A New Way to Quantify the Effect of Uncertainty*

Alexander W. Richter Nathaniel A. Throckmorton

April 24, 2017

ABSTRACT

This paper develops a new method to quantify the effects of uncertainty using estimates from a nonlinear New Keynesian model. The model includes an occasionally binding zero lower bound constraint on the nominal interest rate, which creates time-varying endogenous uncertainty, and two exogenous types of time-varying uncertainty—a volatility shock to technology growth and a volatility shock to the risk premium. A filtered third-order approximation of the Euler equation shows consumption uncertainty on average reduced consumption by about 0.06% and the peak effect was 0.15% during the Great Recession. Other higher-order moments such as inflation uncertainty, technology growth uncertainty, consumption skewness, and inflation skewness had smaller effects. Technology growth volatility explained most of the changes in uncertainty, but risk premium volatility had a major role in the last two recessions.

Keywords: Bayesian Estimation; Uncertainty; Stochastic Volatility; Zero Lower Bound
JEL Classifications: C11; D81; E32; E58

*Richter, Research Department, Federal Reserve Bank of Dallas, and Auburn University, 2200 N. Pearl Street, Dallas, TX 75201 (alex.richter@dal.frb.org); Throckmorton, Department of Economics, College of William & Mary, P.O. Box 8795, Williamsburg, VA 23187 (nathrockmorton@wm.edu). We thank Oliver De Groot, Ben Johannsen, Campbell Leith, and Mike Plante for helpful comments that improved the paper. We also thank Eric Walter for supporting the SciClone computer cluster at the College of William and Mary and Chris Stackpole for assisting with the Bull cluster at the Federal Reserve Bank of Kansas City. The views expressed in this 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

Intuition tells us that economic uncertainty can amplify the business cycle through a precautionary savings motive.[1] The question is whether that channel is quantitatively significant. Unfortunately, it is difficult to answer for two main reasons. One, uncertainty is unobserved so there is disagreement on what constitutes a good measure. Until recently, researchers have primarily relied on various proxies for economic 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 and only loosely connected with the theoretical definition of uncertainty.[2] Two, uncertainty is endogenous. Not only can uncertainty affect economic activity, as intuition suggests, what is happening in the economy can also affect uncertainty. A recent example is when the Great Recession forced the Fed to cut its policy rate to its zero lower bound (ZLB) in December 2008. Plante et al. (2016) find the constraint made real GDP more responsive to shocks that hit the economy and therefore increased the uncertainty surrounding future real GDP growth because it restricted the Fed's ability to stabilize the economy. As a consequence, it is very difficult to determine the causal effect of uncertainty using reduced-form empirical methods. To deal with the identification problem, many researchers have turned to dynamic stochastic general equilibrium (DSGE) models.

A large segment of the literature that uses DSGE models introduces stochastic volatility, which makes it easier to determine the effects of uncertainty by computing impulse responses to temporary increases in volatility.[3] While that exercise addresses the identification problem, it is unable to determine the effect of the uncertainty that arises endogenously in the model. This paper develops a new approach to quantifying uncertainty that addresses both problems. First, we estimate a nonlinear New Keynesian model with an occasionally binding ZLB constraint on the nominal interest rate and two stochastic volatility shocks—one to the rate of return on a 1-period nominal bond and the other to the growth rate of technology. A key aspect of our model is that it contains endogenous and exogenous sources of time-varying uncertainty. The ZLB constraint endogenously generates uncertainty that depends on how severely the Fed is constrained, the volatility shock on the nominal return introduces an exogenous source of financial or demand uncertainty, and the growth volatility shock adds an exogenous form of supply uncertainty. Second, we derive a third-order approximation of the Euler equation to decompose consumption into eight key terms—expected consumption next period, the *ex-ante* real interest rate, and the variance and skewness of consumption, technology growth, and inflation next period. We then iterate on the Euler equation to remove the effect of expected consumption. Finally, we filter the data to create a time series for each term in our decomposition and then quantify the effects of uncertainty and skewness over different horizons.

Another key aspect of our estimation procedure is that we include the macro uncertainty index developed by Jurado et al. (2015) and the financial uncertainty index from Ludvigson et al. (2017) as observables so our model produces the same fluctuations in uncertainty found in the data. We chose those indexes because they are constructed with over 100 macro and financial variables and are based on the same statistic we use to measure uncertainty—expected forecast error volatility—

[1] Other potentially important uncertainty channels include the real options effect and the financial frictions effect.

[2] See, for example, Alexopoulos and Cohen (2009), Bachmann et al. (2013), Baker et al. (2013), and Bloom (2009).

[3] Researchers have used stochastic volatility shocks to examine the effects of uncertainty surrounding a wide range of topics including fiscal policy [Born and Pfeifer (2014); Fernández-Villaverde et al. (2015)], monetary policy [Mumtaz and Zanetti (2013)], investment and firm dynamics [Bloom (2009); Bloom et al. (2014); Christiano et al. (2014); Justiniano and Primiceri (2008)], commodity prices [Plante and Traum (2012)], disaster risk [Gourio (2013)], and the real interest rate in advanced and emerging economies [Basu and Bundick (2016); Fernández-Villaverde et al. (2011)].

which removes the predictable component of a variable in each period, instead of a constant trend.

Over a 24-quarter horizon, consumption uncertainty on average reduced consumption by about 0.06% and the peak effect was 0.15% in 2009Q1, which accounted for 16.3% of the overall decline in consumption in that quarter. We chose that horizon because it is long enough to nearly eliminate the effect of expected future consumption. Inflation uncertainty and both consumption and inflation skewness had a much smaller impact on consumption, even in the Great Recession. Those results are based on a model without capital accumulation. When we extend the model so households can invest in capital, we find uncertainty had a slightly larger effect on consumption. We also conduct counterfactual simulations to decompose uncertainty into its endogenous and exogenous sources. Changes in uncertainty were typically driven by growth volatility shocks. However, risk premium volatility shocks explained about half of the increases in uncertainty in the past two recessions and the ZLB constraint was a major contributor to the uncertainty since the end of the Great Recession.

It is helpful to provide some context for our results. When we estimate a VAR with the uncertainty series ordered first using data from 1986Q1 to 2016Q2, a Cholesky decomposition implies a one standard deviation increase in financial uncertainty causes consumption to decline by about 0.25%, while macro uncertainty decreases consumption by 0.3%. It appears the reduced-form effects of uncertainty are larger than we find using our Euler equation decomposition, but there are three important caveats. One, the responses from the VAR are based on a specific second moment shock, whereas our Euler equation decomposition determines the combined effect of the shocks that best explain the data each period. Two, the responses in the VAR are based on the historical relationship between uncertainty and consumption, whereas the Euler equation decomposition is forward-looking since it depends on the expected variance of future consumption. Three, the reduced-form estimates depend on the ordering of the uncertainty series in the VAR. If the uncertainty series are ordered first, then the state of the economy has no contemporaneous effect on the impulse responses. If they are ordered last, then uncertainty has no impact effect on consumption and the responses to a financial and macro uncertainty shock are significantly smaller. We show the impulse responses from our DSGE model invalidate both of those identification assumptions.

Several papers use the same recursive identification scheme to determine the effects of uncertainty [Bachmann et al. (2013); Basu and Bundick (2016); Bekaert et al. (2013); Bloom (2009); Jurado et al. (2015)]. A notable exception is Ludvigson et al. (2017) who use event and correlation constraints to restrict the set of identified impulse responses and determine if uncertainty causes or is caused by real activity. They find financial uncertainty shocks cause a sharp decline in real activity while macro uncertainty does not. Furthermore, negative shocks to real activity increase macro uncertainty but have no clear effect on financial uncertainty. Those results invalidate the timing assumptions used in recursive identification and the more general identification method represents a substantial improvement over previous approaches, but it is unable to determine how the effects of uncertainty change over time or the underlying sources of the shocks. We account for the potential endogeneity of uncertainty through the lens of an estimated nonlinear DSGE model that allows for various shocks whose effects on uncertainty and real activity depend on the state of the economy.

Structural models used to examine the effects of uncertainty usually include stochastic volatility or a feature that endogenously generates uncertainty, but rarely both. The stochastic volatility literature typically finds an exogenous increase in volatility decreases economic activity, but the quantitative effects are often small unless the shock is large or the volatilities of multiple shocks increase. One exception is Christiano et al. (2014) whose estimates indicate that uncertainty is a major contributor to the business cycle. The risk shock in their model, however, is linked to the

credit spread, which is observable and can fluctuate for reasons besides time-varying risk. We find uncertainty can have a meaningful effect on economic activity, even if exogenous volatility shocks have a small impact. Specifically, the effect of consumption uncertainty from our Euler equation decomposition is an order of magnitude larger than the responses to either volatility shock. Those results stress the importance of accounting for how real activity endogenously affects uncertainty.

Since uncertainty can also stem from what is going on in the economy, a small but growing literature has developed models where uncertainty arises endogenously. Researchers have focused on two mechanisms—incomplete information [Bachmann and Moscarini (2012); Fajgelbaum et al. (2014); Straub and Ulbricht (2015); Van Nieuwerburgh and Veldkamp (2006)] and firm default [Gourio (2014); Navarro (2014)]. The drawback with those models is that they focus on one type of uncertainty and, due to their complexity, there is no tractable way to take them to the data using likelihood based methods. One exception is Plante et al. (2016) who estimate a DSGE model with a ZLB constraint, but they focus on the correlation between uncertainty and economic activity rather than causation. This paper bridges the gap between the stochastic volatility and endogenous uncertainty literatures by estimating a nonlinear DSGE model with both exogenous and endogenous sources of time-varying uncertainty to determine the causal effect of uncertainty.

The paper proceeds as follows. Section 2 estimates two VARs to show the reduced-form effects of uncertainty. Section 3 describes our structural model and uncertainty sources. 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 consumption. Section 6 adds capital accumulation to our model to show how it affects our decomposition. Section 7 concludes.

2 REDUCED-FORM EVIDENCE

To provide a reference for the empirical results from our DSGE model, we first show the effects of macro and financial uncertainty on consumption in a VAR model. The structural model is given by

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

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

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

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 $j = 1, \ldots, p$, $\varepsilon_t = A_0^{-1} u_t$ 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.

We include 9 variables similar to those in our DSGE model and the S&P 500 index to control for the level effects of equity prices, as in Bloom (2009) and Jurado et al. (2015). The quarterly data consists of the financial uncertainty series in Ludvigson et al. (2017), $FinU$, the macro uncertainty series in Jurado et al. (2015), $MacroU$, per capita real GDP, per capita real consumption (nondurables + services), the GDP implicit price deflator, real wages (production and nonsupervisory employees), the risk premium (BAA corporate bond yield − 10 year treasury yield), utilization-adjusted productivity (Fernald (2012)), the federal funds rate, and the S&P 500 index. A detailed description of our data sources is provided in Appendix A. All variables except productivity enter the VAR in logs and, aside from the uncertainty series, the ordering follows Christiano et al. (2005). We estimate the model with up to four lags and calculate the Bayesian information

criterion (BIC) to determine the appropriate lag. For example, when $p = 4$ the parameters are $\beta = \text{vec}(b_0, B_1, B_2, B_3, B_4)$ and the regressors are $X_{t-1} = I_{10} \otimes [1, y'_{t-1}, y'_{t-2}, y'_{t-3}, y'_{t-4}]$. We write the model as $y_t = X_{t-1}\beta + A_0^{-1}u_t$ and calculate the least squares estimates, $\hat{\beta}$ and $\hat{\Sigma}$. The structural shocks are identified by a Cholesky decomposition, $\hat{\Sigma} = (\hat{A}_0^{-1})'\hat{A}_0^{-1}$. According to the BIC, the data prefers a model with one lag, so we use that specification for the impulse responses.

