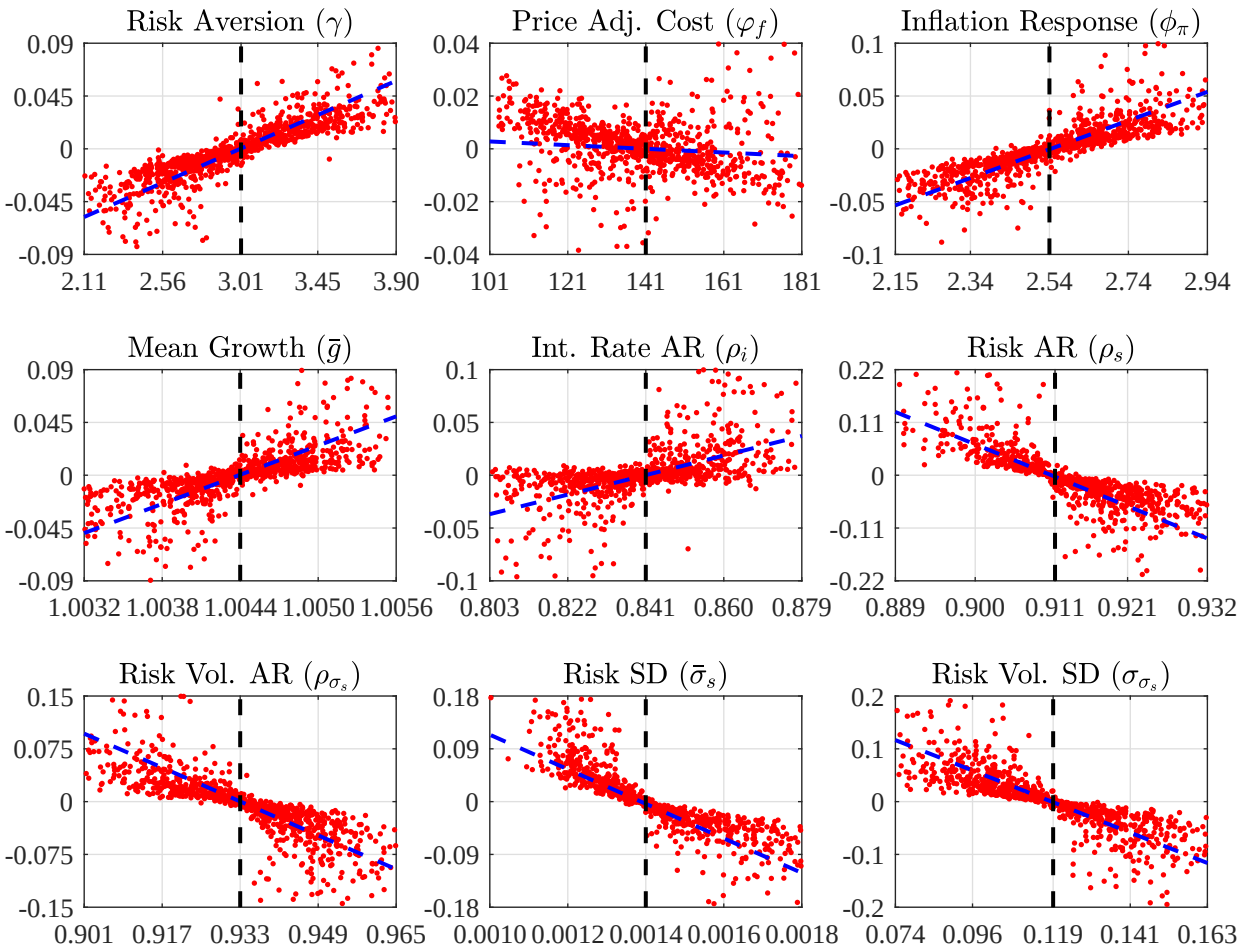

Figure 7: Scatter plots of the effect of output uncertainty on current output over a 24-quarter horizon in 2008Q4 at each posterior draw, conditional on fixing a parameter at its posterior mean. The values on the vertical axes are deviations from the effect with the posterior mean parameters. The horizontal axes denote the parameter values corresponding to the number of standard deviations away from the posterior mean. The dashed vertical lines are the posterior means.

The results depend on how each parameter affects expected volatility. The diagonal line is the linear trend. When the parameters governing the risk premium (ρ_s, ρ_{σ_s}, $\bar{\sigma}_s$, and σ_{σ_s}) are above their posterior means (vertical line) output uncertainty has a larger adverse effect because the variance of the exogenous process and hence expected volatility increase. For example, the posterior mean persistence of the risk premium, ρ_s, is 0.911. When that value is two standard deviations higher (0.932), output uncertainty reduces output by 0.13 percentage points more than at its posterior mean. A higher price adjustment cost parameter also causes output uncertainty to have a larger effect because stickier prices make households more sensitive to changes in the nominal interest rate.

Larger values of the other parameters reduce the effect of output uncertainty on output. An increase in the coefficient of relative risk aversion makes households less willing to substitute across time, which makes output less volatile. Thus, output uncertainty has a smaller adverse effect, even though households react more strongly to expected volatility. A higher monetary response to infla-

tion has a similar effect because it also reduces expected future volatility. Interest rate smoothing is a form of commitment by the Fed to reduce future inflation volatility, so the higher persistence reduces expected volatility. Finally, a higher average growth rate raises the steady-state nominal interest rate, which decreases the likelihood of ZLB events and therefore expected future volatility.

These results are particularly useful given the degree of parameter uncertainty in the literature. By extrapolating from the trend line, it is easy to obtain a rough estimate for the effects of output uncertainty and the likelihood of that outcome given any parameterization of the model. It is also possible to conduct a similar exercise for the other moments in the Euler equation decomposition.

5.6 WELFARE The cost of business cycles literature provides an alternative way to quantify the effects of uncertainty than our Euler equation decomposition. That literature uses welfare analysis to determine the consequences of different levels of volatility. The main difference between the two methods is that our Euler equation decomposition quantifies the effects of different higher order moments—including uncertainty—within a *particular* model, whereas the welfare analysis quantifies the effects of volatility by comparing *different* models. Specifically, the cost of business cycles literature measures the compensating variation of switching from a low to a high volatility model.

Given the household's constant relative risk aversion utility function in our baseline model, the compensating variation between models L (lower volatility) and H (higher volatility) is given by

$$\lambda_t = 1 - \left[\frac{E_t W_c(\tilde{c}^H) + 1/((1-\gamma)(1-\beta)) - E_t W_n(n^H) + E_t W_n(n^L)}{E_t W_c(\tilde{c}^L) + 1/((1-\gamma)(1-\beta))}\right]^{1/(1-\gamma)}, \tag{17}$$

where

$$E_t W_c(\tilde{c}^\vartheta) = E[\textstyle\sum_{j=t}^{\infty} \beta^{j-t}[((\tilde{c}_j^\vartheta)^{1-\gamma} - 1)/(1-\gamma)]|\hat{\Omega}_t],$$
$$E_t W_n(n^\vartheta) = E[\textstyle\sum_{j=t}^{\infty} \beta^{j-t}[\chi(n_j^\vartheta)^{1+\eta}/(1+\eta)]|\hat{\Omega}_t],$$

are the expected present-value of the household's utility from consumption and disutility from labor conditional on its information set at time t, $\hat{\Omega}_t$, which contains the median filtered state and the posterior mean parameters. Also, $\tilde{c}^\vartheta$ and n^ϑ are the optimal choices of detrended consumption and labor conditional on model $\vartheta \in \{H, L\}$. We denote the higher (lower) volatility economy with an H (L), where the expected path of consumption is lower (higher) due to precautionary saving.

We approximate W_c and W_n by integrating across $1{,}000$ simulations of $10{,}000$ quarters. Each simulation is conditional on the state of the economy in a particular period and the posterior mean parameters. λ_t is the fraction of consumption goods in the low volatility economy that would compensate the household for the lower consumption path in the higher volatility economy. When $\lambda_t > 0$ the household is better off in the lower volatility economy. Appendix F shows how to derive λ_t.

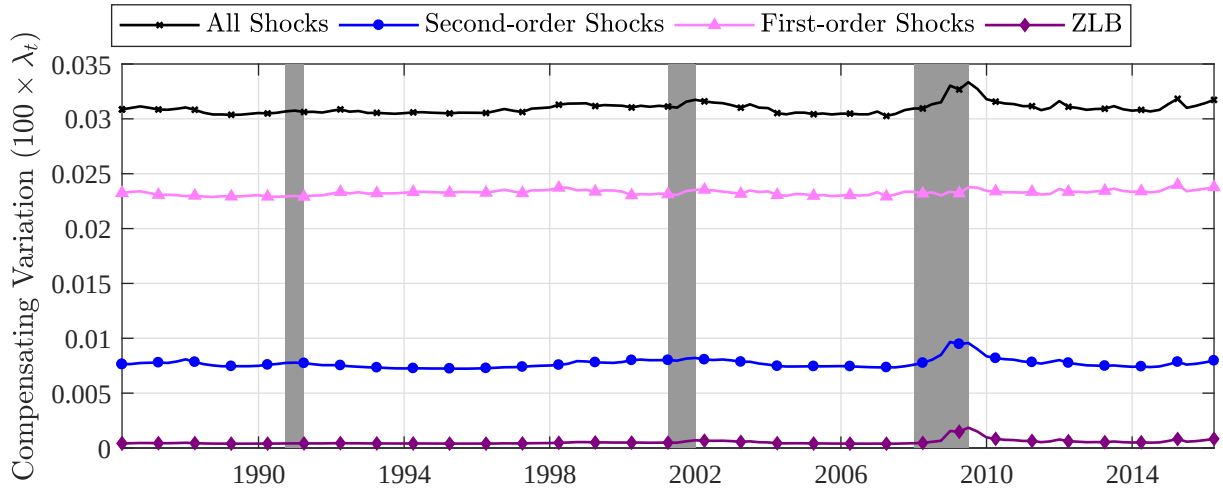

Figure 8: Percent of consumption goods under lower volatility needed to compensate the household for higher volatility. In each period, the welfare cost is conditional on the median filtered state from the posterior mean parameterization.

Figure 8 shows four estimated compensating variations: the effect of all shocks (models (1) and (4), x markers); the effect of only the stochastic volatility shocks (models (1) and (3), circles markers); the effect of only the first-moment shocks (models (3) and (4), triangle markers); and the effect of only the ZLB constraint (models (1) and (2), diamond markers). The compensating variation is shown as the percent of consumption goods in the lower volatility or no ZLB models.

In the baseline model, the household requires compensation of about 0.03% in every period to be indifferent to a world in which there is no volatility (i.e., the constant path of consumption and labor in the deterministic steady state), similar to the value in Lucas (1987) with $\gamma = 2$. There is a small increase in the welfare cost during recessions. Across the sample, about 75% of the compensation stems from the volatility induced by the first-moment shocks to productivity growth, the risk premium, and the interest rate. The remainder is due to the second-moment shocks and the endogenous amplification of both first- and second-moment shocks by the ZLB. Compensation for the uncertainty coming from second-moment shocks to productivity growth and the risk premium is higher than the compensation required for the endogenous uncertainty induced by the ZLB. Also, the higher welfare cost at the end of the Great Recession comes mostly from the interaction of second-moment shocks with the ZLB rather than first-moment shocks interacting with the ZLB.

6 THE EFFECT OF CAPITAL ACCUMULATION

In our baseline model without capital, output is equal to consumption and the only way households can save is by investing in a 1-period nominal bond, which is in zero net supply. This section extends the model so households can also invest in capital. In the data, investment is more volatile than real GDP, especially during recessions, so it is important to add capital to the model since it allows output, consumption, and investment to have different, potentially time-varying, volatilities.