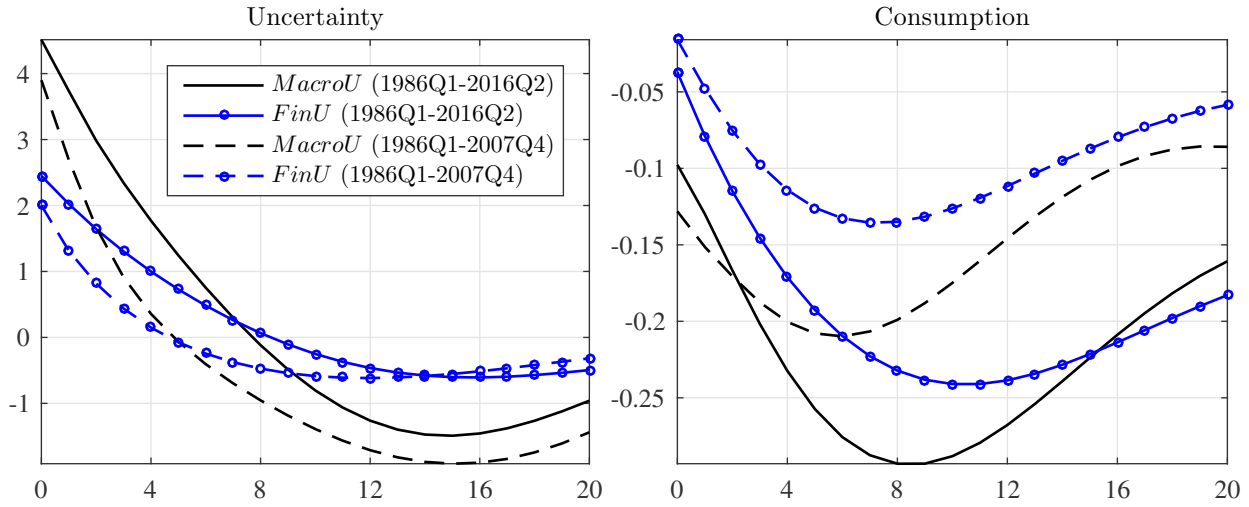

(a) VAR model where the uncertainty series are ordered first.

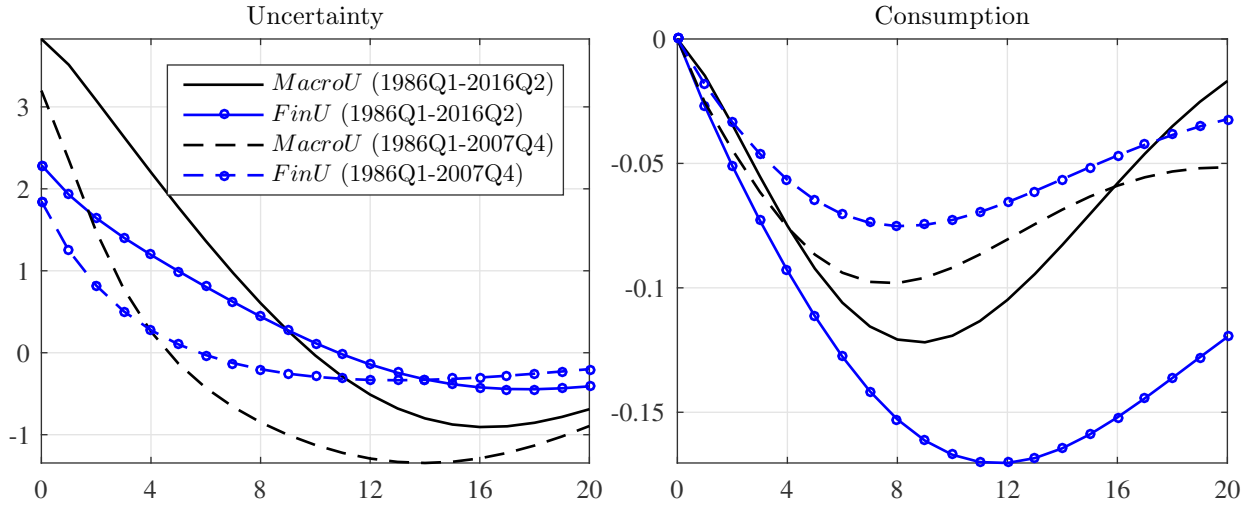

(b) VAR model where the uncertainty series are ordered last.

Figure 1: Impulse responses to a one standard deviation increase in macro and financial uncertainty.

Figure 1 shows the responses of uncertainty and consumption to a one standard deviation shock to the macro (no circles) and financial (circles) uncertainty indexes. Since there is no consensus on where to place the uncertainty series in a VAR, we estimate two models—one with the uncertainty series ordered first (figure 1a) and one with them ordered last (figure 1b). We estimate those models using two samples: 1986Q1-2016Q2 (solid lines) and 1986Q1-2007Q4 (dashed lines). The first sample is the one we use to estimate our DSGE model and the second sample removes the influence of the ZLB, which is an endogenous source of uncertainty according to our DSGE model. The estimates show both shocks have a smaller effect when the uncertainty series are ordered last,

since changes in the state of the economy explain some of the movements in uncertainty. In the full sample, an increase in macro (financial) uncertainty reduces consumption by about 0.3% (0.25%) when uncertainty is ordered first, whereas it declines by only 0.12% (0.17%) when it is ordered last. When we remove the ZLB period from our sample the peak effect on consumption also declines, even though the uncertainty response is similar. Those results indicate that the ZLB amplifies the effect of uncertainty shocks, as our DSGE model predicts. Since the responses are not state dependent, the VAR model overstates the effect of uncertainty in the pre-ZLB period when using estimates based on the entire sample. Interestingly, we will show when uncertainty is ordered last, the responses in the pre-ZLB sample are similar to the average effect of consumption uncertainty in our DSGE model when the ZLB does not bind, despite significant differences in methodology.[4]

3 STRUCTURAL MODEL AND MEASURES OF UNCERTAINTY

Our reduced-form estimates demonstrate that the effects of uncertainty are heavily dependent on the specification of the VAR (sample, variables, and order) and the specific shocks. This section develops a nonlinear DSGE model that can identify the marginal effects of uncertainty and skewness over time. We use a New Keynesian similar to An and Schorfheide (2007), except it is augmented to include a ZLB constraint, which introduces an endogenous source of time-varying uncertainty, and two exogenous sources of time-varying uncertainty: a volatility shock to technology growth (supply uncertainty) and a volatility shock to the return on a nominal bond (demand uncertainty).

3.1 FIRMS

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

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

$$\sigma_{\varepsilon,t} = \bar{\sigma}_\varepsilon(\sigma_{\varepsilon,t-1}/\bar{\sigma}_\varepsilon)^{\rho_{\sigma_\varepsilon}} \exp(\sigma_\xi \xi_t),\ 0 \leq \rho_{\sigma_\varepsilon} < 1,\ \xi \sim \mathbb{N}(0,1), \tag{2}$$

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

The final goods firm purchases $y_t(f)$ units from each intermediate firm to produce the final good, $y_t \equiv [\int_0^1 y_t(f)^{(\theta-1)/\theta} df]^{\theta/(\theta-1)}$, according to a Dixit and Stiglitz (1977) aggregator, where $\theta > 1$ controls the elasticity of substitution between any two goods. It then maximizes dividends to determine its demand function for intermediate good f, $y_t(f) = (p_t(f)/p_t)^{-\theta} y_t$, where $p_t = [\int_0^1 p_t(f)^{1-\theta} df]^{1/(1-\theta)}$ is the price level. Following Rotemberg (1982), each intermediate firm pays a cost to adjust its price level, $adj_t(f) \equiv \varphi_f[p_t(f)/(\bar{\pi} p_{t-1}(f)) - 1]^2 y_t/2$, where $\varphi_f > 0$ scales the size of the cost and $\bar{\pi}$ is the gross inflation rate along the balanced growth path.[5] Therefore, firm f chooses $n_t(f)$ and $p_t(f)$ to maximize the expected discounted present value of future dividends, $E_t \sum_{k=t}^{\infty} q_{t,k} d_k(f)$, 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 $t+1$, $q_{t,k} \equiv \prod_{j=t+1}^{k>t} q_{j-1,j}$, $d_t(f) = p_t(f) y_t(f)/p_t - w_t n_t(f) - adj_t(f)$, and a tilde denotes a variable relative to the level of technology ($\tilde{x} = x/z$). In symmetric equilibrium, all firms make identical decisions (i.e., $p_t(f) =$

[4] Appendix B shows our qualitative results are robust to using a longer data sample similar to Jurado et al. (2015).

[5] Richter and Throckmorton (2016) show the parameter estimates with Calvo and Rotember pricing are very similar.

p_t, $n_t(f) = n_t$, and $y_t(f) = y_t$), so the production function and the optimality conditions reduce to

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

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

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

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, $E_0 \sum_{t=0}^{\infty} \beta^t[(c_t/z_t)^{1-\gamma}/(1-\gamma) - \chi n_t^{1+\eta}/(1+\eta)]$, 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. (2016), s is a shock to the rate of return on the nominal bond and it evolves according to

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

$$\sigma_{\upsilon,t} = \bar{\sigma}_{\upsilon}(\sigma_{\upsilon,t-1}/\bar{\sigma}_{\upsilon})^{\rho_\sigma} \exp(\sigma_\zeta \zeta_t),\ 0 \leq \rho_{\sigma_\upsilon} < 1,\ \zeta \sim \mathbb{N}(0,1), \tag{7}$$

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

The first order conditions to the household's 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 consumption Euler equation, which we use to identify the effect of uncertainty.

3.3 MONETARY POLICY The central bank sets the gross nominal interest rate according to

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

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

where y^{gdp} is real GDP (i.e., the level of output minus the resources lost due to price adjustment costs), i^n is the gross notional interest rate, $\bar{\imath}$ and $\bar{\pi}$ are the inflation and interest rate targets, which equal their values along the balanced growth path, and ϕ_π and ϕ_y are the responses to deviations of inflation from the target rate and deviations of real GDP growth from the balanced growth rate.[6]

[6] We also estimated the model with the potential output gap, y_t^{gdp}/y_t^p, instead of the growth rate gap in the Taylor rule. We decided to use the specification in (11) because that model had a significantly higher marginal data density.

3.4 COMPETITIVE EQUILIBRIUM

The aggregate resource constraint is given by

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

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

In order 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, $\{\tilde{c}_t, \tilde{y}_t, \tilde{y}_t^{gdp}, n_t\}_{t=0}^{\infty}$, prices, $\{\tilde{w}_t, mc_t, i_t, i_t^n, \pi_t^{gap}\}_{t=0}^{\infty}$, and exogenous variables, $\{s_t, g_t, \sigma_{\varepsilon,t}, \sigma_{\upsilon,t}\}_{t=0}^{\infty}$, that satisfy the detrended equilibrium system, (1)-(13), given the initial conditions, $\{c_{-1}, i_{-1}^n, s_0, a_0, \nu_0, \sigma_{\varepsilon,0}, \sigma_{\upsilon,0}\}$, and the five sequences of shocks, $\{\varepsilon_t, \upsilon_t, \nu_t, \xi_t, \zeta_t\}_{t=1}^{\infty}$.

3.5 MEASURES OF UNCERTAINTY

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

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

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

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 models. Following Plante et al. (2016), the endogenous time-varying uncertainty surrounding trended consumption growth, $c_t^g \equiv g_t \tilde{c}_t/\tilde{c}_{t-1}$, is given by

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

which is the same statistic we use to measure exogenous uncertainty, except it is calculated with an endogenous variable. In contrast with the exogenous types of uncertainty, this form of uncertainty also fluctuates due to events that happen in the economy. For example, when the nominal interest rate is stuck at its ZLB, the economy is more sensitive to first moment shocks that adversely affect the economy, which increases the endogenous uncertainty about consumption growth. The ZLB also generates uncertainty by amplifying the effect of the exogenous volatility shocks. When the ZLB does not bind, first and second moment shocks still affect $U_{c^g,t}$ but the magnitudes are smaller.

4 NUMERICAL METHODS AND EULER EQUATION 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 shocks 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, υ, and ν, 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 nonlinear policy functions for consumption and inflation. To estimate the model, we also create a policy function for consumption growth uncertainty, (14), by interpolating the policy function for consumption given the updated state and then numerically integrating using the Rouwenhorst weights. See Appendix C for a detailed 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 deflator, DEF, the average federal funds rate, FFR, the macro uncertainty series in Jurado et al. (2015), $MacroU$, and the financial uncertainty series in Ludvigson et al. (2017), $FinU$, from 1986Q1 to 2016Q2. The vector of observables is given by

$$\hat{\mathbf{x}}^{data} \equiv \begin{bmatrix} \log(RGDP_t/CNP_t) - \log(RGDP_{t-1}/CNP_{t-1}) \\ \log(DEF_t/DEF_{t-1}) \\ \log(1 + FFR_t/100)/4 \\ (MacroU - \mu_{MacroU})/\sigma_{MacroU} \\ (FinU - \mu_{FinU})/\sigma_{FinU} \end{bmatrix},$$

where μ and σ denote mean and standard deviation across time. We use the $MacroU$ and $FinU$ uncertainty series as observables to help inform the stochastic volatility parameters in our model and ensure that our model produces the same fluctuations in uncertainty that are found in the data. The benefit of these particular uncertainty series is that they are calculated the exact same way as the uncertainty measures in our model, so they distinguish between uncertainty and conditional volatility, and they reflect the common variation in more than 100 macro and financial time series.