The final goods firm's problem is unchanged. Intermediate firm i produces goods according to $y_t(i) = k_{t-1}(i)^{\alpha}(z_t n_t(i))^{1-\alpha}$. It chooses its capital and labor inputs, $n_t(i)$ and $k_{t-1}(i)$, and its price, $p_t(i)$, to maximize its profit function. In symmetric equilibrium, the optimality conditions imply

$$\tilde{y}_t^f = (\tilde{k}_{t-1}/g_t)^{\alpha} n_t^{1-\alpha}, \tag{18}$$

$$\alpha \tilde{w}_t n_t = (1-\alpha) r_t^k (\tilde{k}_{t-1}/g_t), \tag{19}$$

$$mc_t = \tilde{w}_t^{1-\alpha} (r_t^k)^{\alpha} / ((1-\alpha)^{1-\alpha} \alpha^{\alpha}), \tag{20}$$

and the Phillips curve, (5), which is identical except for the change in the marginal cost definition.

The household chooses $\{c_t, n_t, b_t, x_t, k_t\}_{t=0}^{\infty}$ to maximize the same utility function subject to

$$c_t + x_t + b_t/(i_t s_t) = w_t n_t + r_t^k k_{t-1} + b_{t-1}/\pi_t + d_t,$$
$$k_t = (1-\delta) k_{t-1} + x_t (1 - \varphi_x (x_t^g - 1)^2/2),$$

where x is investment in physical capital, $x_t^g \equiv x_t/(\bar{g} x_{t-1})$ is the growth rate of investment relative to the balanced growth rate, $\varphi_x > 0$ scales the size of the cost to adjusting investment, and k is the capital stock, which earns a real return r^k and depreciates at rate δ. In addition to the first-order conditions in the model without capital, (8) and (9), there are two new optimality conditions given by

$$q_t = \beta E_t [(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma} (r_{t+1}^k + q_{t+1}(1-\delta))/g_{t+1}], \tag{21}$$

$$1 = q_t [1 - \varphi_x (\tilde{x}_t^g - 1)^2 - \varphi_x \tilde{x}_t^g (\tilde{x}_t^g - 1)] + \beta \varphi_x \bar{g} E_t [q_{t+1} (\tilde{c}_t/\tilde{c}_{t+1})^{\gamma} (\tilde{x}_{t+1}^g)^2 (\tilde{x}_{t+1}^g - 1)/g_{t+1}]. \tag{22}$$

The detrended law of motion for capital and the aggregate resource constraint are given by

$$\tilde{k}_t = (1-\delta)(\tilde{k}_{t-1}/g_t) + \tilde{x}_t (1 - \varphi_x (\tilde{x}_t^g - 1)^2/2), \tag{23}$$

$$\tilde{c}_t + \tilde{x}_t = \tilde{y}_t. \tag{24}$$

Once again, we redefined variables that grow along the balanced growth path in terms of technology. A competitive equilibrium includes infinite sequences of quantities, $\{\tilde{c}_t, \tilde{y}_t, \tilde{y}_t^f, n_t, \tilde{x}_t, \tilde{k}_t\}_{t=0}^{\infty}$, prices, $\{\tilde{w}_t, i_t, i_t^n, \pi_t^{gap}, mc_t, q_t, r_t^k\}_{t=0}^{\infty}$, and exogenous variables, $\{s_t, g_t, \sigma_{g,t}, \sigma_{s,t}\}_{t=0}^{\infty}$, that satisfy the detrended equilibrium system, (1), (2), (5)-(11), (13), and (18)-(24), given the initial conditions, $\{c_{-1}, i_{-1}^n, x_{-1}, k_{-1}, g_0, s_0, \varepsilon_{i,0}, \sigma_{g,0}, \sigma_{s,0}\}$, and sequences of shocks, $\{\varepsilon_{g,t}, \varepsilon_{s,t}, \varepsilon_{i,t}, \varepsilon_{\sigma_g,t}, \varepsilon_{\sigma_s,t}\}_{t=1}^{\infty}$.

The model is numerically too costly to estimate, so we calibrate the three new parameters. The capital depreciation rate, δ, is calibrated to 0.025. The cost share of capital, α, and the investment adjustment cost parameter, φ_x, are set to 0.19 and 4.06, respectively, which equal the mean posterior estimates in Gust et al. (2017). Although there are some differences between our model and the one in Gust et al. (2017) (e.g., their model includes sticky wages and variable capital uti-

lization, whereas our model has stochastic volatility), we believe these parameter values provide a good approximation of what we would obtain if we estimated the model with Bayesian methods.

Fortunately, introducing capital does not change the consumption Euler equation we used to construct the decomposition in the model without capital. We generate policy functions for each term in the decomposition in the same way as the model without capital, except we filter the data with per capita real fixed investment growth in addition to the five observables we previously used.

Figure 9 shows the influence of the different types of uncertainty. The left panel plots the effects of consumption uncertainty over a 24-quarter horizon in the models with and without capital. In the capital model, consumption uncertainty on average decreases current consumption by 0.08%, which is only 0.02 percentage points more than in our baseline model without capital. The difference is more pronounced when the ZLB first binds. In 2008Q4 consumption uncertainty lowered consumption by 0.22% compared with only 0.14% in the baseline model, but that discrepancy quickly dissipated. Furthermore, the median effect in the capital model is typically in the left tail of the credible set implied by the baseline model. The right panel shows the impact of technology growth and inflation uncertainty in the model with capital—the two other types of uncertainty in the bond Euler equation. Both terms have nearly identical effects to those in the baseline model.

A major benefit of the capital model is that it provides a new Euler equation, (21), that we can use to quantify the effects of the uncertainty about the real rental rate of capital and Tobin's q on current consumption. Using the method in section 4.3, a third-order Taylor approximation implies

$$\begin{aligned}
\gamma\hat{c}_t \approx\ & \gamma E_t\hat{c}_{t+1} - ((\beta/\bar{g})\bar{r}^k E_t\hat{r}^k_{t+1} + (\beta/\bar{g})(1-\delta)E_t\hat{q}_{t+1} - \hat{q}_t - E_t\hat{g}_{t+1}) \\
& - \tfrac{1}{2}(\gamma^2 \operatorname{var}_t \hat{c}_{t+1} + \operatorname{var}_t \hat{g}_{t+1} + (\beta/\bar{g})\bar{r}^k \operatorname{var}_t \hat{r}^k_{t+1} + (\beta/\bar{g})(1-\delta)\operatorname{var}_t \hat{q}_{t+1}) \\
& - \gamma \operatorname{cov}_t(\hat{c}_{t+1}, \hat{g}_{t+1}) + \gamma(\beta/\bar{g})\bar{r}^k \operatorname{cov}_t(\hat{c}_{t+1}, \hat{r}^k_{t+1}) + \gamma(\beta/\bar{g})(1-\delta)\operatorname{cov}_t(\hat{c}_{t+1}, \hat{q}_{t+1}) \\
& + (\beta/\bar{g})\bar{r}^k \operatorname{cov}_t(\hat{g}_{t+1}, \hat{r}^k_{t+1}) + \beta((1-\delta)/\bar{g})\operatorname{cov}_t(\hat{g}_{t+1}, \hat{q}_{t+1}) \\
& + \tfrac{1}{6}(\gamma^3 \operatorname{skew}_t \hat{c}_{t+1} + \operatorname{skew}_t \hat{g}_{t+1} - (\beta/\bar{g})\bar{r}^k \operatorname{skew}_t \hat{r}^k_{t+1} - \beta((1-\delta)/\bar{g})\operatorname{skew}_t \hat{q}_{t+1}),
\end{aligned} \tag{25}$$

which we can once again iterate forward to eliminate the influence of expected future consumption. Several terms enter the same way as our previous decomposition. For example, the *ex-ante* variance of consumption and technology growth appear in (15) and (25), so they have the exact same effect on consumption. The rental rate and Tobin's q variance terms replace the inflation variance term.

The right panel also plots the new uncertainty terms over a 24-quarter horizon. Rental rate uncertainty has a similarly small effect as inflation uncertainty. Unlike the other higher-order terms, uncertainty about Tobin's q has almost half as large of an effect on consumption as consumption uncertainty, which shows the importance of capital adjustment costs for the transmission of uncertainty. Overall, uncertainty about the return on capital (rental rate and Tobin's q) has a larger influence on consumption than uncertainty about the real return on a risk-free nominal bond (inflation).

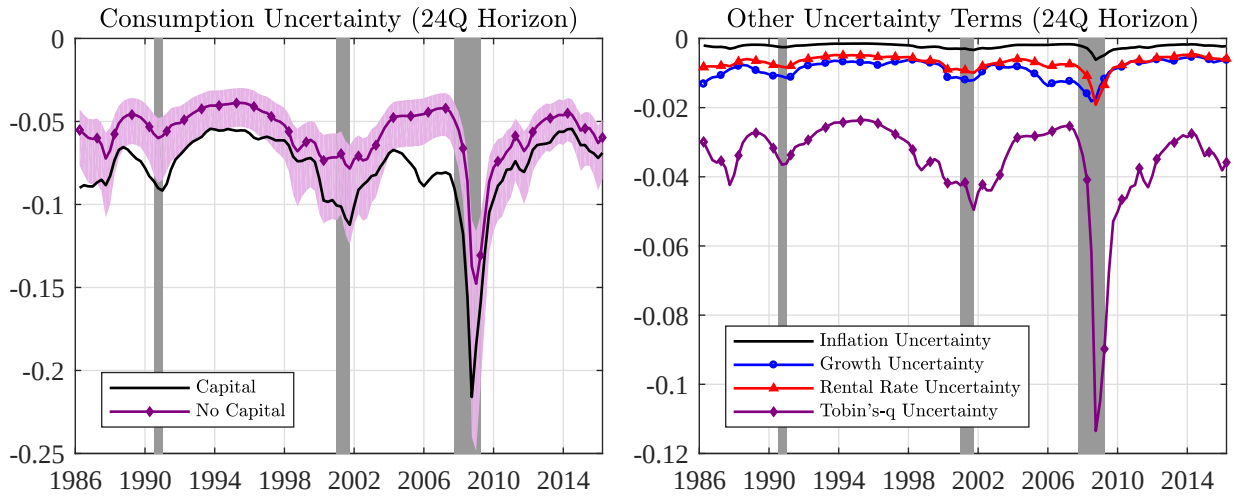

Figure 9: Filtered decomposition of the effects on current consumption. The shaded regions denote NBER recessions. The vertical axes are the contribution to the percentage point deviation of detrended consumption from its steady state.

7 COMPARISON WITH THE TRADITIONAL VAR APPROACH

The literature often adds a measure of uncertainty to the variables in a structural VAR and computes impulse responses using a recursive identification scheme.[4] The responses depend on where uncertainty is ordered in the list of variables. If uncertainty is ordered first, then subsequent variables in the VAR, which reflect information about the state of the economy, have no contemporaneous effect on the responses to an uncertainty shock. If it is ordered last, then none of the preceding variables in the VAR contemporaneously depend on uncertainty, so an uncertainty shock has no effect on impact. Therefore, the modeler must specify whether the uncertainty series is exogenous or endogenous. The challenges are even greater when accounting for multiple sources of uncertainty.

Due to the nonlinearities introduced by stochastic volatility and the ZLB constraint, we are interested in whether a *linear* VAR, commonly employed in the literature on uncertainty, can recover the dynamic relationship between uncertainty and real activity predicted by our baseline *nonlinear* model. We focus on the impulse response of output growth to a change in financial uncertainty, $U_{s,t} = \sqrt{E_t[\sigma^2_{v,t+1}]}$. Since financial uncertainty is exogenous in our structural model, it is easy for us to compare its effects to those in a VAR model and assess accuracy. The shocks in the VAR are identified recursively and the variables—financial uncertainty, output growth, inflation, wage growth, the risk premium, and the interest rate—are ordered from first to last in the same way as Christiano et al. (2005).[5] Appendix E provides additional information about our structural VAR.