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	$\sigma_{me,macrou}$	0.44721
Number of Particles	N_p	40,000	Financial Uncertainty ME SD	$\sigma_{me,finu}$	0.44721

Table 1: Calibrated parameters for the DSGE model and estimation procedure.

We calibrate four parameters that are not well-informed by our data (table 1). The discount factor along the balanced growth path, $\bar{\beta}$, is calibrated to 0.9987 to match $(1/T)\sum_{t=1}^{T}(1+G_t/400)(1+\Pi_t)/(1+FFR_t/100)^{1/4}$, where T is the sample size, G_k is the annual utilization-adjusted growth rate of technology from Fernald (2012) and $\Pi_k = \log(DEF_k/DEF_{k-1})$. 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 intermediate 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 macro estimate in Peterman (2016).

We use Bayesian methods to estimate the remaining parameters. 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-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).

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

$$\hat{\mathbf{x}}_t^{model} = \left[\log(g_t \tilde{y}_t^{gdp}/\tilde{y}_{t-1}^{gdp}) \quad \log(\pi_t) \quad \log(i_t) \quad (U_{c^g,t} - \mu_{U_{c^g}})/\sigma_{U_{c^g}} \quad (U_{s,t} - \mu_{U_s})/\sigma_{U_s}\right],$$

$\xi \sim \mathbb{N}(0, \Sigma)$ is a vector of MEs, and $\Sigma = \text{diag}([\sigma^2_{me,y^g}, \sigma^2_{me,\pi}, \sigma^2_{me,i}, \sigma^2_{me,macrou}, \sigma^2_{me,finu}])$. Following Herbst and Schorfheide (2016), the ME variance of real GDP growth, the inflation rate, and both uncertainty series is set to 20% of their variance in the data. However, we set the ME variance for the policy rate to 2% of its variance in the data because the federal funds rate it is less noisy and it affects the level of uncertainty predicted by the model. We chose to link consumption growth uncertainty to the macro uncertainty index and risk premium uncertainty to the financial uncertainty index because Ludvigson et al. (2017) show 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, consumption growth uncertainty is equal to real GDP uncertainty and it is determined endogenously, 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. 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 final analysis is based on a sample of $1{,}000$ draws.

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

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

where var_t, skew_t, and cov_t denote the variance, third moment, and covariance of a future 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 deviations from the balanced growth path. The decomposition shows consumption depends on its expected value next period, the *ex-ante* real interest rate, the

variance and skewness of consumption, technology growth, and inflation next period, and the covariances between those variables next period. We omitted higher-order covariance terms, such as $\text{cov}_t(\hat{\pi}_{t+1}^2, \hat{c}_{t+1})$, as well as fourth-order and higher terms because we found they had almost no effect on consumption in our sample. The variance terms quantify the effect of the three types of uncertainty and the skewness terms determine the effects of upside and downside risk. The first two covariance terms capture the risk premia associated with unexpected inflation or technology growth, which are partially offset by the covariance between those variables. Higher risk aversion means households are less willing to intertemporally substitute consumption goods, which makes them more sensitive to changes in risk premia or the variance and skewness of future consumption.

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

$$\begin{aligned}\hat{c}_t \approx\ & E_t\hat{c}_{t+q} - \tfrac{1}{\gamma}E_t\textstyle\sum_{j=1}^{q}\hat{r}_{t+j-1} \\ & - \textstyle\sum_{j=1}^{q}(\text{cov}_t(\hat{\pi}_{t+j},\hat{c}_{t+j}) + \text{cov}_t(\hat{g}_{t+j},\hat{c}_{t+j}) + \tfrac{1}{\gamma}\text{cov}_t(\hat{\pi}_{t+j},\hat{g}_{t+j})) \\ & - \tfrac{1}{2\gamma}\textstyle\sum_{j=1}^{q}(\text{var}_t\,\hat{g}_{t+j} + \text{var}_t\,\hat{\pi}_{t+j} + \gamma^2\,\text{var}_t\,\hat{c}_{t+j}) \\ & + \tfrac{1}{6\gamma}\textstyle\sum_{j=1}^{q}(\text{skew}_t\,\hat{g}_{t+j} + \text{skew}_t\,\hat{\pi}_{t+j} + \gamma^3\,\text{skew}_t\,\hat{c}_{t+j}),\end{aligned} \tag{15}$$

where $q \geq 1$ is the forecast horizon. The sum over q quarters of each variance term represents the marginal effect of a given type of uncertainty, conditional on expected consumption in quarter q. As q becomes sufficiently large, the conditional expectation drops out of the decomposition and we can identify the unconditional effects of each type of uncertainty. Over a 1-quarter horizon, expected consumption is a strong indicator of current consumption, but that result hides the effect higher order moments in future quarters have on expectations. By decomposing expected future consumption, we can show how uncertainty, skewness, and the risk premia affect economic activity over longer horizons, while taking into account the distributions of all five shocks and the endogenous dynamics that stem from those shocks. To quantify those effects, we first create a policy function for each term in the decomposition using $100{,}000$ q-quarter simulations initialized at each node in the discretized state space. In total there are $10q+1$ policy functions. We then filter the data and use the median states and shocks to interpolate each term in the decomposition for the entire sample, which produces a time series that shows how each component affects consumption.

5 ESTIMATION RESULTS

This section provides our main results. We first report the posterior parameter estimates and then show the results of our Euler equation decomposition. To better understand the mechanisms, we conclude by showing the impulse responses and the results of a nonlinear variance decomposition.

5.1 PRIOR AND POSTERIOR DISTRIBUTIONS

The first four columns of table 2 display the estimated parameters and information about the prior distributions. The means and standard deviations are in line with Gust et al. (2016), who were the first to estimate a nonlinear DSGE model with an occasionally binding ZLB constraint. The prior for the coefficient of relative risk aversion is taken from An and Schorfheide (2007). The priors for the balanced growth rates of technology and inflation are based on the average per capita GDP growth rate and the GDP deflator 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

	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)	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 ($\rho_{\sigma_\varepsilon}$)	Beta	0.6000	0.2000	0.95721	0.01890	0.92614	0.98109
Risk SV Persistence (ρ_{σ_υ})	Beta	0.6000	0.2000	0.93308	0.01617	0.90404	0.95725
Int. Rate Shock SD (σ_ν)	IGam	0.0025	0.0025	0.00127	0.00017	0.00102	0.00157
Tech Shock SD (σ_ε)	IGam	0.0075	0.0075	0.00371	0.00054	0.00288	0.00463
Risk Shock SD (σ_υ)	IGam	0.0025	0.0025	0.00139	0.00022	0.00107	0.00177
Tech SV Shock SD (σ_ξ)	IGam	0.1000	0.0250	0.11216	0.02350	0.07647	0.15372
Risk SV Shock SD (σ_ζ)	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.

well as stronger responses that could explain data during the ZLB period. The persistence parameters are diffuse, but most of the means are set to 0.6 since we expect a modest degree of persistence is needed to explain the data, especially in the ZLB period. 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 linear models without a ZLB constraint.

The last four columns display the posterior means, standard deviations, and 90% credible sets for the estimated parameters. Many of the parameter estimates differ from Gust et al. (2016) and Plante et al. (2016), who estimate similar nonlinear models but without stochastic volatility. The low frequency movements in the $MacroU$ and $FinU$ uncertainty series coupled with the sharp increases in both series during the Great Recession produce large stochastic volatility shocks that are highly persistent. 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 of interest rate smoothing and a strong response to the real GDP growth and inflation gaps, which are necessary to explain the long ZLB period because real GDP growth and inflation were only slightly below their targets throughout most of that 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 under-represent the effects of the ZLB period. The mean coefficient of relative risk aversion is within the range of estimates typically found in the literature (e.g., An and Schorfheide (2007)), but the price adjustment cost parameter is larger than Gust et al. (2016) and Plante et al. (2016) estimate. Appendix E provides further information about the estimation procedure including the trace plots, the prior and posterior kernel densities of the parameters, the median filtered shocks and observables, the unconditional moments, and the conditional and unconditional ZLB durations.

5.2 EULER EQUATION DECOMPOSITION Figure 2 shows the filtered time series of the terms in the Euler equation decomposition in (15) over different forecast horizons. The values on the vertical axis of each subplot denote the marginal effects on current consumption in percentage point deviations from the balanced growth path. The top panel shows the decomposition over a 1-quarter horizon. We separate the first order terms (left panel) and the higher order terms (right panel) so the marginal effects of expected consumption and the real interest rate do not drown out the effects of the uncertainty and skewness terms. We also plot current consumption in the top left panel so it is easier to see the contribution of each term. Over a 1-quarter horizon, the changes in consumption are almost entirely driven by expectations about consumption next quarter. The real interest rate had a smaller role, typically reducing consumption by less than 0.14%. The peak effect was about -0.37% during the Great Recession, but that effect declined as the economy rebounded.

The higher-order terms show uncertainty about consumption in the next quarter had time-varying adverse effects on current consumption and it never had a positive effect, in contrast with the empirical results in Ludvigson et al. (2017). The effect of consumption uncertainty peaked during each of the last two recessions, but it was much larger during the Great Recession because the ZLB constraint made the economy more sensitive to adverse shocks and increased the expected volatility of future consumption. Even during the Great Recession, however, the quantitative effect of consumption uncertainty was small. The peak increase in uncertainty reduced consumption by less than 0.01%. Furthermore, consumption skewness and both inflation uncertainty and inflation skewness had very little effect on current consumption. However, these magnitude are misleading because they hide the effects that uncertainty and skewness have on expected future consumption.

The middle left panel shows how expected future consumption affected current consumption over horizons up to 24 quarters. Once again, we plot the filtered time series for current consumption as a percent deviation from the balanced growth path. In most periods, the differences between current and expected consumption were much larger over horizons beyond 1-quarter, which indicates that other factors, such as the real interest rate and uncertainty, explain a larger fraction of the fluctuations in consumption. We decided to focus on a 24-quarter horizon because it is long enough that expected consumption barely matters for current consumption. For example, in 2009Q2—the last quarter of the Great Recession—expected consumption in 2009Q3 explained 74.3% of the decline in current consumption, whereas expected consumption in 2015Q3 explained only 1.6% of the decline. Over those same horizons, the contribution of consumption uncertainty increased from 2.9% to 11.9%, which shows that it is crucial to decompose the expected value of future consumption in order to uncover the underlying sources of the fluctuations in current consumption.[7]

The middle right panel shows the effect of consumption uncertainty over the same forecast horizons shown in the left panel, except the values on the vertical axis are cumulative effects (i.e., the sum of the impact in each quarter over a given horizon). Although the effect of consumption uncertainty is small when it is conditional on expected consumption over a 1-quarter horizon, it is more significant over longer horizons that reduce the contribution of expected future consumption. Over a 24-quarter horizon, consumption uncertainty on average decreases current consumption by 0.06% and the peak effect was about 0.15% in 2009Q1, which accounted for 16.6% of the overall decline in consumption in that quarter. Those results suggest that uncertainty had a similar average effect on consumption as the VAR indicates when uncertainty is ordered last. However, it

[7] We would like to show the effect of each term as a percentage of the total change of consumption every period, but the value is uninformative when consumption near its steady state or when some of the terms have competing effects.