[4] Alexopoulos and Cohen (2009) develop a proxy based on the number of *New York Times* articles on uncertainty. Bachmann et al. (2013) use forecaster disagreement from the Business Outlook Survey. Leduc and Liu (2016) create a measure based on respondents from Michigan Survey of Consumers who report uncertainty as a reason why it is a bad time to purchase vehicles. Basu and Bundick (2017), Bekaert et al. (2013), and Bloom (2009) use the VIX. The effect of an increase in these proxies varies widely, ranging from close to 0% to over 1%, depending on the shock size.

[5] We obtain very similar results using bivariate VARs with uncertainty ordered first and output growth second.

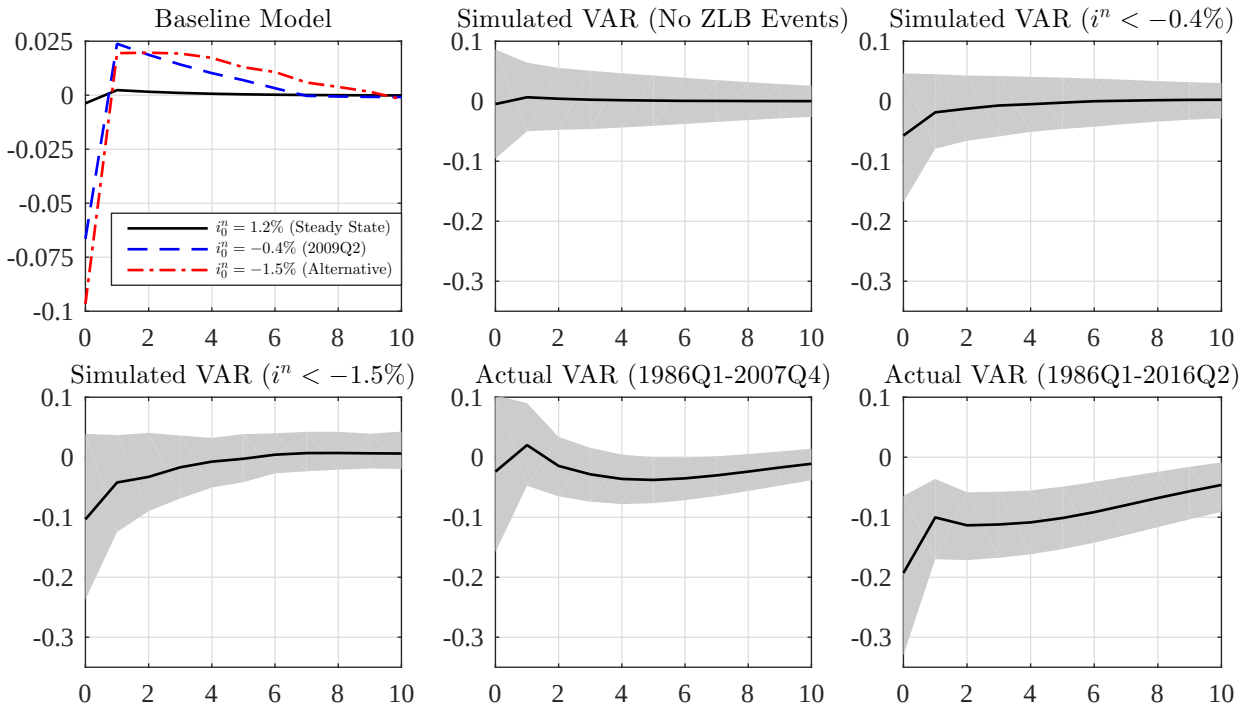

Figure 10: Impulse responses of output growth to a 2 standard deviation increase in financial uncertainty. The solid lines are the median responses and the shaded regions denote the 1 standard deviation (16%-84%) credible sets.

Figure 10 shows the responses to a 2 standard deviation financial uncertainty shock. The first subplot shows the predictions of our baseline model given different initializations of the state. When the response is initialized at the stochastic steady state (solid line), where the notional interest rate, i^n, is 1.2%, the effect of financial uncertainty on output growth is negligible across the whole horizon. However, when the response is initialized at the median filtered state corresponding to 2009Q2 (dashed line), where $i^n = -0.4\%$ initially, output growth declines by 0.07%. We alternatively initialize the response at an average state vector across simulated quarters at the ZLB such that $i^n = -1.5\%$ initially (dashed-dotted line). In that case, the financial uncertainty shock leads to a 0.1% decrease in output growth on impact. In summary, our baseline model predicts the impact effect of financial uncertainty on output growth depends on the initial state of the economy.

The simulated VARs in the next three subplots are estimated with data from short-sample simulations of the baseline model conditional on the posterior mean parameterization. The solid lines represent the median response and the shaded regions represent the 16%-84% credible sets. The first simulated VAR is estimated using artificial data without any ZLB events (i.e., $i^n > 0$ always). The response of output growth to a financial uncertainty shock is close to zero across the whole horizon, which is very similar to the prediction of our baseline model initialized at steady state.

The next two simulated VARs are estimated with artificial data where the notional rate falls below -0.4% or -1.5% for at least one quarter, so the responses represent averages across quarters

when the ZLB does and does not bind. Given these initial states, the median impact effects of financial uncertainty on output growth are -0.06% and -0.10%, respectively. Although the responses are not significantly different from zero, the median impact of financial uncertainty shocks identified by the VAR decreases as the quantity and severity of ZLB events increase in the simulated data and it is quantitatively similar to our structural model. Therefore, the *linear* VAR does a good job capturing the endogenous amplification of exogenous financial uncertainty shocks at the ZLB.

Finally, we estimate the same VAR with U.S. data. The second to last subplot excludes the Great Recession and subsequent ZLB period (1986Q1-2007Q4), while the last subplot is based on the sample used to estimate our baseline model (1986Q1-2016Q2). The results are similar to the predictions of our structural model. The effect of uncertainty is small and statistically insignificant in the pre-ZLB period, whereas it is much larger and significant when the ZLB period is included.

Interestingly, the structural VAR estimates based on U.S. data are also very similar to the effects of output uncertainty from our Euler equation decomposition.[6] Given the stark differences in methodology, the similarity between the two sets of results is noteworthy. They provide strong evidence that aggregate uncertainty has limited effects on the economy, even when the the ZLB binds.

A major advantage of our Euler equation decomposition is that it does not require us to take a stand on whether a given type of uncertainty is endogenous or exogenous. It can also account for multiple forms of uncertainty and how they nonlinearly interact with the economy. In other words, our decomposition is able to quantify the overall effect of all types of uncertainty in each period by accounting for the first and second moment shocks that best explain macro and uncertainty data. It also has the added advantage of being able to quantify the effects of other higher-order moments.

8 CONCLUSION

The literature often uses exogenous volatility shocks to examine the effects of aggregate uncertainty. We develop a new way to quantify the effects of uncertainty that accounts for both exogenous and endogenous uncertainty sources. We first estimate a nonlinear New Keynesian model with a ZLB constraint and stochastic volatility, while linking to both macro and uncertainty data. We then decompose output into its expected mean, variance, and skewness with the consumption Euler equation. Our decomposition reveals that uncertainty had a relatively small impact. Despite the nonlinearities in the model, output uncertainty never reduced output by more than 0.22%, even during the Great Recession, and other higher-order moments had much smaller effects on output.

A major benefit of our method is its flexibility to examine the effects of uncertainty in a broad class of models. While some models are too costly to estimate, it is possible to calculate the expected volatility or skewness surrounding any endogenous variable in the model. After calibrating

[6] We obtain very similar results if we use the macro uncertainty index instead of the financial uncertainty index.

the model, one could then filter the data while linking to an empirical measure of any higher-order moment. Using our results as a benchmark, future research could examine other sources of endogenous uncertainty, such as borrowing constraints, firm default, limited information, irreversible investment, search frictions, or heterogeneity to examine whether uncertainty has significantly larger effects. Alternative models may endogenously explain a larger fraction of the changes in uncertainty and the effects of uncertainty may become larger in some periods. However, our results suggest the peak effects of uncertainty may not significantly increase in other models, given that uncertainty had limited effects in our model despite the significant nonlinearity that occurs at the ZLB.

REFERENCES

ALEXOPOULOS, M. AND J. COHEN (2009): "Uncertain Times, Uncertain Measures," University of Toronto Working Paper 352.

AN, S. AND F. SCHORFHEIDE (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113–172.

ARELLANO, C., Y. BAI, AND P. J. KEHOE (2016): "Financial Frictions and Fluctuations in Volatility," NBER Working Paper 22990.

BACHMANN, R., S. ELSTNER, AND E. SIMS (2013): "Uncertainty and Economic Activity: Evidence from Business Survey Data," *American Economic Journal: Macroeconomics*, 5, 217–49.

BASU, S. AND B. BUNDICK (2015): "Endogenous Volatility at the Zero Lower Bound: Implications for Stabilization Policy," FRB Kansas City Working Paper 15-1.

——— (2017): "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 85, 937–958.

BEKAERT, G., M. HOEROVA, AND M. LO DUCA (2013): "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics*, 60, 771–788.

BERNDT, A., H. LUSTIG, AND ŞEVIN YELTEKIN (2012): "How Does the US Government Finance Fiscal Shocks?" *American Economic Journal: Macroeconomics*, 4, 69–104.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

BORN, B. AND J. PFEIFER (2014): "Policy Risk and the Business Cycle," *Journal of Monetary Economics*, 68, 68–85.

BRUNNERMEIER, M. K. AND Y. SANNIKOV (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379–421.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

CHUNG, H. AND E. M. LEEPER (2007): "What Has Financed Government Debt?" NBER Working Paper 13425.

COLEMAN, II, W. J. (1991): "Equilibrium in a Production Economy with an Income Tax," *Econometrica*, 59, 1091–1104.

DE GROOT, O., A. RICHTER, AND N. THROCKMORTON (2018): "Uncertainty Shocks in a Model of Effective Demand: Comment," *Econometrica*, forthcoming.

FAJGELBAUM, P., M. TASCHEREAU-DUMOUCHEL, AND E. SCHAAL (2017): "Uncertainty Traps," *The Quarterly Journal of Economics*, 132, 1641–1692.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, K. KUESTER, AND J. F. RUBIO-RAMÍREZ (2015): "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, 105, 3352–84.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–61.

FERNÁNDEZ-VILLAVERDE, J. AND J. F. RUBIO-RAMÍREZ (2007): "Estimating Macroeconomic Models: A Likelihood Approach," *Review of Economic Studies*, 74, 1059–1087.

GORDON, N. J., D. J. SALMOND, AND A. F. M. SMITH (1993): "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation," *IEE Proceedings F - Radar and Signal Processing*, 140, 107–113.

GOURIO, F. (2014): "Financial Distress and Endogenous Uncertainty," Manuscript, Federal Reserve Bank of Chicago.

GUERRÓN-QUINTANA, P. A. AND J. M. NASON (2013): "Bayesian Estimation of DSGE Models," in *Handbook of Research Methods and Applications in Empirical Macroeconomics*, Edward Elgar Publishing, chap. 21, 486–512.

GUST, C., E. HERBST, D. LÓPEZ-SALIDO, AND M. E. SMITH (2017): "The Empirical Implications of the Interest-Rate Lower Bound," *American Economic Review*, 107, 1971–2006.

HALL, G. J. AND T. J. SARGENT (2011): "Interest Rate Risk and Other Determinants of Post-WWII US Government Debt/GDP Dynamics," *American Economic Journal: Macroeconomics*, 3, 192–214.

HE, Z. AND A. KRISHNAMURTHY (2014): "A Macroeconomic Framework for Quantifying Systemic Risk," NBER Working Paper 19885.

HERBST, E. P. AND F. SCHORFHEIDE (2016): *Bayesian Estimation of DSGE Models*, Princeton, NJ: Princeton University Press.

ILUT, C. L. AND H. SAIJO (2016): "Learning, Confidence, and Business Cycles," NBER Working Paper 22958.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

KITAGAWA, G. (1996): "Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models," *Journal of Computational and Graphical Statistics*, 5, pp. 1–25.

KOPECKY, K. AND R. SUEN (2010): "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 13, 701–714.

LEDUC, S. AND Z. LIU (2016): "Uncertainty Shocks are Aggregate Demand Shocks," *Journal of Monetary Economics*, 82, 20–35.

LESTER, R., M. PRIES, AND E. SIMS (2014): "Volatility and Welfare," *Journal of Economic Dynamics and Control*, 38, 17–36.

LUCAS, JR, R. E. (1987): *Models of Business Cycles*, Oxford: Basil Blackwell.

LUDVIGSON, S. C., S. MA, AND S. NG (2017): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" NBER Working Paper 21803.

MASON, J. W. AND A. JAYADEV (2014): ""Fisher Dynamics" in US Household Debt, 1929-2011," *American Economic Journal: Macroeconomics*, 6, 214–234.

MENDOZA, E. G. (2010): "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100, 1941–1966.

MUMTAZ, H. AND F. ZANETTI (2013): "The Impact of the Volatility of Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 45, 535–558.

NAKATA, T. (2017): "Uncertainty at the Zero Lower Bound," *American Economic Journal: Macroeconomics*, 9, 186–221.

NAVARRO, G. (2014): "Financial Crises and Endogenous Volatility," Manuscript, New York University.

OTROK, C. (2001): "On Measuring the Welfare Cost of Business Cycles," *Journal of Monetary Economics*, 47, 61–92.

PARKER, J. A. AND B. PRESTON (2005): "Precautionary Saving and Consumption Fluctuations," *American Economic Review*, 95, 1119–1143.

PETERMAN, W. B. (2016): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," *Economic Inquiry*, 54, 100–120.

PLANTE, M., A. W. RICHTER, AND N. A. THROCKMORTON (2018): "The Zero Lower Bound and Endogenous Uncertainty," *Economic Journal*, forthcoming.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

ROUWENHORST, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton, NJ: Princeton University Press, 294–330.

SAIJO, H. (2017): "The Uncertainty Multiplier and Business Cycles," *Journal of Economic Dynamics and Control*, 78, 1–25.

SIMS, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

SMETS, F. AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

STEWART, L. AND P. MCCARTY, JR (1992): "Use of Bayesian Belief Networks to Fuse Continuous and Discrete Information for Target Recognition, Tracking, and Situation Assessment," *Proc. SPIE*, 1699, 177–185.

STRAUB, L. AND R. ULBRICHT (2015): "Endogenous Uncertainty and Credit Crunches," Toulouse School of Economics Working Paper 15-604.

TALLARINI, JR., T. D. (2000): "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45, 507–532.

TAYLOR, J. B. (1993): "Discretion Versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

VAN NIEUWERBURGH, S. AND L. VELDKAMP (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53, 753–772.

XU, S. (2017): "Volatility Risk and Economic Welfare," *Journal of Economic Dynamics and Control*, 80, 17–33.

A DATA SOURCES

We drew from the following data sources to estimate our New Keynesian and VAR models:

1. **Financial Uncertainty index**: Monthly. Source: Ludvigson et al. (2017), $h = 3$ (1-quarter forecast horizon). Data available from `http://www.sydneyludvigson.com/`.
2. **Macro Uncertainty Index**: Monthly. Source: Jurado et al. (2015), $h = 3$ (1-quarter forecast horizon). Data available from `http://www.sydneyludvigson.com/`.
3. **Real GDP**: Quarterly, chained 2009 dollars, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 1.1.6 (FRED ID: GDPC1).
4. **GDP Deflator**: Quarterly, seasonally adjusted, index 2009=100. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 1.1.9 (FRED ID: GDPDEF).
5. **Average Hourly Earnings**: Monthly, production and nonsupervisory employees, dollars per hour, seasonally adjusted. Source: Bureau of Labor Statistics (FRED ID: AHETPI).
6. **Interest Rate Spread (Risk Premium)**: Monthly, Moody's seasoned Baa corporate bond yield relative to the yield on 10-Year treasury bond. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates, H.15 (FRED ID: BAA10YM)
7. **Effective Federal Funds Rate:** Daily. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates, H.15 (FRED ID: FEDFUNDS).
8. **Civilian Noninstitutional Population**: Monthly. Source: U.S. Bureau of Labour Statistics, Current Population Survey (FRED ID: CNP16OV).

9. **Fixed Investment**: Quarterly, billions of dollars, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 1.1.5 (FRED ID: FPI).

We applied the following transformations to the above series:

10. **Per Capita Real GDP**: $1{,}000{,}000\times$Real GDP$/$Population.
11. **Real Wage**: $100\times$Average Hourly Earnings$/$Price Index.
12. **Real Investment**: Average FPI in 2009$\times$(FPI Quantity Index$/100$). Quantity Index FRED ID: A007RA3Q086SBEA.
13. **Per Capita Real Investment**: $1{,}000{,}000\times$Real Investment$/$Population.

We converted the monthly or daily time series to a quarterly frequency by applying time averages over each quarter. In order, the variables used to estimate our VAR model are series 1, 10, 4, 11, 6, and 7. The observables used to estimate our nonlinear model without capital include series 10, 4, 7, 1, and 2. When we filter the data using the model with capital, we add series 13 as an observable.

B EULER EQUATION DECOMPOSITION

The bond Euler equation is given by

$$1 = \beta E_t[(\tilde{y}_t/\tilde{y}_{t+1})^{\gamma}(s_t i_t/(g_{t+1}\pi_{t+1}))] = E_t[\exp(\hat{\imath}_t + \hat{s}_t - \hat{\pi}_{t+1} - \hat{g}_{t+1} + \gamma(\hat{y}_t - \hat{y}_{t+1}))],$$

where a hat denotes log deviation from the balanced growth path. After reorganizing, we obtain

$$\begin{aligned}
-(\hat{\imath}_t + \hat{s}_t + \gamma\hat{y}_t) &= \log(\ E_t[\exp(-\hat{\pi}_{t+1} - \hat{g}_{t+1} - \gamma\hat{y}_{t+1})]) \\
&= \log(\ E_t[1 - (\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma\hat{y}_{t+1}) + \tfrac{1}{2}(\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma\hat{y}_{t+1})^2 - \\
&\qquad \tfrac{1}{6}(\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma\hat{y}_{t+1})^3 + \cdots]) \\
&= \log\Big(1 - (E_t[\hat{\pi}_{t+1}] + E_t[\hat{g}_{t+1}] + \gamma E_t[\hat{y}_{t+1}]) + \\
&\qquad \tfrac{1}{2}\Big(E_t[\hat{\pi}^2_{t+1}] + E_t[\hat{g}^2_{t+1}] + \gamma^2 E_t[\hat{y}^2_{t+1}] + \\
&\qquad\quad 2E_t[\hat{\pi}_{t+1}\hat{g}_{t+1}] + 2\gamma E_t[\hat{g}_{t+1}\hat{y}_{t+1}] + 2\gamma E_t[\hat{\pi}_{t+1}\hat{y}_{t+1}]\Big) - \\
&\qquad \tfrac{1}{6}\Big(E_t[\hat{\pi}^3_{t+1}] + E_t[\hat{g}^3_{t+1}] + \gamma^3 E_t[\hat{y}^3_{t+1}] + 6\gamma E_t[\hat{\pi}_{t+1}\hat{g}_{t+1}\hat{y}_{t+1}] + \\
&\qquad\quad 3E_t[\hat{\pi}^2_{t+1}\hat{g}_{t+1}] + 3\gamma E_t[\hat{\pi}^2_{t+1}\hat{y}_{t+1}] + 3E_t[\hat{g}^2_{t+1}\hat{\pi}_{t+1}] + \\
&\qquad\quad 3\gamma E_t[\hat{g}^2_{t+1}\hat{y}_{t+1}] + 3\gamma^2 E_t[\hat{y}^2_{t+1}\hat{\pi}_{t+1}] + 3\gamma^2 E_t[\hat{y}^2_{t+1}\hat{g}_{t+1}]\Big) + \cdots\Big),
\end{aligned}$$

where the second equality follows from the Maclaurin series for $e^x = 1 + x + x^2/2 + x^3/6 + \cdots$.

Subsequently applying a third-order Maclaurin series to $\log(1-x) \approx -x - x^2/2 - x^3/3$ implies

$$\begin{aligned}\hat{\imath}_t + \hat{s}_t + \gamma\hat{y}_t \approx\ & E_t[\hat{\pi}_{t+1}] + E_t[\hat{g}_{t+1}] + \gamma E_t[\hat{y}_{t+1}] \\ & - \tfrac{1}{2}((E_t[\hat{\pi}_{t+1}^2] - (E_t[\hat{\pi}_{t+1}])^2) + (E_t[\hat{g}_{t+1}^2] - (E_t[\hat{g}_{t+1}])^2) + \gamma^2(E_t[\hat{y}_{t+1}^2] - (E_t[\hat{y}_{t+1}])^2)) \\ & - \gamma(E_t[\hat{\pi}_{t+1}\hat{y}_{t+1}] - E_t[\hat{\pi}_{t+1}]E_t[\hat{y}_{t+1}]) - \gamma(E_t[\hat{g}_{t+1}\hat{y}_{t+1}] - E_t[\hat{g}_{t+1}]E_t[\hat{y}_{t+1}]) \\ & - (E_t[\hat{\pi}_{t+1}\hat{g}_{t+1}] - E_t[\hat{\pi}_{t+1}]E_t[\hat{g}_{t+1}]) \\ & + \tfrac{1}{6}(E_t[\hat{\pi}_{t+1}^3] - 3E_t[\hat{\pi}_{t+1}]E_t[\hat{\pi}_{t+1}^2] + 2(E_t[\hat{\pi}_{t+1}])^3) \\ & + \tfrac{1}{6}(E_t[\hat{g}_{t+1}^3] - 3E_t[\hat{g}_{t+1}]E_t[\hat{g}_{t+1}^2] + 2(E_t[\hat{g}_{t+1}])^3) \\ & + \tfrac{1}{6}\gamma^3(E_t[\hat{y}_{t+1}^3] - 3E_t[\hat{y}_{t+1}]E_t[\hat{y}_{t+1}^2] + 2(E_t[\hat{y}_{t+1}])^3),\end{aligned}$$

after dropping the higher-order terms. Therefore, current output is approximated by

$$\begin{aligned}\gamma\hat{y}_t \approx\ & \gamma E_t\hat{y}_{t+1} - \hat{r}_t - \tfrac{1}{2}(\operatorname{var}_t\hat{\pi}_{t+1} + \operatorname{var}_t\hat{g}_{t+1} + \gamma^2\operatorname{var}_t\hat{y}_{t+1}) \\ & - (\gamma\operatorname{cov}_t(\hat{\pi}_{t+1},\hat{y}_{t+1}) + \gamma\operatorname{cov}_t(\hat{\pi}_{t+1},\hat{g}_{t+1}) + \operatorname{cov}_t(\hat{\pi}_{t+1},\hat{y}_{t+1})) \\ & + \tfrac{1}{6}(\operatorname{skew}_t\hat{\pi}_{t+1} + \operatorname{skew}_t\hat{g}_{t+1} + \gamma^3\operatorname{skew}_t\hat{y}_{t+1}),\end{aligned}$$

where $\hat{r}_t = \hat{\imath}_t + \hat{s}_t - E_t\hat{\pi}_{t+1} - E_t\hat{g}_{t+1}$ is the *ex-ante* real rate, $\operatorname{var}_t(x_{t+1}) = E_t[\hat{x}_{t+1}^2] - (E_t[\hat{x}_{t+1}])^2$ is the variance of x, $\operatorname{skew}_t\hat{x}_{t+1} = E_t[\hat{x}_{t+1}^3] - 3E_t[\hat{x}_{t+1}]E_t[\hat{x}_{t+1}^2] + 2(E_t[\hat{x}_{t+1}])^3$ is the third moment of x, and $\operatorname{cov}_t(x_{t+1},y_{t+1}) = E_t[x_{t+1}y_{t+1}] - E_t[x_{t+1}]E_t[y_{t+1}]$ is the covariance between x and y.