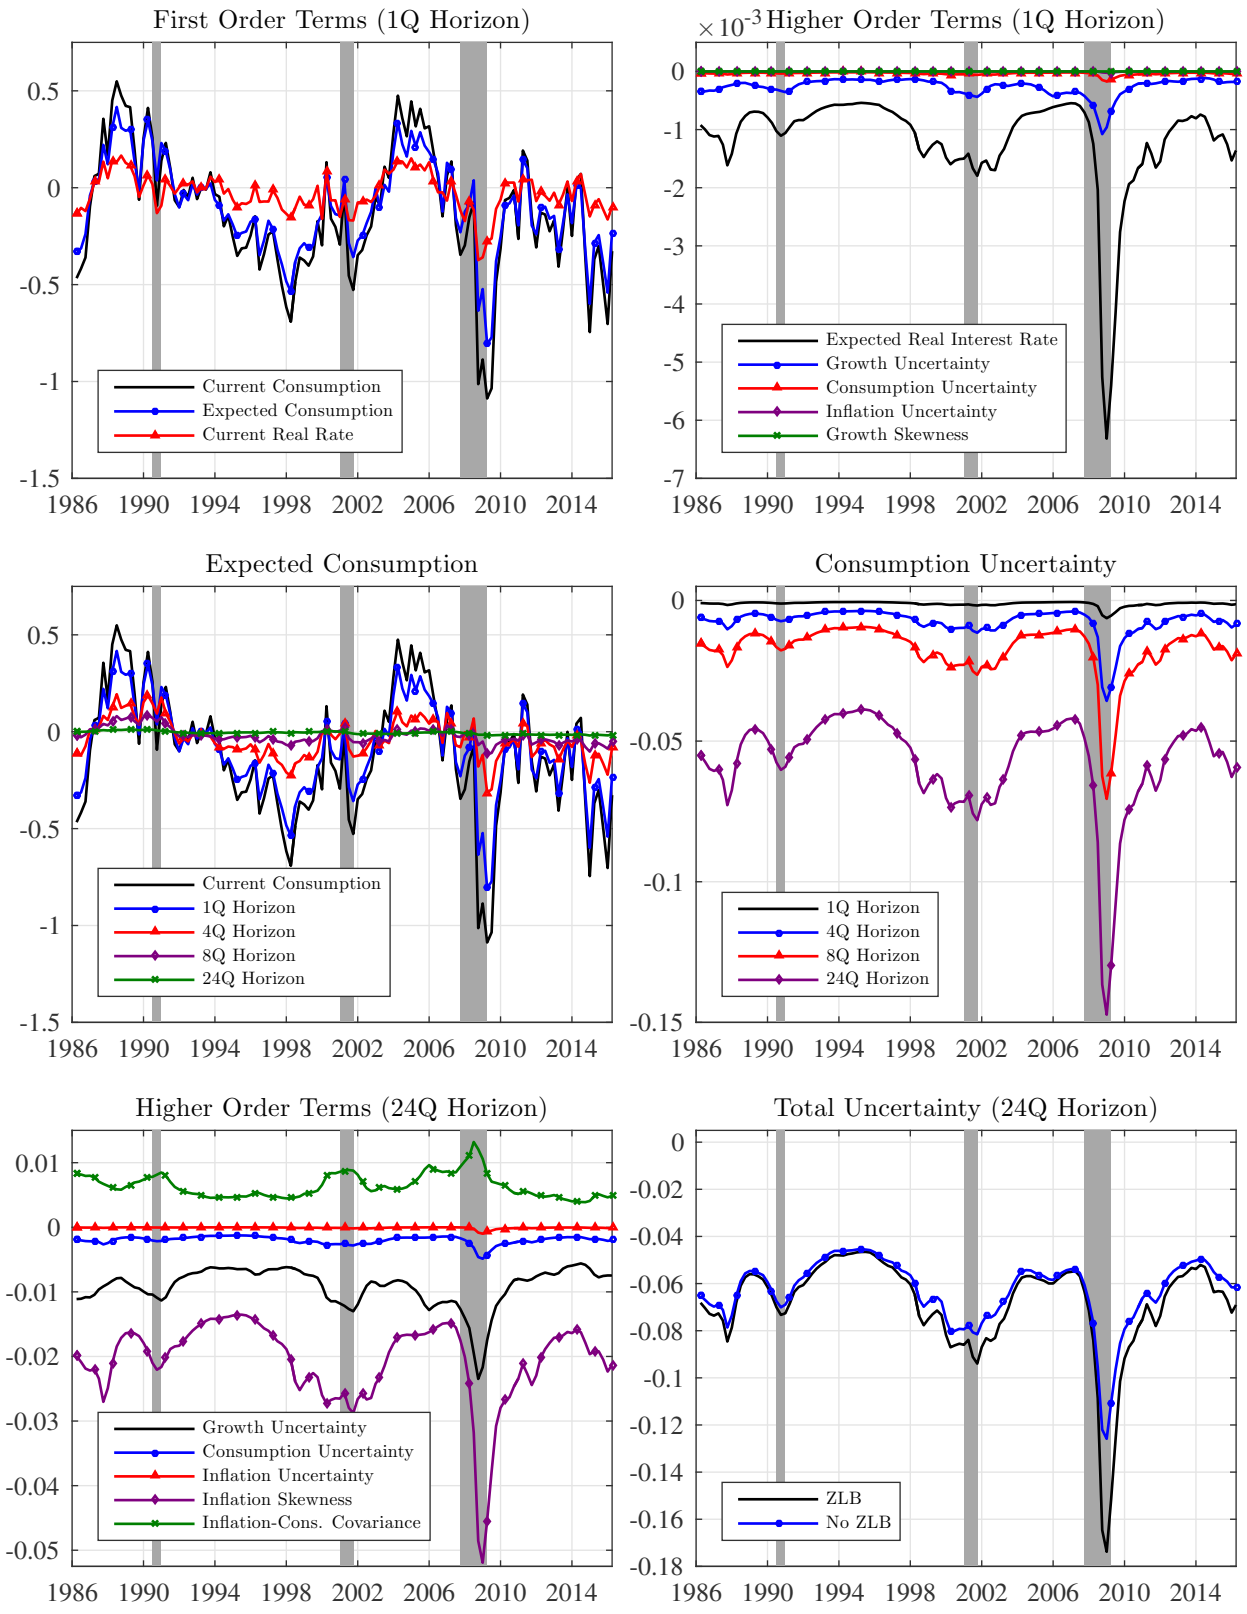

Figure 2: Filtered decomposition of the effects on current consumption. The shaded regions denote NBER recessions. The vertical axis is the contribution to the percentage point deviation of detrended consumption from its steady state.

is important to emphasize that the estimates from the VAR are based on past data (i.e., backward-looking), they only show the effect of one second moment shock with the other shocks turned off, and they are independent of the state of the economy. In contrast, the estimates from our Euler equation decomposition account for all possible realizations future shocks (i.e., forward-looking) and they are based on the combined effect of the shocks that best explain the data in each period.

The other higher order moments are shown in the bottom left panel. During the Great Recession, the peak effects of inflation uncertainty, technology growth uncertainty, and consumption skewness over a 24-quarter horizon were -0.023%, -0.005%, and -0.001%, respectively, and the average effects were an order of magnitude smaller. We do not show the effect of inflation skewness because it is near zero throughout the sample. It is not surprising that inflation uncertainty and skewness had such small effects on consumption because the Fed aggressively targeted inflation throughout our sample. However, we expected a larger effect of consumption skewness, especially during the Great Recession. The ZLB constraint creates downside risk because it prevents the Fed from responding to adverse shocks through conventional channels. Evidently, those effects on economic activity are very small when controlling for other moments. Interestingly, the inflation risk premium had the second largest adverse effect on consumption out of all the higher order terms in the decomposition and the growth risk premium had a small positive effect on economic activity.

The bottom right panel shows the total effect of uncertainty (i.e., the sum of consumption, growth, and inflation uncertainty) on consumption over a 24-quarter horizon. On average, total uncertainty reduces consumption by about 0.07% with a maximum decline of 0.17% in the Great Recession. As a counterfactual, we also plot the effect of uncertainty after removing the influence of the ZLB constraint using the solution to the unconstrained nonlinear model. The differences between the two paths show how much the ZLB constraint 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. Significant differences between the two paths occurred from 2008Q4 to 2009Q4, when the notional rate was well below zero and there was a strong expectation of staying at the ZLB. During that period, uncertainty reduced consumption by about 0.048 percentage points more (-0.174% versus -0.126%) than it would have if the Fed was not constrained by the ZLB.

The ZLB also amplified the marginal effect of total uncertainty around the time of the 2001 recession. That happened even though the annual federal funds rate exceeded 4% before the recession. Macro uncertainty increased before the recession and the filtered volatility series for technology growth and the risk premium were elevated. Households expected volatility to remain elevated since the uncertainty series are very persistent. Hence, first-moment shocks were amplified over the 24-quarter forecast horizon, making it more likely that the ZLB will bind. It became even more likely that the ZLB would bind when the federal funds rate was reduced during the recession and remained below 2%. These results demonstrate that endogenous events, particularly during recessions, can have a meaningful effect on how people respond to sharp increases in uncertainty.

5.3 IMPULSE RESPONSES The Euler equation decomposition shows consumption uncertainty had the largest effect on consumption at the end of the Great Recession when the Fed was constrained by the ZLB. Impulse responses initialized at different states of the economy are a useful way to highlight which states and shocks contributed to that result. Figure 3 plots the responses to a 2 standard deviation positive risk premium (first column), risk premium volatility (second column), growth (third column), and growth volatility (fourth column) shock. The parameters are set to their posterior means, and the simulations are initialized at two different states. As a benchmark,

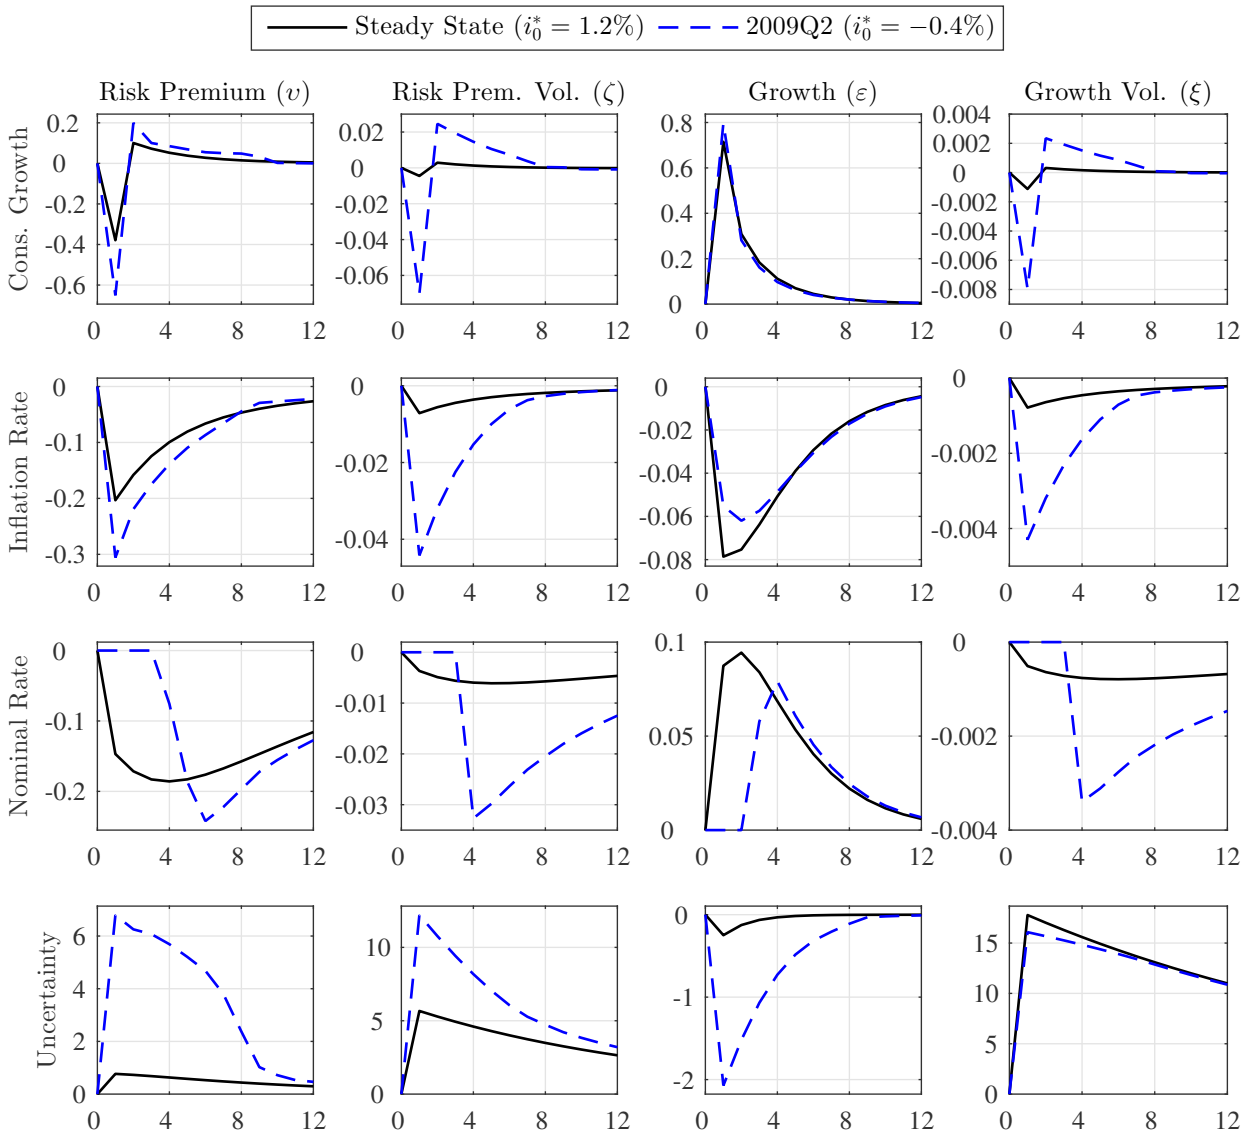

Figure 3: Impulse responses to a 2 standard deviation positive shock at and away from the ZLB. The steady-state simulation (solid line) is initialized at the stochastic steady state. The other simulation (dashed line) is initialized at the filtered state corresponding to 2009Q2 so the ZLB binds. The vertical axis is in percentage point deviations from the baseline simulation, except uncertainty which is a percent change. The horizontal axis is the time period in quarters.