The derivation of (25) follows very similar steps, although it contains significantly more terms.

C SOLUTION METHOD

C.1 BASELINE MODEL

We begin by compactly writing the detrended equilibrium system as

$$E[f(\mathbf{v}_{t+1},\mathbf{v}_t,\varepsilon_{t+1})|\mathbf{z}_t,\vartheta] = 0,$$

where f is a vector-valued function, $\mathbf{v} = (g, s, \sigma_g, \sigma_s, \tilde{c}, \tilde{y}, \tilde{y}^f, n, \tilde{w}, mc, i, i^n, \pi^{gap})$ is a vector of variables, $\varepsilon \equiv [\varepsilon_g, \varepsilon_s, \varepsilon_i, \varepsilon_{\sigma_g}, \varepsilon_{\sigma_s}]'$ is a vector of shocks, $\mathbf{z}_t = (\varepsilon_{i,t}, \log(\sigma_{g,t}), \log(\sigma_{s,t}), g_t, s_t, mp_{t-1})$, and ϑ are the parameters. Since i_{t-1}^n and $\tilde{y}_{t-1}$ only appear in the policy rule, we defined $mp_{t-1} = (i_{t-1}^n)^{\rho_i}(\tilde{y}_{t-1})^{\phi_y(\rho_i-1)}$ and rewrote the rule as $i_t^n = mp_{t-1}(\bar{\imath}(\pi_t^{gap})^{\phi_\pi}(g_t\tilde{y}_t/\bar{g})^{\phi_y})^{1-\rho_i}\exp(\sigma_i\varepsilon_{i,t})$.

There are many ways to discretize the exogenous state variables, $\varepsilon_{i,t}$, $\log(\sigma_{g,t})$, and $\log(\sigma_{s,t})$. We use the Markov chain in Rouwenhorst (1995), which Kopecky and Suen (2010) show outperforms other methods for approximating autoregressive processes. The bounds on g_t, s_t, and mp_{t-1} are set to $\pm 3\%$, $\pm 2\%$, and $\pm 2\%$ of the deterministic steady state to contain the filtered state variables. We discretize the state variables into $(4, 9, 7, 7, 7, 7)$ evenly-spaced points. There are

$D = 86{,}436$ nodes in the state space, and the realization of $\mathbf{z}_t$ on node d is denoted $\mathbf{z}_t(d)$. The Rouwenhorst method provides integration nodes for $[\varepsilon_{i,t+1}(m), \log(\sigma_{g,t+1}(m)), \log(\sigma_{s,t+1}(m))]$ that are the same as the respective state variables. However, the processes for g_{t+1} and s_{t+1} do not have a standard autoregressive form because of the stochastic volatility. Thus, the first moment shocks, $[\varepsilon_{g,t+1}, \varepsilon_{s,t+1}]$, are discretized separately from the volatility processes. The policy functions are then interpolated at realizations of $g_{t+1}(m)$ and $s_{t+1}(m)$ that can occur between nodes in the state space. We use the same number of interpolation nodes as the state variables, $(4, 9, 7, 7, 7)$, or $M = 12{,}348$, and the Rouwenhorst method provides weights, $\phi(m)$, for $m \in \{1, \ldots, M\}$.

For the policy functions, we approximate $\tilde{c}_t(\mathbf{z}_t)$ and $\pi_t^{gap}(\mathbf{z}_t)$. Our choice of policy functions, while not unique, simplifies solving for the other variables in the nonlinear system of equations given $\mathbf{z}_t$. The following steps outline our global policy function iteration algorithm:

1. Use Sims's (2002) `gensys` algorithm to solve the log-linear model without the ZLB imposed. Then map that solution onto the discretized state space to initialize $\tilde{\mathbf{c}}_0$ and $\boldsymbol{\pi}_0^{gap}$.
2. On iteration $j \in \{1, \ldots, N_{iter}\}$ and for each $d \in \{1, \ldots, D\}$, use Chris Sims's `csolve` to find $\tilde{c}_t$ and π_t^{gap} to satisfy $\mathbb{E}[f(\cdot)|\mathbf{z}_t(d), \vartheta] \approx 0$, where N_{iter} is the number of iterations. Guessing $\tilde{c}_t = \tilde{\mathbf{c}}_{j-1}(d)$ and $\pi_t^{gap} = \boldsymbol{\pi}_{j-1}^{gap}(d)$, approximate $\mathbb{E}[f(\cdot)|\mathbf{z}_t(d), \vartheta]$ as follows:
   (a) Solve for $\{\tilde{y}_t, \tilde{y}_t^f, i_t^n, i_t, \tilde{w}_t, mp_t\}$ given $\tilde{c}_t$, π_t^{gap}, and $\mathbf{z}_t(d)$.
   (b) Linearly interpolate the policy functions, $\tilde{\mathbf{c}}_{j-1}$ and $\boldsymbol{\pi}_{j-1}^{gap}$, at the updated state variables, $\mathbf{z}_{t+1}(m)$, to obtain $\tilde{c}_{t+1}(m)$ and $\pi_{t+1}^{gap}(m)$ on every integration node, $m \in \{1, \ldots, M\}$.
   (c) Given $\{\tilde{c}_{t+1}(m), \pi_{t+1}^{gap}(m)\}_{m=1}^M$, solve for the other elements of $\mathbf{v}_{t+1}(m)$ and compute:

$$\mathbb{E}[f(\mathbf{v}_{t+1}, \mathbf{v}_t(d), \varepsilon_{t+1})|\mathbf{v}_t(d), \vartheta] \approx \textstyle\sum_{m=1}^M \phi(m) f(\mathbf{v}_{t+1}(m), \mathbf{v}_t(d), \varepsilon_{t+1}(m)).$$

   When `csolve` converges, set $\tilde{\mathbf{c}}_j(d) = \tilde{c}_t$ and $\boldsymbol{\pi}_j^{gap}(d) = \pi_t^{gap}$.
3. Repeat step 2 until $\text{maxdist}_j < 10^{-6}$, where $\text{maxdist}_j \equiv \max\{|\tilde{\mathbf{c}}_j - \tilde{\mathbf{c}}_{j-1}|, |\boldsymbol{\pi}_j^{gap} - \boldsymbol{\pi}_{j-1}^{gap}|\}$. When that occurs, the algorithm has converged to an approximate nonlinear solution.

Figure 11 shows the distribution of the absolute value of the errors in base 10 logarithms for the consumption Euler equation and the Phillips curve. For example, an error of -3 means there is a mistake of 1 consumption good for every $1{,}000$ goods. The mean Euler equation error is -3.96 and the mean Phillips curve error is -2.32. By construction, the errors on nodes used in the solution algorithm are less than the convergence criterion, 10^{-6}. The larger average errors are due to linear interpolation of the policy functions with respect to the (g_t, s_t, mp_{t-1}) states. To measure the errors between the nodes, we created a new grid with a total of $D = 850{,}500$ nodes by increasing the number of points in the (g_t, s_t, mp_{t-1}) dimensions to $(15, 15, 15)$. We used the same number of points in the $(\varepsilon_{i,t}, \log(\sigma_{g,t}), \log(\sigma_{s,t}))$ dimensions since they are discretized

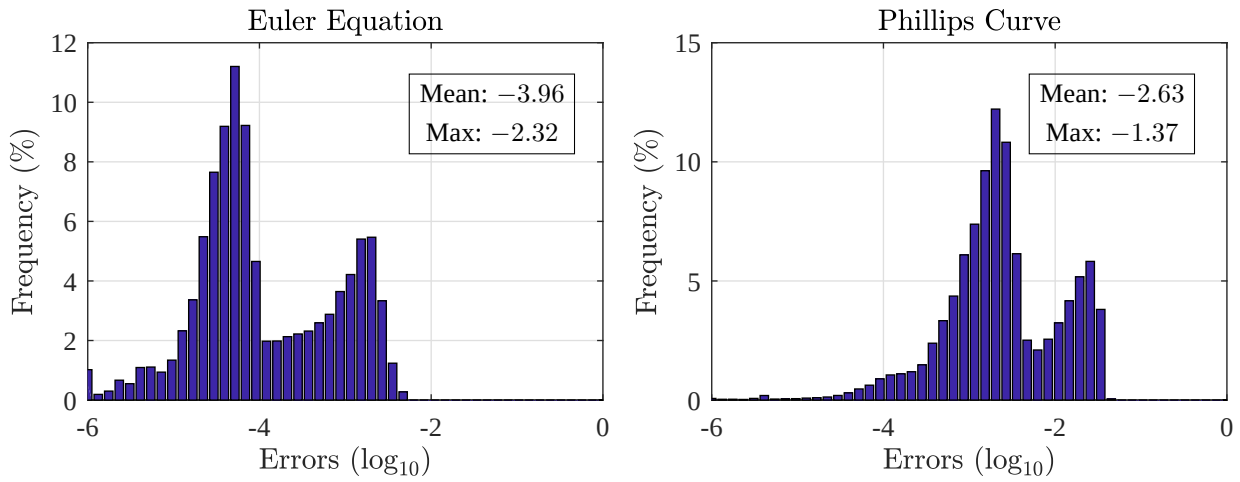

Figure 11: Distribution of Euler equation and Phillips curve errors in base 10 logarithms

with the Rouwenhorst method, which means the corresponding integration weights and nodes are state dependent. Therefore, the reported errors are consistent with the accuracy of the integral calculated when solving the model. Calculating the errors between the nodes corresponding to the exogenous state variables would require changing the numerical integration method (e.g., Gauss-Hermite quadrature). We decided not to show those errors because then the accuracy of the integral used to compute the errors would be inconsistent with the methods used to compute the solution.

C.2 CAPITAL MODEL

We solve the model with capital in the same way as the baseline model without capital. The state vector is the same as the baseline model, except it includes two additional endogenous state variables, x_{t-1} and k_{t-1}. The bounds on g_t, s_t, mp_{t-1}, x_{t-1} and k_{t-1} are set to $\pm 3\%$, $\pm 1.5\%$, $\pm 2\%$, $\pm 10\%$, and $\pm 7\%$ of steady state. We discretize the state variables into $(4, 7, 7, 7, 7, 7, 7, 11)$ points respectively, so there are $D = 5{,}176{,}556$ nodes in the state space. We use the most points on the capital dimension because it has the widest grid. Once again, we set the number of points on each shock equal to the number of points on the corresponding state variable.

D ESTIMATION ALGORITHM

We use a random walk Metropolis-Hastings algorithm to estimate our model with quarterly data from 1986Q1 to 2016Q2. To measure how well the model fits the data, we use the adapted particle filter described in Algorithm 12 in Herbst and Schorfheide (2016), which modifies the filter in Stewart and McCarty (1992) and Gordon et al. (1993) to better account for the outliers in the data.

D.1 METROPOLIS-HASTINGS ALGORITHM

The following steps outline the algorithm:

1. Specify the prior distributions, means, variances, and bounds of each element of the vector of N_e estimated parameters, $\theta \equiv \{\gamma, \varphi_f, \phi_\pi, \phi_y, \bar{g}, \bar{\pi}, \rho_i, \rho_g, \rho_s, \rho_{\sigma_g}, \rho_{\sigma_s}, \sigma_i, \bar{\sigma}_g, \bar{\sigma}_s, \sigma_{\sigma_g}, \sigma_{\sigma_s}\}$.