the steady-state simulation (solid lines) is initialized at the stochastic steady state. We compare the benchmark responses to the responses when the notional rate is negative (dashed lines) 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 difference for uncertainty) from a counterfactual simulation without a shock in the first quarter. Uncertainty is measured by the expected volatility of the 1-quarter-ahead forecast error for consumption growth. 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 con-

sumption, which lowers consumption 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 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 consumption. The result is a larger increase in uncertainty since households expect a wider range of future realization of consumption 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 consumption 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 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 consumption 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 consumption growth and decreases inflation like a typical supply shock, so the Fed faces a tradeoff between stabilizing inflation and real GDP growth unlike with a risk premium shock. In steady state, the policy rate immediately increases since the response to the real GDP gap dominates the response to the inflation gap. In 2009Q2 the ZLB initially binds, 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 consumption growth and a smaller decline in inflation. In contrast with the risk premium shock, a positive growth shock causes uncertainty to decline because it reduces the probability that the ZLB binds next period. However, the responses are smaller in magnitude.

Growth volatility shocks generate a larger change in uncertainty than level shocks. Similar to a risk premium volatility shock, a positive growth volatility shock (fourth column) reduces consumption growth and inflation, which leads to a lower nominal rate. However, the responses differ from a risk premium volatility shock in a few ways. One, growth volatility directly affects consumption volatility. Therefore, uncertainty increases more than it does in response to risk premium volatility shocks. Two, since the central bank faces a tradeoff between stabilizing consumption and inflation in response to a growth shock, the response of uncertainty is similar in both initial states. Higher volatility directly creates more uncertainty by prolonging the expected duration at the ZLB, but it also places a higher probability on positive growth shocks that would push the notional rate into positive territory. Evidentally those two effects offset each other. Three, while the uncertainty response is roughly twice as large as the response to a risk premium volatility shock when the ZLB binds, the differences are much larger away from the ZLB. Therefore, growth volatility shocks play a much larger role in explaining the fluctuations in uncertainty when the ZLB does not bind.

5.4 SOURCES OF UNCERTAINTY The top panel of figure 4 plots the standardized macro uncertainty series (solid line) from Jurado et al. (2015) and the standardized financial uncertainty series (circles) from Ludvigson et al. (2017) that we use as observables in our estimation. Given the mean parameter estimates from our model, we are able to determine how much of the fluctuations in uncertainty are due to exogenous sources and how much are caused by the state of the economy.

The middle panel shows the decomposition of the sources of uncertainty in our model. We first plot the filtered time series of consumption growth uncertainty (solid line) to show the total amount of uncertainty across time. The measurement error equation directly links the macro uncertainty

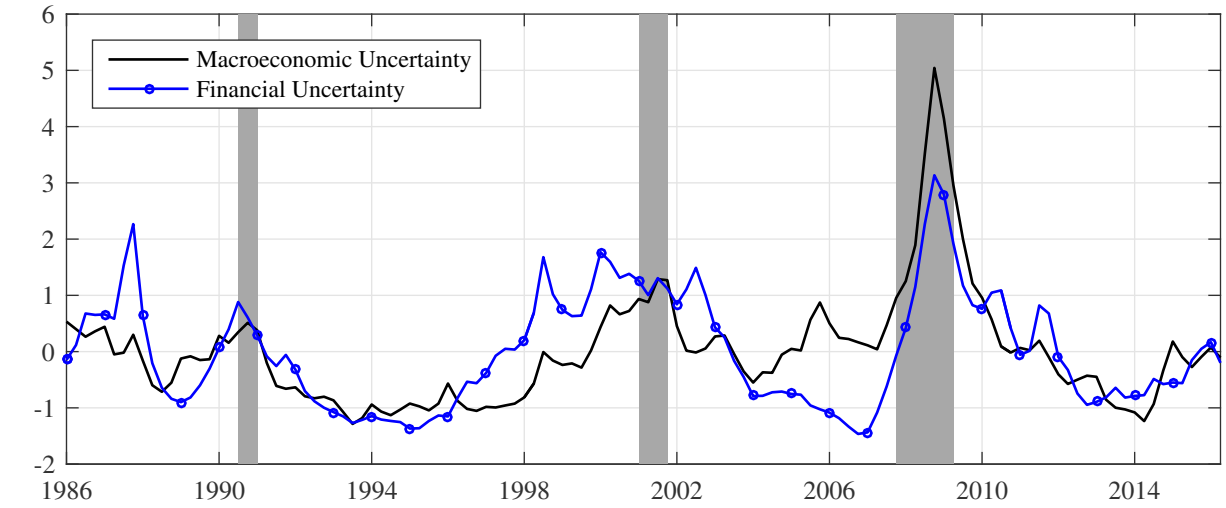

(a) Uncertainty measures in the data.

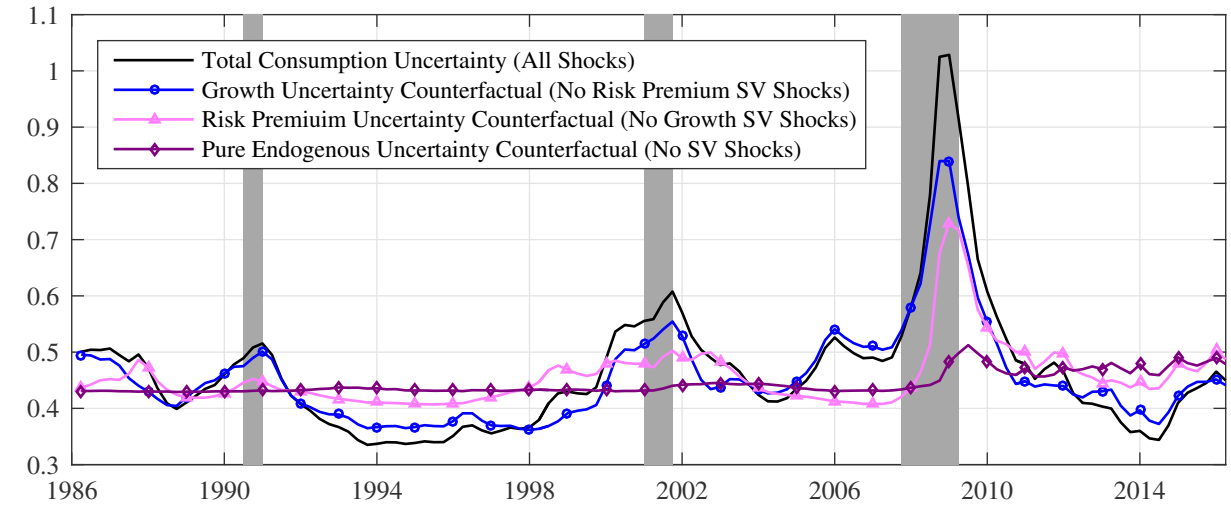

(b) Consumption uncertainty estimates from the DSGE model.

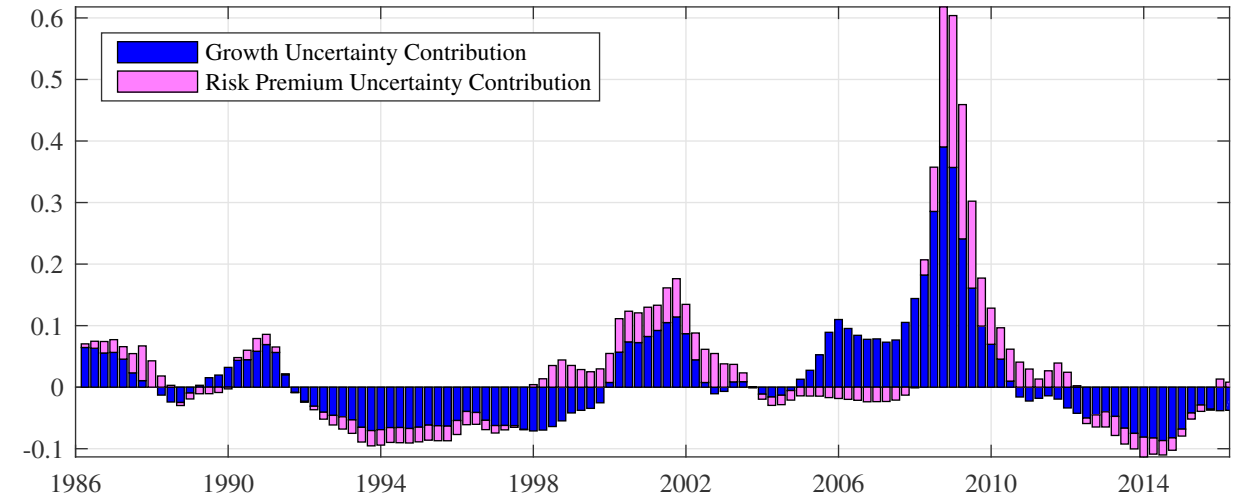

(c) Relative contribution of the exogenous sources of consumption uncertainty.

Figure 4: Time series of the uncertainty measures in the data and the sources of uncertainty in DSGE model.

series in the data to consumption uncertainty, so the filtered uncertainty series is very close to the data, except the units are consistent with the model. We then conduct counterfactual simulations that highlight the contribution of each exogenous source of uncertainty to total uncertainty. To isolate the effect of growth uncertainty, we turn off the risk premium volatility shocks (circles). Similarly, we zero out the growth volatility shocks to identify the effect of risk premium uncertainty (triangles). When we turn off both volatility shocks, the model still generates time-varying uncertainty due to the ZLB, which reflects one aspect of the endogenous uncertainty (diamonds).[8]

The decomposition of consumption uncertainty reveals a number of interesting findings. First, consumption uncertainty with (solid line) and without (circles) the risk premium volatility shocks is positively correlated. In other words, the data indicate that most of the fluctuations in uncertainty are explained by growth volatility shocks, which are independent of the state of the economy, rather that risk premium volatility shocks. Second, risk premium uncertainty (triangles), which is linked to the financial uncertainty series in the data, is an important driver of consumption uncertainty only in certain parts of our sample. For example, total consumption uncertainty begins to rise in the late 1990s due to elevated risk premium volatility and it remains elevated for about two years after the Great Recession partly due the higher persistence in risk premium volatility. Some of the fluctuations in uncertainty are also explained by time-varying endogenous uncertainty (diamonds). For most of the example, 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.

Despite some nonlinear interactions between the 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 the bottom panel represent the growth counterfactual relative to the pure endogenous uncertainty counterfactual (circles minus diamonds) and the light bars represent the risk premium counterfactual relative to the pure endogenous uncertainty counterfactual (triangles minus diamonds), which is approximately equal to total uncertainty relative to the pure endogenous uncertainty counterfactual (solid minus diamonds). The results indicate that growth uncertainty is typically the biggest contributor to total uncertainty and the two sources of exogenous uncertainty often act together. However, there are two notable exceptions. One, the model predicts that risk premium uncertainty precedes the 2001 recession. Two, growth uncertainty precedes the rise in risk premium uncertainty during the Great Recession, but the effects of risk premium uncertainty linger while the impact of growth uncertainty is negligible for a few years immediately after the recession. Interestingly, growth and risk premium shocks have nearly equal roles in explaining the past two recessions. Toward the end of the sample, total uncertainty is near its lowest point, which is primarily due to the reduced influence of growth volatility shocks.

6 THE EFFECT OF CAPITAL ACCUMULATION

In our baseline DSGE model, real GDP 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 physical 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.

[8] One drawback of our approach to decomposing the sources of uncertainty is that the endogenous amplification of first and second moment shocks is embedded in the growth and risk premium uncertainty counterfactuals, which reduces the relative importance of endogenous uncertainty. Unfortunately, there is no easy way to address this problem.

The final goods firm's problem is unchanged. Intermediate firm $f \in [0,1]$ produces a differentiated good, $y_t(f)$, according to $y_t(f) = k_{t-1}(f)^{\alpha}(z_t n_t(f))^{1-\alpha}$. It then chooses its capital and labor inputs, $n_t(f)$ and $k_{t-1}(f)$, and its price, $p_t(f)$, to maximize the same profit function as in the baseline model. In a symmetric equilibrium, the production function and optimality conditions are

$$\tilde{y}_t = (\tilde{k}_{t-1}/g_t)^{\alpha} n_t^{1-\alpha}, \tag{16}$$

$$\alpha \tilde{w}_t n_t = (1-\alpha) r_t^k (\tilde{k}_{t-1}/g_t), \tag{17}$$

$$mc_t = \tilde{w}_t^{1-\alpha} (r_t^k)^{\alpha} / ((1-\alpha)^{1-\alpha} \alpha^{\alpha}), \tag{18}$$

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 level of investment in physical capital, $x_t^g \equiv x_t/(\bar{g} x_{t-1})$ is the growth rate of investment from time $t-1$ to t relative to the balanced growth rate, $\varphi_x > 0$ scales the size of the cost to adjusting investment at a rate that differs from the balanced growth rate, and k is the capital stock, which earns a real rate of return r^k and depreciates at rate δ. In addition to the original optimality conditions, (8) and (9), there is an equation for Tobin's q and a second Euler equation, given by,

$$q_t = \beta E_t[(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma}(r_{t+1}^k + q_{t+1}(1-\delta))/g_{t+1}], \tag{19}$$

$$1 = q_t[1 - \varphi_x(\tilde{x}_t^g - 1)^2 - \varphi_x \tilde{x}_t^g (\tilde{x}_t^g - 1)] + \beta \varphi_x \bar{g} E_t[q_{t+1}(\tilde{c}_t/\tilde{c}_{t+1})^{\gamma}(\tilde{x}_{t+1}^g)^2(\tilde{x}_{t+1}^g - 1)/g_{t+1}]. \tag{20}$$

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

$$\tilde{k}_t = (1-\delta)(\tilde{k}_{t-1}/g_t) + \tilde{x}_t(1 - \varphi_x(\tilde{x}_t^g - 1)^2/2), \tag{21}$$

$$\tilde{c}_t + \tilde{x}_t = \tilde{y}_t^{gdp}. \tag{22}$$

Once again, we redefined variables that grow along the balanced growth path in terms of technology. A competitive equilibrium includes infinite sequences of quantities, $\{\tilde{c}_t, \tilde{y}_t, \tilde{y}_t^{gdp}, n_t, \tilde{x}_t, \tilde{k}_t\}_{t=0}^{\infty}$, 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, $\{s_t, g_t, \sigma_{\varepsilon,t}, \sigma_{\upsilon,t}\}_{t=0}^{\infty}$, that satisfy the detrended equilibrium system, (1), (2), (5)-(11), (13), and (16)-(22), given the initial conditions, $\{c_{-1}, i_{-1}^n, x_{-1}, k_{-1}, s_0, a_0, \nu_0, \sigma_{\varepsilon,0}, \sigma_{\upsilon,0}\}$, and the five sequences of shocks, $\{\varepsilon_t, \upsilon_t, \nu_t, \xi_t, \zeta_t\}_{t=1}^{\infty}$.

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. (2016). Although there are some differences between our model and the one in Gust et al. (2016) (e.g., their model includes sticky wages and variable capital utilization, 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.[9]

Fortunately, introducing capital does not change the consumption Euler equation we used to construct the decomposition in the baseline model. We generate policy functions for each term

[9] We decided not to add other important features, such as a marginal efficiency of investment shock or irreversible investment, because those changes could affect our parameter estimates and the accuracy of our quantitative results.

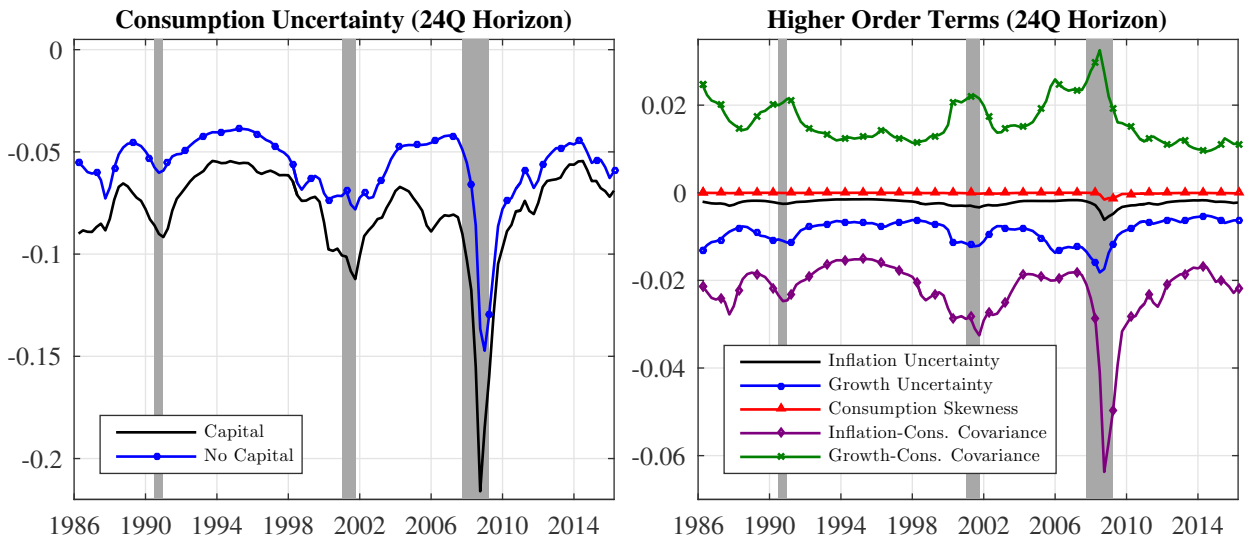

Figure 5: Filtered decomposition of the effects on current consumption. The shaded regions denote NBER recessions. The vertical axis is the contribution to the percentage point deviation of detrended consumption from its steady state.

in the decomposition in the same way as our baseline model, except we filter the data with per capita real investment growth, in addition to the five observables we previously used. Figure 5 shows the effects of the same higher-order moments shown in figure 2. The left panel compares the effects of consumption uncertainty over a 24-quarter horizon in the models with and without capital. Consumption uncertainty on average decreases current consumption by 0.08%, which is only 0.02 percentage points more than in the baseline model. The differences are more pronounced when the ZLB first binds. For example, in 2008Q4 consumption declines by 0.22% compared with only 0.14% in the baseline model, but those differences last for only a couple quarters. The right panel shows the effects of other higher order terms over the same 24-quarter horizon in the model with capital. Similar to the baseline model, these terms had smaller effects on current consumption than consumption uncertainty. When comparing with the baseline model (see the bottom left panel of figure 2), we find technology growth uncertainty had a relatively larger effect in the model with capital whereas the role of inflation uncertainty was smaller. Overall, our results indicate that the addition of capital had a relatively small impact on how higher-order moments affect the economy.

7 CONCLUSION

This paper develops a new method for quantifying the effects of macro uncertainty and its underlying sources using estimates from a nonlinear New Keynesian model. The model includes a ZLB constraint, which introduces time-varying endogenous uncertainty, and two exogenous sources of time-varying uncertainty—a volatility shock to technology growth and a volatility shock to the risk premium. The filtered time series of the terms in a third-order approximation of the Euler equation reveal that consumption uncertainty on average reduces consumption by about 0.06% and the peak effect was 0.14% during the Great Recession. Adding capital accumulation to the baseline model has a fairly small impact—the average effect increases to -0.08% and the peak effect becomes -0.22%. Other higher-order terms, such as inflation uncertainty, consumption skewness, and inflation skewness had smaller effects in our models with and without capital. Technology growth

volatility usually explains the majority of the fluctuations in macro uncertainty in our sample, but during and around the last two recessions risk premium volatility played a more significant role.

A major benefit of our methodology is its flexibility. It can be used to examine the causal effects of any type of uncertainty in a broad class of models. Although there are major computational challenges to adding complexity to the model, it would be interesting to examine additional or alternative forms of endogenous and exogenous uncertainty to determine whether there is a larger effect on consumption or whether a specific type of uncertainty plays a bigger role in explaining the changes in uncertainty. For example, one could introduce irreversible investment to add another source of endogenous uncertainty and examine the magnitude of the real options effect. It is also worthwhile to examine the effects of shocks to the volatilities of various fiscal instruments or the marginal efficiency of investment. We view these extensions as important areas for future research.

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.

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

BACHMANN, R. AND G. MOSCARINI (2012): "Business Cycles and Endogenous Uncertainty," Manuscript, RWTH Aachen University.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2013): "Measuring Economic Policy Uncertainty," Manuscript, Stanford University.

BASU, S. AND B. BUNDICK (2016): "Uncertainty Shocks in a Model of Effective Demand," FRB Kansas City Working Paper 14-15.

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

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

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2014): "Really Uncertain Business Cycles," Center for Economic Studies Working Paper 14-18.

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

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.

CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

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

DIXIT, A. K. AND J. E. STIGLITZ (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308.

FAJGELBAUM, P., M. TASCHEREAU-DUMOUCHEL, AND E. SCHAAL (2014): "Uncertainty Traps," NBER Working Paper 19973.

FERNALD, J. G. (2012): "A quarterly, Utilization-Adjusted Series on Total Factor Productivity," Federal Reserve Bank of San Francisco Working Paper 2012-19.

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. (2013): "Credit Risk and Disaster Risk," *American Economic Journal: Macroeconomics*, 5, 1–34.

——— (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 (2016): "The Empirical Implications of the Interest-Rate Lower Bound," Finance and Economics Discussion Series 2012-083.

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

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

JUSTINIANO, A. AND G. E. PRIMICERI (2008): "The Time-Varying Volatility of Macroeconomic Fluctuations," *American Economic Review*, 98, 604–41.

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.

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

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

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

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 (2016): "The Zero Lower Bound and Endogenous Uncertainty," *Economic Journal*, forthcoming.

PLANTE, M. AND N. TRAUM (2012): "Time-varying Oil Price Volatility and Macroeconomic Aggregates," Federal Reserve Bank of Dallas Working Paper 1201.

RICHTER, A. W. AND N. A. THROCKMORTON (2016): "Is Rotemberg Pricing Justified by Macro Data?" *Economics Letters*, 149, 44–48.

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.

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.

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.

A DATA SOURCES

We drew from the following data sources to estimate our VAR and DSGE models, which are available on the Federal Reserve Economic Database (FRED) unless specified otherwise:

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. **Personal Consumption Expenditures, Nondurable Goods**: Monthly, billions of dollars, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 2.8.5. (FRED ID: PCEND).
5. **Personal Consumption Expenditures, Services**: Monthly, billions of dollars, seasonally adjusted. Source: Bureau of Economic Analysis, National Income and Product Accounts, Table 2.8.5. (FRED ID: PCES).
6. **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).
7. **Average Hourly Earnings**: Monthly, production and nonsupervisory employees, dollars per hour, seasonally adjusted. Source: Bureau of Labor Statistics (FRED ID: AHETPI).
8. **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)

9. **Utilization Adjusted Total Factor Productivity**: Quarterly. Source: Fernald (2012), measure dtfp_util. Data available from http://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/
10. **Effective Federal Funds Rate:** Daily. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates, H.15. (FRED ID: FEDFUNDS).
11. **S&P500 Index**: Daily. Source: Yahoo Finance. Data available from https://finance.yahoo.com/quote%5EGSPC/history?p=%5EGSPC
12. **Civilian Noninstitutional Population**: Monthly. Source: U.S. Bureau of Labour Statistics, Current Population Survey. (FRED ID: CNP16OV).
13. **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:

14. **Per Capita Real GDP**: $1{,}000{,}000 \times$Real GDP/Population.
15. **Real PCE, Nondurable Goods**: Average PCE Nondurables in $2009 \times$(PCE Nondurables Quantity Index/100). Quantity Index FRED ID: DNDGRA3M086SBEA.
16. **Real PCE, Services**: Average PCE Nondurables in $2009 \times$(PCE Services Quantity Index/100). Quantity Index FRED ID: DSERRA3M086SBEA.
17. **Per Capita Real PCE**: $1{,}000{,}000 \times$(Real PCE Nondurables+Real PCE Services)/Population.
18. **Real Wage**: $100 \times$Average Hourly Earnings/Price Index.
19. **Real Investment**: Average FPI in $2009 \times$(FPI Quantity Index/100). Quantity Index FRED ID: A007RA3Q086SBEA.
20. **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, 2, 14, 17, 6, 18, and 8-11. The observables used to estimate our baseline DSGE model include series 14, 6, 10, 1, and 2. When we filter the data using the model with capital, we add series 20 as an observeable.