2. We match data on per capita real GDP, $RGDP/CNP$, the GDP deflator, DEF, the federal funds rate, FFR, the macro uncertainty series in Jurado et al. (2015), UM, and the financial uncertainty series in Ludvigson et al. (2017), UF. The vector of observables is given by

$$\hat{\mathbf{x}}_t^{data} \equiv [\Delta \log(RGDP_t/CNP_t),\ \Delta \log(DEF_t),\ \log(1 + FFR_t/100)/4,\ z(UM_t),\ z(UF_t)],$$

where Δ denotes a difference, $z(\cdot)$ is a standardized variable, and $t \in \{1, \ldots, T\}$. When we filter the data using the model with capital, we add per capita real investment, RI/CNP, to the vector of observables, so $\hat{\mathbf{x}}_t^{data}$ also includes $\Delta \log(RI_t/CNP_t)$.

3. Find the posterior mode to initialize the preliminary Metropolis-Hastings step.
   (a) For all $i \in \{1, \ldots, N_m\}$, where $N_m = 5{,}000$, apply the following steps:
      i. Draw $\hat{\theta}_i$ from the joint prior distribution and calculate its density value:

$$\log \ell_i^{prior} = \textstyle\sum_{j=1}^{N_e} \log p(\hat{\theta}_{i,j} | \mu_j, \sigma_j^2),$$

      where p is the prior density function of parameter j with mean μ_j and variance σ_j^2.
      ii. Given $\hat{\theta}_i$, solve the model according to Appendix C. If the algorithm converges, then compute the stochastic steady state, otherwise repeat step 3(a)i and redraw $\hat{\theta}_i$.
      iii. If the stochastic steady state exists, then use the particle filter in section D.2 to obtain the log-likelihood value for the model, $\log \ell_i^{model}$, otherwise repeat step 3(a)i.
      iv. The posterior log-likelihood is $\log \ell_i^{post} = \log \ell_i^{prior} + \log \ell_i^{model}$
   (b) Calculate $\max(\log \ell_1^{post}, \ldots, \log \ell_{N_m}^{post})$ and find the corresponding parameter vector, $\hat{\theta}_0$.
4. Approximate the covariance matrix for the joint posterior distribution of the parameters, Σ, which is used to draw candidates during the preliminary Metropolis-Hastings step.
   (a) Locate the draws with a likelihood in the top decile. Stack the $N_{m,sub} = (1-p)N_m$ draws in a $N_{m,sub} \times N_e$ matrix, $\hat{\Theta}$, and define $\tilde{\Theta} = \hat{\Theta} - \sum_{i=1}^{N_{m,sub}} \hat{\theta}_{i,j}/N_{m,sub}$.
   (b) Calculate $\Sigma = \tilde{\Theta}'\tilde{\Theta}/N_{m,sub}$ and verify it is positive definite, otherwise repeat step 3.
5. Perform an initial run of the random walk Metropolis-Hastings algorithm.
   (a) For all $i \in \{0, \ldots, N_d\}$, where $N_d = 25{,}000$, perform the following steps:
      i. Draw a candidate vector of parameters, $\hat{\theta}_i^{cand}$, where

$$\hat{\theta}_i^{cand} \sim \begin{cases} \mathbb{N}(\hat{\theta}_0, c_0\Sigma) & \text{for } i = 0, \\ \mathbb{N}(\hat{\theta}_{i-1}, c\Sigma) & \text{for } i > 0. \end{cases}$$

      We set $c_0 = 0$ and tune c to target an overall acceptance rate of roughly 30%.

ii. Calculate the prior density value, $\log \ell_i^{prior}$, of the candidate draw, $\hat{\theta}_i^{cand}$ as in 3(a)i.

iii. Given $\hat{\theta}_i^{cand}$, solve the model according to Appendix C. If the algorithm converges, compute the stochastic steady state, otherwise repeat 5(a)i and draw a new $\hat{\theta}_i^{cand}$.

iv. If the stochastic steady state exists, then use the particle filter in section D.2 to obtain the log-likelihood value for the model, $\log \ell_i^{model}$, otherwise repeat 5(a)i.

v. Accept or reject the candidate draw according to

$$(\hat{\theta}_i, \log \ell_i) = \begin{cases} (\hat{\theta}_i^{cand}, \log \ell_i^{cand}) & \text{if } i = 0, \\ (\hat{\theta}_i^{cand}, \log \ell_i^{cand}) & \text{if } \min(1, \ell_i^{cand}/\ell_{i-1}) > \hat{u}, \\ (\hat{\theta}_{i-1}, \log \ell_{i-1}) & \text{otherwise,} \end{cases}$$

where $\hat{u}$ is a draw from a uniform distribution, $\mathbb{U}[0,1]$, and the posterior log-likelihood associated with the candidate draw is $\log \ell_i^{cand} = \log \ell_i^{prior} + \log \ell_i^{model}$.

(b) Burn the first $N_b = 5000$ draws and use the remaining sample to calculate the mean draw, $\bar{\theta}^{preMH} = \sum_{i=N_b+1}^{N_{\text{preMH}}} \hat{\theta}_i$, and the covariance matrix, Σ^{preMH}. We follow step 4 to calculate Σ^{preMH} but use all $N_d - N_b$ draws instead of just the upper pth percentile.

6. Following the procedure in step 5, perform a final run of the Metropolis-Hastings algorithm, where $\hat{\theta}_0 = \bar{\theta}^{preMH}$ and $\Sigma = \Sigma^{preMH}$. We set $N_d = 100{,}000$ and keep every 100th draw. The remaining $1{,}000$ draws form a representative sample from the joint posterior density.

D.2 ADAPTED PARTICLE FILTER Henceforth, our definition of $\mathbf{v}_t$ from Appendix C is referred to as the state vector, which should not be confused with the state variables for the nonlinear model.

1. Initialize the filter by drawing $\{\varepsilon_{t,p}\}_{t=-24}^{0}$ for all $p \in \{0, \ldots, N_p\}$ and simulating the model, where N_p is the number of particles. We initialize the filter with the final state vector, $\mathbf{v}_{0,p}$, which is approximately a draw from the model's ergodic distribution. We set $N_p = 40{,}000$.

2. For $t \in \{1, \ldots, T\}$, sequentially filter the data with the linear or nonlinear model as follows:

(a) For $p \in \{1, \ldots, N_p\}$, draw shocks from an adapted distribution, $\varepsilon_{t,p} \sim \mathbb{N}(\bar{\varepsilon}_t, I)$, where $\bar{\varepsilon}_t$ maximizes $p(\xi_t|\mathbf{v}_t)p(\mathbf{v}_t|\bar{\mathbf{v}}_{t-1})$ and $\bar{\mathbf{v}}_{t-1} = \sum_{p=1}^{N_p} \mathbf{v}_{t-1,p}/N_p$ is the mean state vector.

i. Use the model solution to update the state vector, $\mathbf{v}_t$, given $\bar{\mathbf{v}}_{t-1}$ and a guess for $\bar{\varepsilon}_t$. Define $\mathbf{v}_t^h \equiv H\mathbf{v}_t$, where H selects the observable variables from the state vector.

ii. Calculate the measurement error (ME), $\xi_t = \mathbf{v}_t^h - \hat{\mathbf{x}}_t^{data}$, which is assumed to be multivariate normally distributed, $p(\xi_t|\mathbf{v}_t) = (2\pi)^{-3/2}|R|^{-1/2}\exp(-\xi_t' R^{-1}\xi_t/2)$, with covariance matrix, $R \equiv \text{diag}(\sigma_{me,y^g}, \sigma_{me,\pi}, \sigma_{me,i}, \sigma_{me,um}, \sigma_{me,uf})^2$.

iii. The probability of observing $\mathbf{v}_t$, given $\bar{\mathbf{v}}_{t-1}$, is $p(\mathbf{v}_t|\bar{\mathbf{v}}_{t-1}) = (2\pi)^{-3/2}\exp(-\bar{\varepsilon}_t'\bar{\varepsilon}_t/2)$.

iv. Maximize $p(\xi_t|\mathbf{v}_t)p(\mathbf{v}_t|\bar{\mathbf{v}}_{t-1}) \propto \exp(-\xi_t' R^{-1}\xi_t/2)\exp(-\bar{\varepsilon}_t'\bar{\varepsilon}_t/2)$ by solving for the optimal $\bar{\varepsilon}_t$. We converted MATLAB's `fminsearch` routine to Fortran.

(b) Use the model solution to predict the state vector, $\mathbf{v}_{t,p}$, given $\mathbf{v}_{t-1,p}$ and $\varepsilon_{t,p}$.

(c) Calculate $\xi_{t,p} = \mathbf{v}_{t,p}^h - \mathbf{x}_t$. The unnormalized weight on particle p is given by

$$\omega_{t,p} = \frac{p(\xi_t|\mathbf{v}_{t,p})p(\mathbf{v}_{t,p}|\mathbf{v}_{t-1,p})}{g(\mathbf{v}_{t,p}|\mathbf{v}_{t-1,p},\mathbf{x}_t)} \propto \frac{\exp(-\xi_{t,p}' R^{-1}\xi_{t,p}/2)\exp(-\varepsilon_{t,p}'\varepsilon_{t,p}/2)}{\exp(-(\varepsilon_{t,p}-\bar{\varepsilon}_t)'(\varepsilon_{t,p}-\bar{\varepsilon}_t)/2)}.$$

Without adaptation, $\bar{\varepsilon}_t = 0$ and $\omega_{t,p} = p(\xi_t|\mathbf{v}_{t,p})$, as in a basic bootstrap particle filter. The time-t contribution to the log-likelihood is $\ell_t^{model} = \sum_{p=1}^{N_p}\omega_{t,p}/N_p$.

(d) Normalize the weights, $W_{t,p} = \omega_{t,p}/\sum_{p=1}^{N_p}\omega_{t,p}$. Then use systematic resampling with replacement from the swarm of particles as described in Kitagawa (1996) to get a set of particles that represents the filter distribution and reshuffle $\{\mathbf{v}_{t,p}\}_{p=1}^{N_p}$ accordingly.

3. The log-likelihood is $\log \ell^{model} = \sum_{t=1}^{T}\log \ell_t^{model}$.

E VECTOR AUTOREGRESSION MODEL

The structural VAR model is given by

$$A_0 y_t = a_0 + A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \quad t = 1,\ldots,T,$$

where $\varepsilon_t \sim N(0, I)$. The reduced-form VAR model is obtained by inverting A_0 and is given by

$$y_t = b_0 + B_1 y_{t-1} + \cdots + B_p y_{t-p} + \upsilon_t, \quad t = 1,\ldots,T,$$

where $b_0 = A_0^{-1}a_0$ is a $K\times 1$ vector of intercepts, $B_j = A_0^{-1}A_j$ are $K\times K$ coefficient matrices for $j = 1,\ldots,p$, $\upsilon_t = A_0^{-1}\varepsilon_t$ is a $K\times 1$ vector of shocks that has a multivariate normal distribution with zero mean and variance-covariance matrix Σ, and y is a $K\times 1$ vector of endogenous variables.

The VAR is estimated with data generated from the baseline model or analogous variables in U.S. data. The variables are ordered as in Christiano et al. (2005). We rewrite the model as $Y_T = \beta X + U$ and calculate the least squares estimates, $\hat{\beta}$ and $\hat{\Sigma}$. For example, when $p = 4$ the parameters are $\beta = [b_0, B_1, B_2, B_3, B_4]$ and the regressors are $X = [\mathbf{1}, Y_{T-1}', Y_{T-2}', Y_{T-3}', Y_{T-4}']'$ where $Y_{T-i} = [y_{1-i},\ldots,y_{T-i}]$ and $U = [\upsilon_1,\ldots,\upsilon_T]$. The data includes the financial uncertainty series in Ludvigson et al. (2017), per capita real output growth, the GDP implicit price deflator inflation rate, real wage growth, the risk premium, and the federal funds rate. When using artificial data, we set $p = 1$, consistent with the structural model. When using U.S. data, we calculate the Bayesian information criterion (BIC). According to the BIC, the data prefers one lag, so we focus on that specification. The structural shocks are identified by a Cholesky decomposition, $\hat{\Sigma} = (\hat{A}_0^{-1})'\hat{A}_0^{-1}$.