B ROBUSTNESS OF VAR RESULTS

In Section 2, our VAR results are based on the same sample period we use to estimate our DSGE model, so we can draw comparisons between the two sets of estimates. Figure 6 plots the same impulse responses as figure 1, except we use data from 1966Q2 to 2016Q2, so our sample period is similar to what others have used in the literature. The BIC still prefers one lag in the VAR, but we report the results based on two lags, since that specification is more common in the literature.

The qualitative results are robust to using the longer sample. The responses of consumption to a one standard deviation positive shock to either macro or financial uncertainty are smaller when the uncertainty series are ordered first or the ZLB period is excluded from the sample. The peak responses to a financial uncertainty shock are slightly larger in the longer sample, whereas the responses to a macro uncertainty shock are typically smaller. Notably, the response of consumption to a macro uncertainty shock when the uncertainty series are ordered last and the ZLB period is excluded from the sample have virtually the same peak effect regardless of when the sample begins.

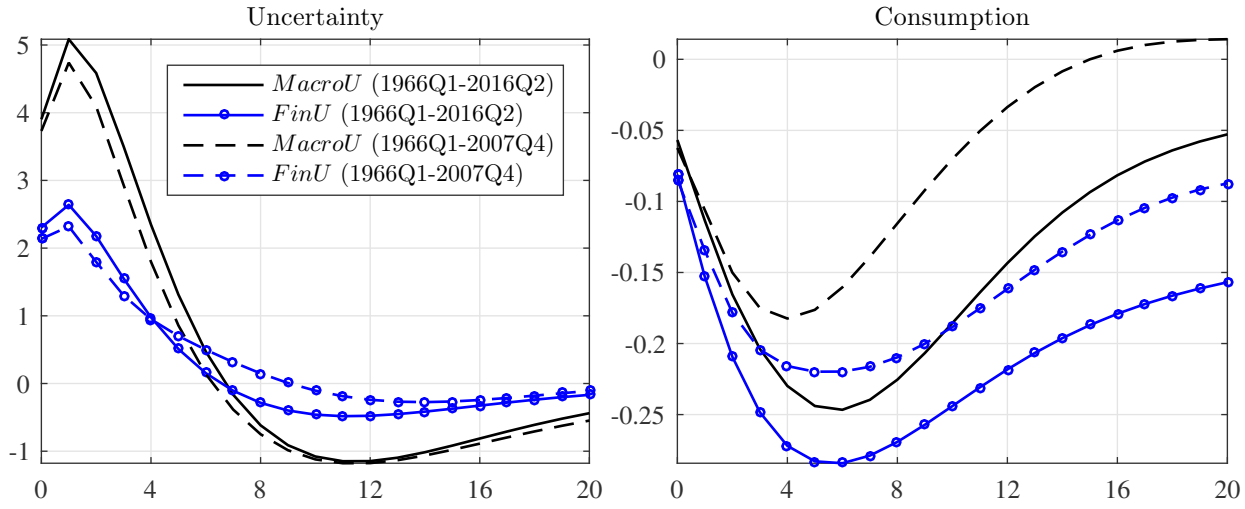

(a) VAR model where the uncertainty series are ordered first.

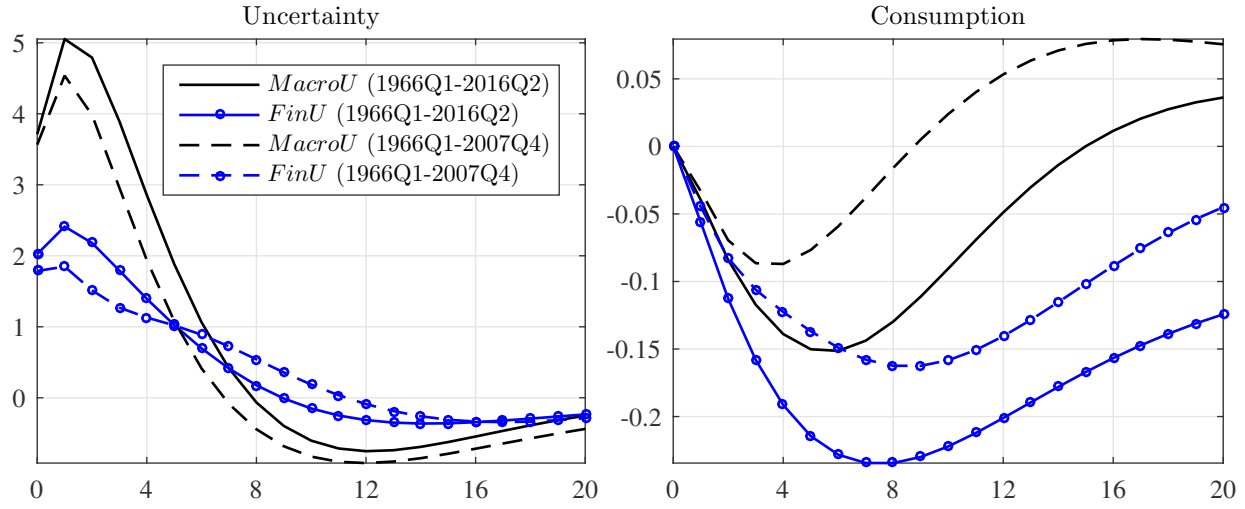

(b) VAR model where the uncertainty series are ordered last.

Figure 6: Impulse responses to a one standard deviation increase in macro and financial uncertainty.

C SOLUTION METHOD

C.1 BASELINE MODEL

We begin by writing the equilibrium system of equations compactly as

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

where f is a vector-valued function, $\mathbf{v} = (s, g, \sigma_\varepsilon, \sigma_\upsilon, \tilde{c}, \tilde{y}, \tilde{y}^{gdp}, n, \tilde{w}, i, i^n, \pi^{gap})$, and $\Omega = \{S, P, \mathbf{z}\}$ is the information set, which contains the structural model, S, its parameters, P, and the state vector, $\mathbf{z}_t = (\nu_t, \log(\sigma_{\varepsilon,t}), \log(\sigma_{\upsilon,t}), g_t, s_t, mp_{t-1})$. Since i^n_{t-1} and $\tilde{y}^{gdp}_{t-1}$ only appear in the policy rule, we eliminate a state variable by defining $mp_{t-1} = (i^n_{t-1})^{\rho_i}(\tilde{y}^{gdp}_{t-1})^{\phi_y(\rho_i-1)}$ and rewriting the rule as

$$i^n_t = mp_{t-1}(\bar{\imath}(\pi^{gap}_t)^{\phi_\pi}(g_t\tilde{y}^{gdp}_t/\bar{g})^{\phi_y})^{1-\rho_i}\exp(\sigma_\nu\nu_t).$$

There are many ways to discretize the normally-distributed and autoregressive exogenous state variables, ν_t, $\log(\sigma_{\varepsilon,t})$, and $\log(\sigma_{\upsilon,t})$. We follow 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 steady state, which are wide enough to contain the filtered state variables given the posterior draws. We discretize the state variables into $(4, 9, 7, 7, 7, 7)$ points respectively, such that they are evenly spaced in each dimension. Therefore, 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 is also used to obtain M integration nodes with weights, $\{\phi(m)\}_{m=1}^{M}$, that correspond to the shocks, $\{\nu_{t+1}(m), \log(\sigma_{\varepsilon,t+1})(m), \log(\sigma_{\upsilon,t+1})(m), \varepsilon_{t+1}(m), \upsilon_{t+1}(m)\}_{m=1}^{M}$. We use the same number of points, $(4, 9, 7, 7, 7)$, as the respective state variables, so $M = 12{,}348$. The processes for g_{t+1} and s_{t+1} do not have a standard autoregressive form because the standard deviations of the shocks are time-varying. Therefore, we chose not use the Rouwenhorst method to discretize the processes for g and s. Instead, the first moment shocks and log volatility processes are discretized separately with the Rouwenhorst method, so $\tilde{c}_{t+1}(m)$ and $\pi_{t+1}^{gap}(m)$ are interpolated at realizations of $g_{t+1}(m)$ and $s_{t+1}(m)$ that can occur in between the nodes in the state space.

The following steps outline our policy function iteration algorithm:

1. Obtain initial conjectures for $\tilde{\mathbf{c}}_0$ and $\boldsymbol{\pi}_0^{gap}$ from the log-linear model without the ZLB imposed using Sims's (2002) `gensys` algorithm and map it to the discretized state space.

2. For iteration j, implement the following steps with the ZLB imposed for $d \in \{1, \ldots, D\}$:

   (a) Solve for $\{\tilde{y}_t, \tilde{y}_t^{gdp}, i_t^n, i_t, \tilde{w}_t, mp_t\}$ given $\tilde{c}_t = \tilde{\mathbf{c}}_{j-1}(d)$, $\pi_t^{gap} = \boldsymbol{\pi}_{j-1}^{gap}(d)$, 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 vector, $\mathbf{z}_{t+1}(m)$, 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 rest of $\mathbf{v}_{t+1}(m)$ and compute

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

   (d) Use Chris Sims' `csolve` to find $\tilde{c}_t$ and π_t^{gap} that satisfy $\mathbb{E}[f(\cdot)|\Omega_t(d)] = 0$.

3. Using the argument of `csolve` on iteration j as an initial conjecture for iteration $j + 1$, 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 7 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 $(\nu_t, \log(\sigma_{\varepsilon,t}), \log(\sigma_{\upsilon,t}))$ dimensions since they are discretized 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.

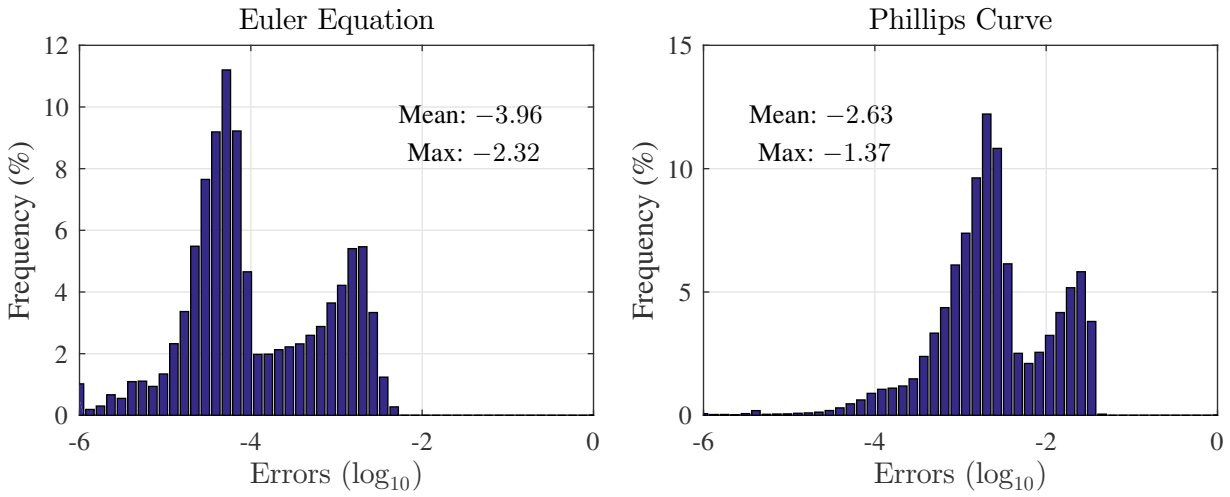

Figure 7: Distribution of Euler equation and Phillips curve errors in base 10 logarithms

C.2 CAPITAL MODEL We solve the model with capital in the same way we solve the baseline model. 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, $\theta \equiv \{\gamma, \varphi, \phi_\pi, \phi_y, \bar{g}, \bar{\pi}, \rho_i, \rho_g, \rho_s, \rho_{\sigma_\varepsilon}, \rho_{\sigma_\upsilon}, \sigma_\nu, \sigma_\varepsilon, \sigma_\upsilon, \sigma_\xi, \sigma_\zeta\}$.
2. The vector of N_x observables consists of per capita real GDP, $RGDP/CNP$, the GDP deflator, DEF, the average federal funds rate, FFR, the macro uncertainty series in Jurado et al. (2015), $MacroU$, and the financial uncertainty series in Ludvigson et al. (2017), $FinU$, from 1986Q1 to 2016Q2. Therefore, $N_x = 5$ and the row vector of observables is given by

$$\hat{\mathbf{x}}_t^{data} \equiv \begin{bmatrix} \log(RGDP_t/CNP_t) - \log(RGDP_{t-1}/CNP_{t-1}) \\ \log(DEF_t/DEF_{t-1}) \\ \log(1 + FFR_t/100)/4 \\ (MacroU_t - \mu_{MacroU})/\sigma_{MacroU} \\ (FinU_t - \mu_{FinU})/\sigma_{FinU} \end{bmatrix}^T,$$

where μ and σ denote mean and standard deviation across time 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 $\log(RI_t/CNP_t) - \log(RI_{t-1}/CNP_{t-1})$.