F WELFARE COST DERIVATION

The representative household's preferences are given by

$$E_t W(\tilde{c}, n) = E_t \sum_{j=t}^{\infty} \beta^{j-t} \left[\frac{\tilde{c}_j^{1-\gamma} - 1}{1-\gamma} - \chi \frac{n_j^{1+\eta}}{1+\eta} \right].$$

When $\gamma \neq 1$, the time-t welfare cost, λ_t, satisfies

$$\begin{aligned}
E_t W(\tilde{c}^H, n^H) &\equiv E_t W((1-\lambda_t)\tilde{c}^L, n^L) \\
&= E_t \sum_{j=t}^{\infty} \beta^{j-t} \left[\frac{((1-\lambda_t)\tilde{c}_j^L)^{1-\gamma} - 1}{1-\gamma} - \chi \frac{(n_j^L)^{1+\eta}}{1+\eta} \right] \\
&= (1-\lambda_t)^{1-\gamma} E_t \sum_{j=t}^{\infty} \beta^{j-t} \frac{(\tilde{c}_j^L)^{1-\gamma}}{1-\gamma} - \sum_{j=t}^{\infty} \frac{\beta^{j-t}}{1-\gamma} - \chi E_t \sum_{j=t}^{\infty} \beta^{j-t} \frac{(n_j^L)^{1+\eta}}{1+\eta} \\
&= (1-\lambda_t)^{1-\gamma} \left(E_t \sum_{j=t}^{\infty} \beta^{j-t} \frac{(\tilde{c}_j^L)^{1-\gamma} - 1}{1-\gamma} + \sum_{j=t}^{\infty} \frac{\beta^{j-t}}{1-\gamma} \right) - \sum_{j=t}^{\infty} \frac{\beta^{j-t}}{1-\gamma} - \chi E_t \sum_{j=t}^{\infty} \beta^{j-t} \frac{(n_j^L)^{1+\eta}}{1+\eta} \\
&= (1-\lambda_t)^{1-\gamma} \left(E_t W^c(\tilde{c}^L) + \frac{1}{(1-\gamma)(1-\beta)} \right) - \frac{1}{(1-\gamma)(1-\beta)} - E_t W^n(n^L).
\end{aligned}$$

Solving for λ_t yields (17) in the main text.

G ESTIMATION DIAGNOSTICS

The section provides additional results related to the nonlinear estimation of our baseline model. Table 3 shows how the unconditional moments for the observables match equivalent statistics in the data. We also show trace plots of our posterior draws (figure 12), kernel densities of the estimated parameters (figure 13), and the median filtered observables (figure 14) and shocks (figure 15).

	Real GDP Growth ($\hat{y}_t$)		Inflation Rate (π_t)		Interest Rate (i_t)	
	Mean	SD	Mean	SD	Mean	SD
Data	1.41	2.40	2.18	0.99	3.68	2.77
Model	1.78	2.27	2.56	0.93	4.83	1.43
	$(1.10, 2.49)$	$(1.56, 3.25)$	$(1.99, 3.11)$	$(0.63, 1.37)$	$(3.59, 6.05)$	$(0.89, 2.16)$
	Autocorrelations			Cross-Correlations		
	$(\hat{y}_t, \hat{y}_{t-1})$	(π_t, π_{t-1})	(i_t, i_{t-1})	$(\hat{y}_t, \pi_t)$	$(\hat{y}_t, i_t)$	(π_t, i_t)
Data	0.31	0.63	0.99	0.03	0.18	0.50
Model	0.27	0.76	0.91	-0.11	0.16	0.32
	$(0.02, 0.51)$	$(0.63, 0.86)$	$(0.83, 0.96)$	$(-0.46, 0.19)$	$(-0.09, 0.44)$	$(-0.16, 0.68)$

Table 3: Unconditional moments. For each draw from the posterior distribution, we run $10{,}000$ simulations with the same length as the data. To compute the moments, we first calculate time averages and then the means and quantiles across the simulations. The values in parentheses are $(5\%, 95\%)$ credible sets. All values are annualized net rates.

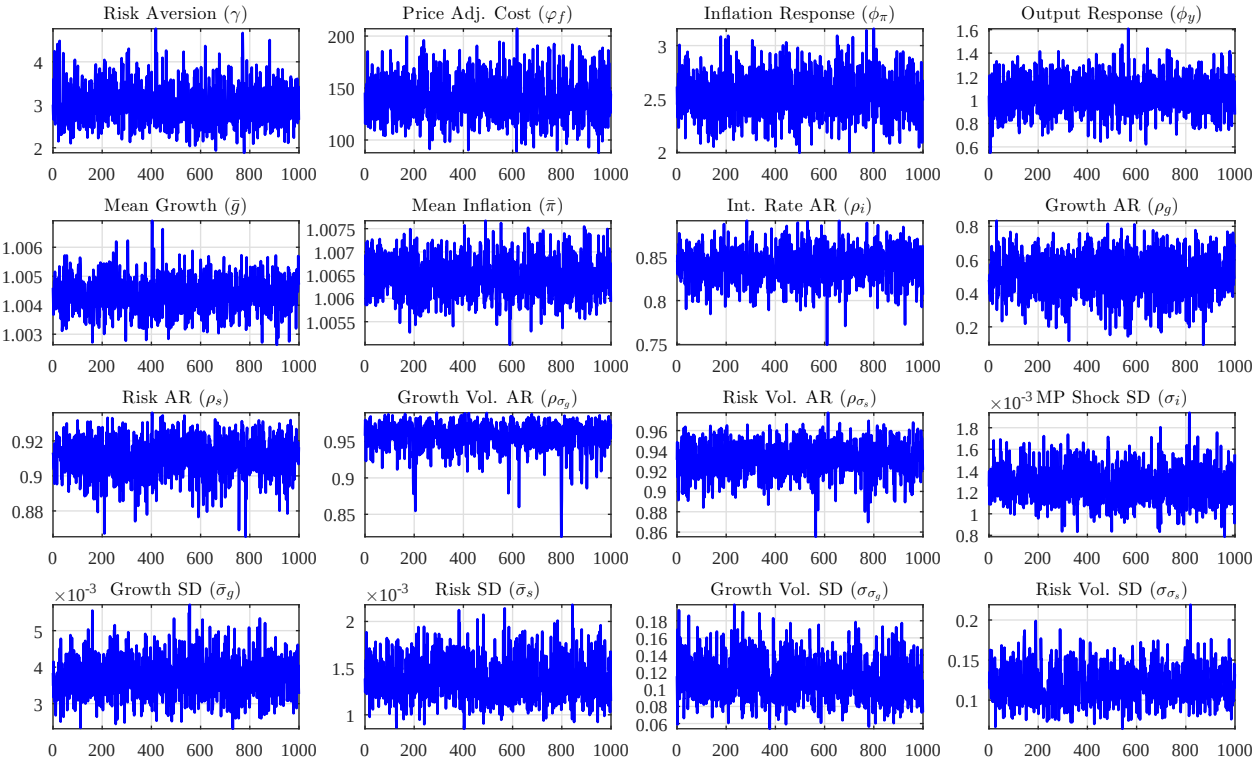

Figure 12: Trace plots. We obtained 100,000 draws from each posterior distribution and kept every 100th draw.

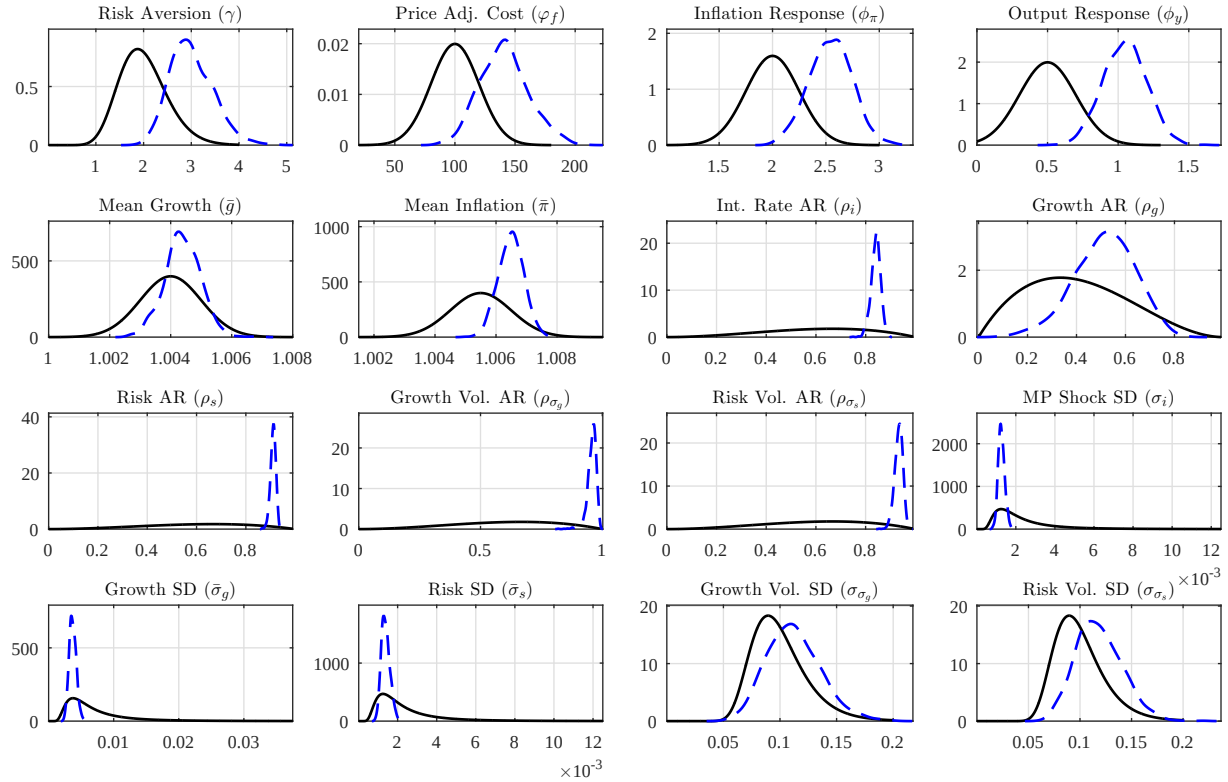

Figure 13: Prior (solid lines) and posterior kernel (dashed lines) densities of the estimated parameters.

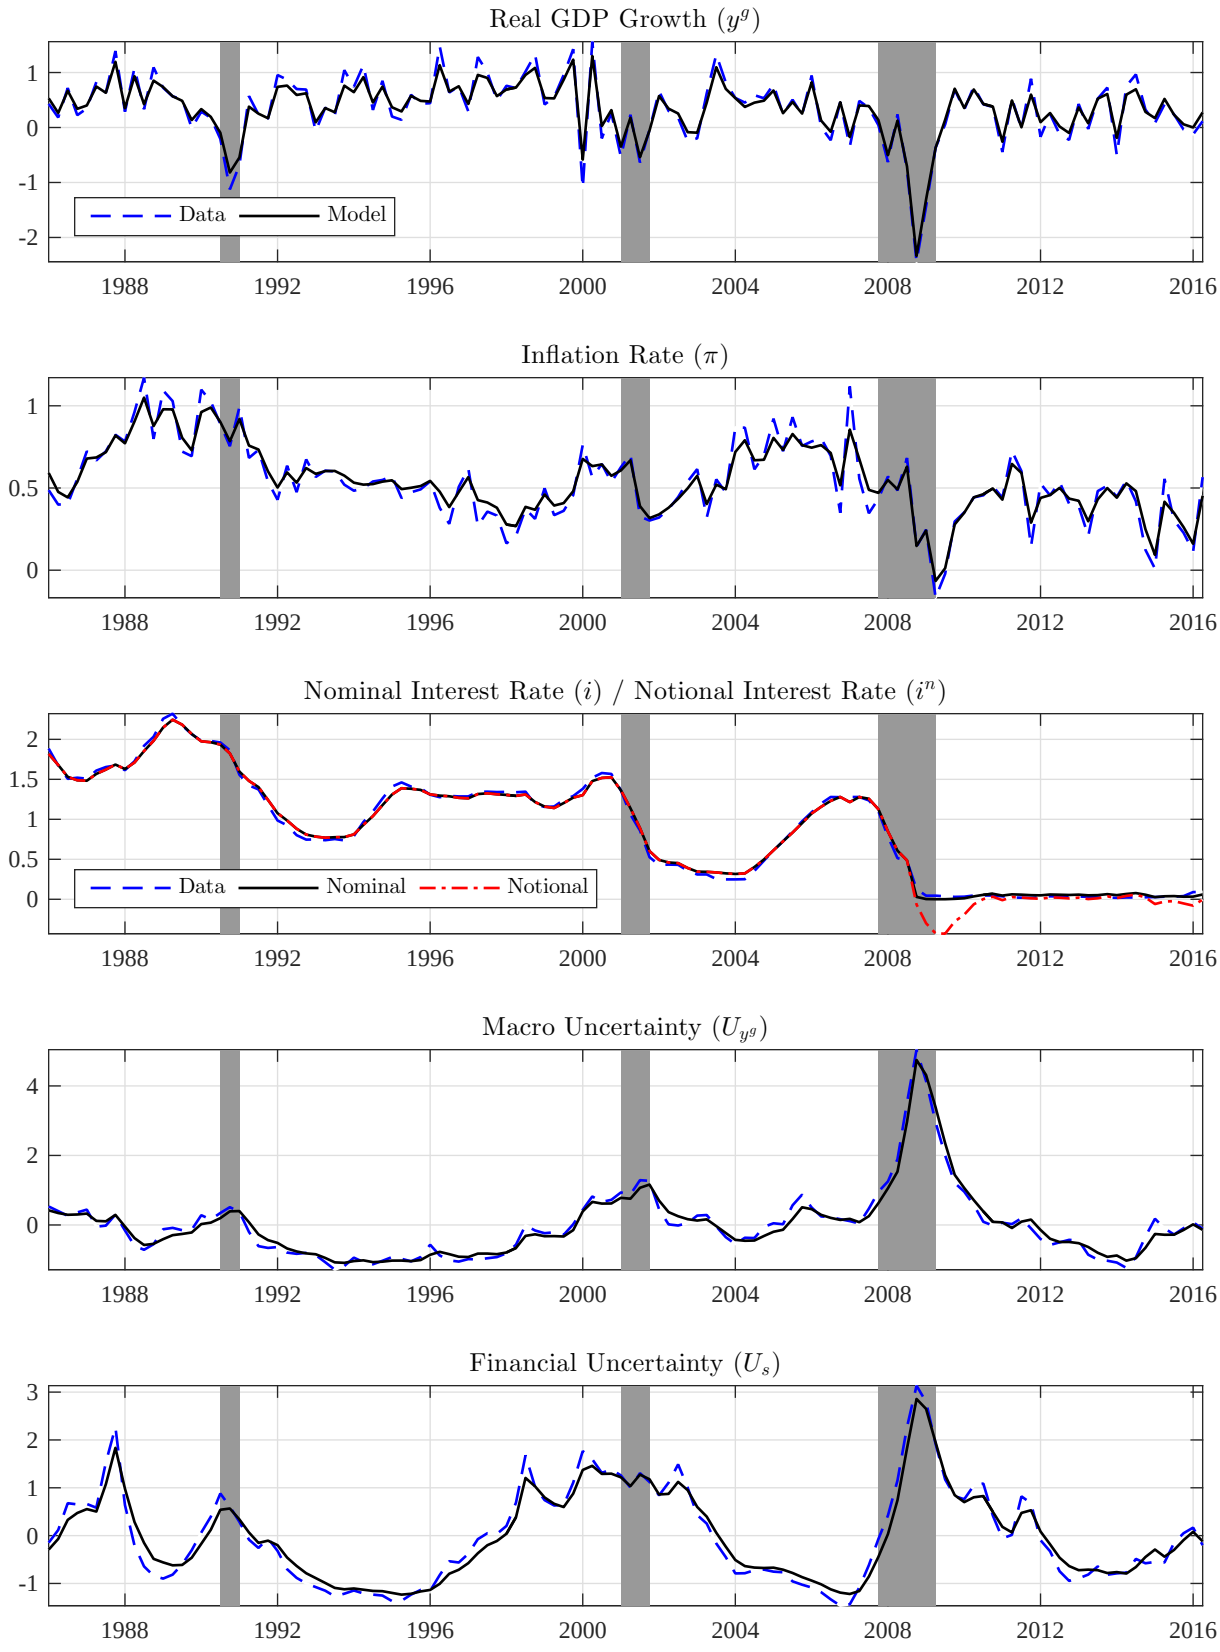

Figure 14: Time paths of the data (dashed line) and the median filtered series from the baseline model (solid line).

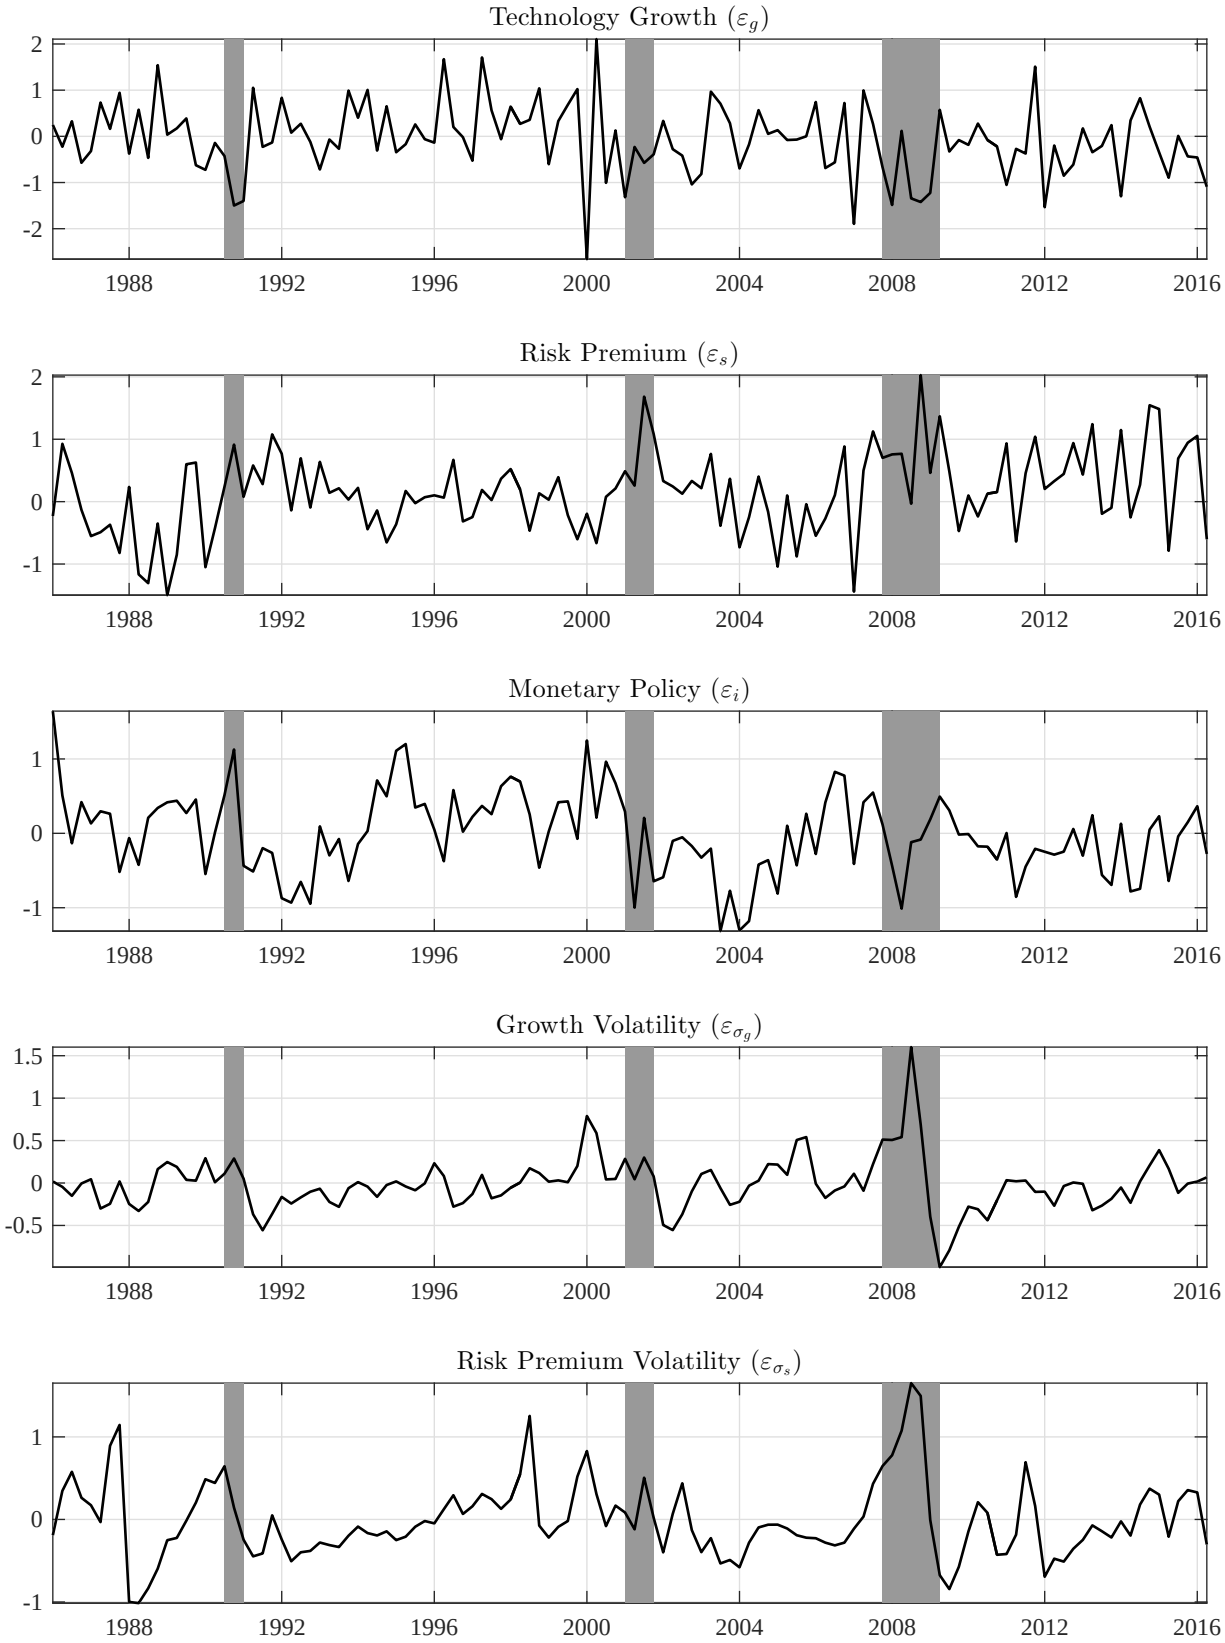

Figure 15: Median paths of the estimated shocks normalized by their respective posterior mean standard deviation.