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:

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

      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 quintile. 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

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

      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 The following steps outline the filter:

1. Initialize the filter by drawing $\mathbf{e}_{t,p} = \{\nu_{t,p}, \varepsilon_{t,p}, \upsilon_{t,p}, \xi_{t,p}, \zeta_{t,p}\}_{t=-24}^{0}$ 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, $\mathbf{z}_{0,p}$, which is a draw from the ergodic distribution. We set $N_p = 40{,}000$.
2. For all $p \in \{1, \ldots, N_p\}$ apply the following steps:
   (a) Draw a vector of shocks from an adapted distribution, $\mathbf{e}_{t,p} \sim \mathbb{N}(\bar{\mathbf{e}}_t, I)$, where $\bar{\mathbf{e}}_t$ is chosen to maximize $p(\mu_t|\mathbf{z}_t)p(\mathbf{z}_t|\mathbf{z}_{t-1})$ and $\mathbf{z}_{t-1} = \sum_{p=1}^{N_p} \mathbf{z}_{t-1,p}/N_p$ is the state vector.
      i. Given $\mathbf{z}_{t-1}$ and a guess for $\bar{\mathbf{e}}_t$, obtain $\mathbf{z}_t$, and the endogenous variables, $\mathbf{w}_t$.
      ii. Transform the predictions for real GDP ($\tilde{y}^{gdp}$), inflation (π), the policy rate (i), consumption growth uncertainty, and risk premium uncertainty according to $\hat{\mathbf{x}}_t^{model} = \left[\log(g_t\tilde{y}_t^{gdp}/\tilde{y}_{t-1}^{gdp}), \log(\pi_t), \log(i_t), (U_{c^g,t} - \mu_{U_{c^g}})/\sigma_{U_{c^g}}, (U_{s,t} - \mu_{U_s})/\sigma_{U_s}\right]$. When we add capital to the baseline model, $\hat{\mathbf{x}}_t^{model}$ also includes $\log(g_t\tilde{x}_t/\tilde{x}_{t-1})$.
      iii. Calculate the difference between the model predictions and the data, $\mu_t = \hat{\mathbf{x}}_t^{model} - \hat{\mathbf{x}}_t^{data}$, which is assumed to be multivariate normally distributed with density:
      $$p(\mu_t|\mathbf{z}_t) = (2\pi)^{-3/2}|H|^{-1/2}\exp(-\mu_t' H^{-1}\mu_t/2),$$
      where $H \equiv \text{diag}(\sigma_{me,c^g}^2, \sigma_{me,\pi}^2, \sigma_{me,i}^2, \sigma_{me,macrou}^2, \sigma_{me,finu}^2)$ is the measurement error covariance matrix. H also includes $\sigma_{me,x}^2$ in the model with capital.
      iv. The probability of observing the current state, $\mathbf{z}_t$, given $\mathbf{z}_{t-1}$, is given by
      $$p(\mathbf{z}_t|\mathbf{z}_{t-1}) = (2\pi)^{-3/2}\exp(-\bar{\mathbf{e}}_t'\bar{\mathbf{e}}_t/2).$$
      v. Maximize $p(\mu_t|\mathbf{z}_t)p(\mathbf{z}_t|\mathbf{z}_{t-1}) \propto \exp(-\mu_t' H^{-1}\mu_t/2)\exp(-\bar{\mathbf{e}}_t'\bar{\mathbf{e}}_t/2)$ by solving for the optimal $\bar{\mathbf{e}}_t$. We converted MATLAB's `fminsearch` routine to Fortran.
   (b) Obtain $\mathbf{z}_{t,p}$, and the vector of endogenous variables, $\mathbf{w}_{t,p}$, given $\mathbf{z}_{t-1,p}$ and $\mathbf{e}_{t,p}$.
   (c) Calculate, $\mu_{t,p} = \hat{\mathbf{x}}_{t,p}^{model} - \hat{\mathbf{x}}_t^{data}$. The unnormalized weight on particle p is given by
   $$\omega_{t,p} = \frac{p(\mu_t|\mathbf{z}_{t,p})p(\mathbf{z}_{t,p}|\mathbf{z}_{t-1,p})}{g(\mathbf{z}_{t,p}|\mathbf{z}_{t-1,p}, \hat{\mathbf{x}}_t^{data})} \propto \frac{\exp(-\mu_{t,p}' H^{-1}\mu_{t,p}/2)\exp(-\mathbf{e}_{t,p}'\mathbf{e}_{t,p}/2)}{\exp(-(\mathbf{e}_{t,p} - \bar{\mathbf{e}}_t)'(\mathbf{e}_{t,p} - \bar{\mathbf{e}}_t)/2)}.$$
   If there was no adaptation, then $\bar{\mathbf{e}}_t = 0$ and $\omega_{t,p} = p(\mu_t|\mathbf{z}_{t,p})$, as it is in a basic filter. The contribution to the model's likelihood in period t is then $\ell_t^{model} = \sum_{p=1}^{N_p} \omega_{t,p}/N_p$.
   (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 $\{\mathbf{z}_{t,p}\}_{p=1}^{N_p}$ accordingly.
3. Apply step 2 for all $t \in \{1, \ldots, T\}$. The log-likelihood is then $\log \ell^{model} = \sum_{t=1}^{T} \log \ell_t^{model}$.

E DSGE MODEL ESTIMATION DIAGNOSTICS

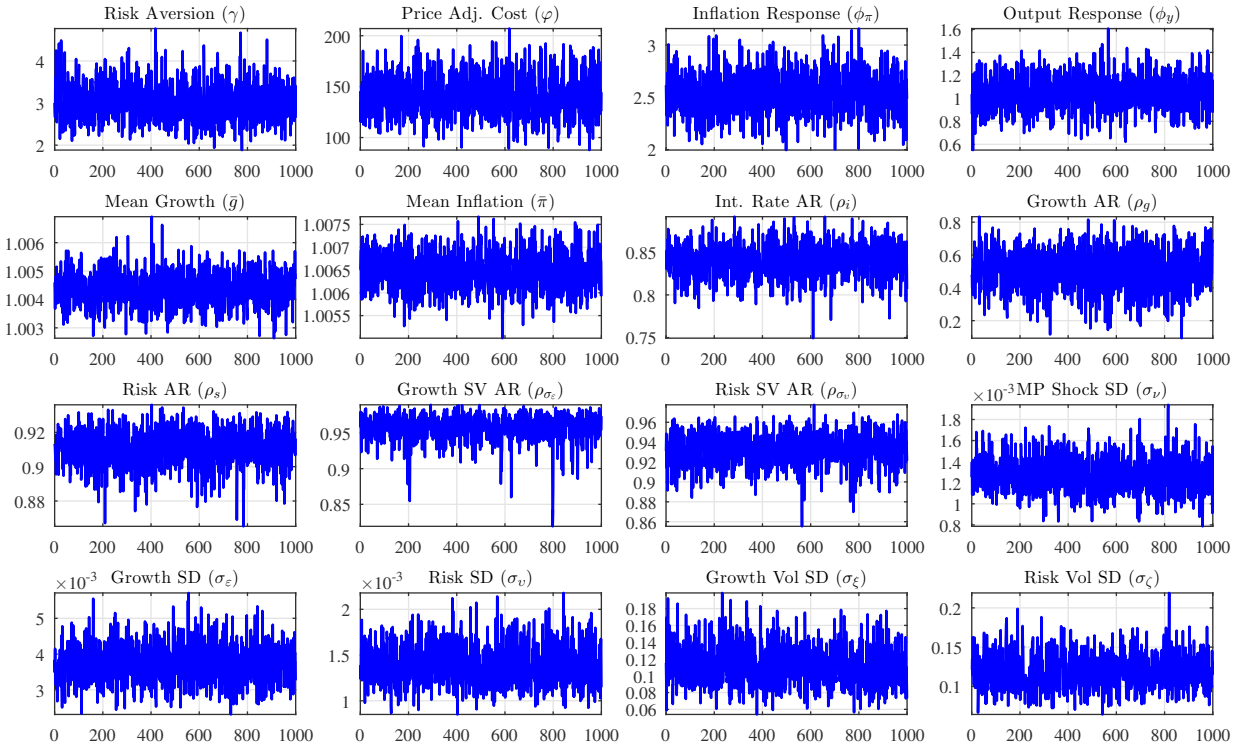

Figure 8: Trace plots. We obtained 100,000 draws from each posterior distribution and kept every 100th draw.

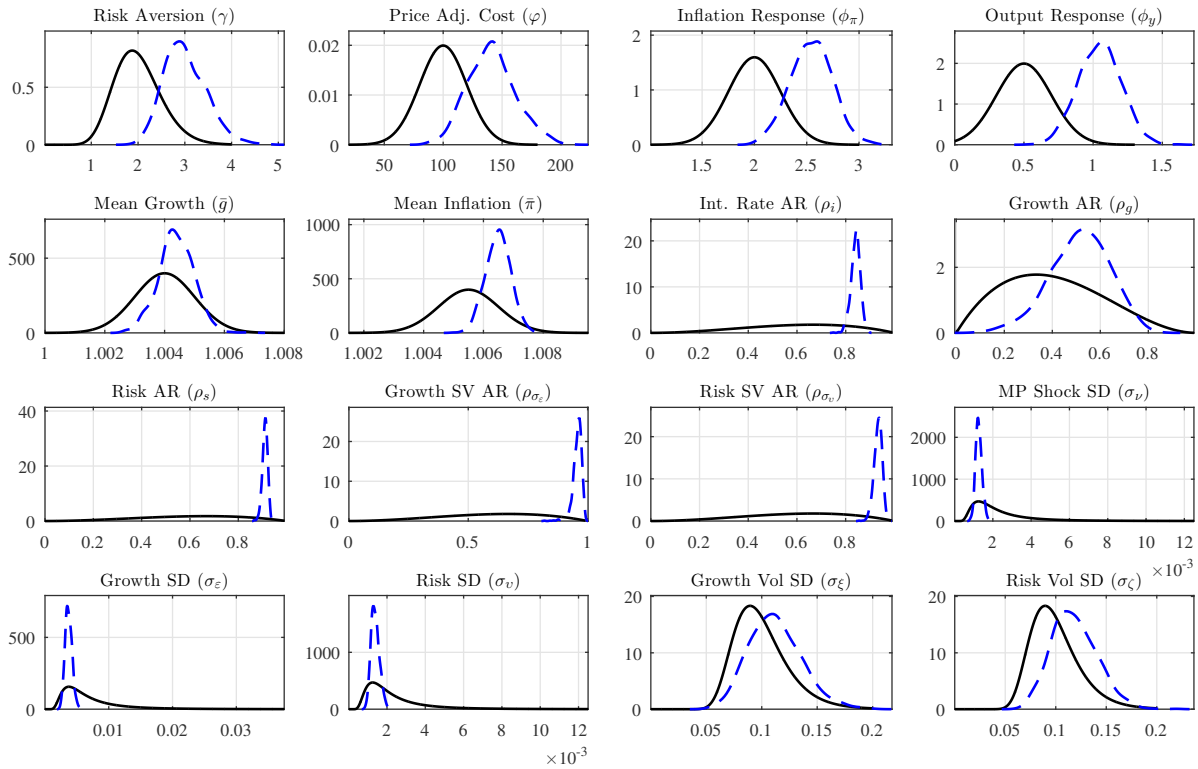

Figure 9: Prior (solid lines) and posterior kernel (dashed lines) densities of the estimated parameters.

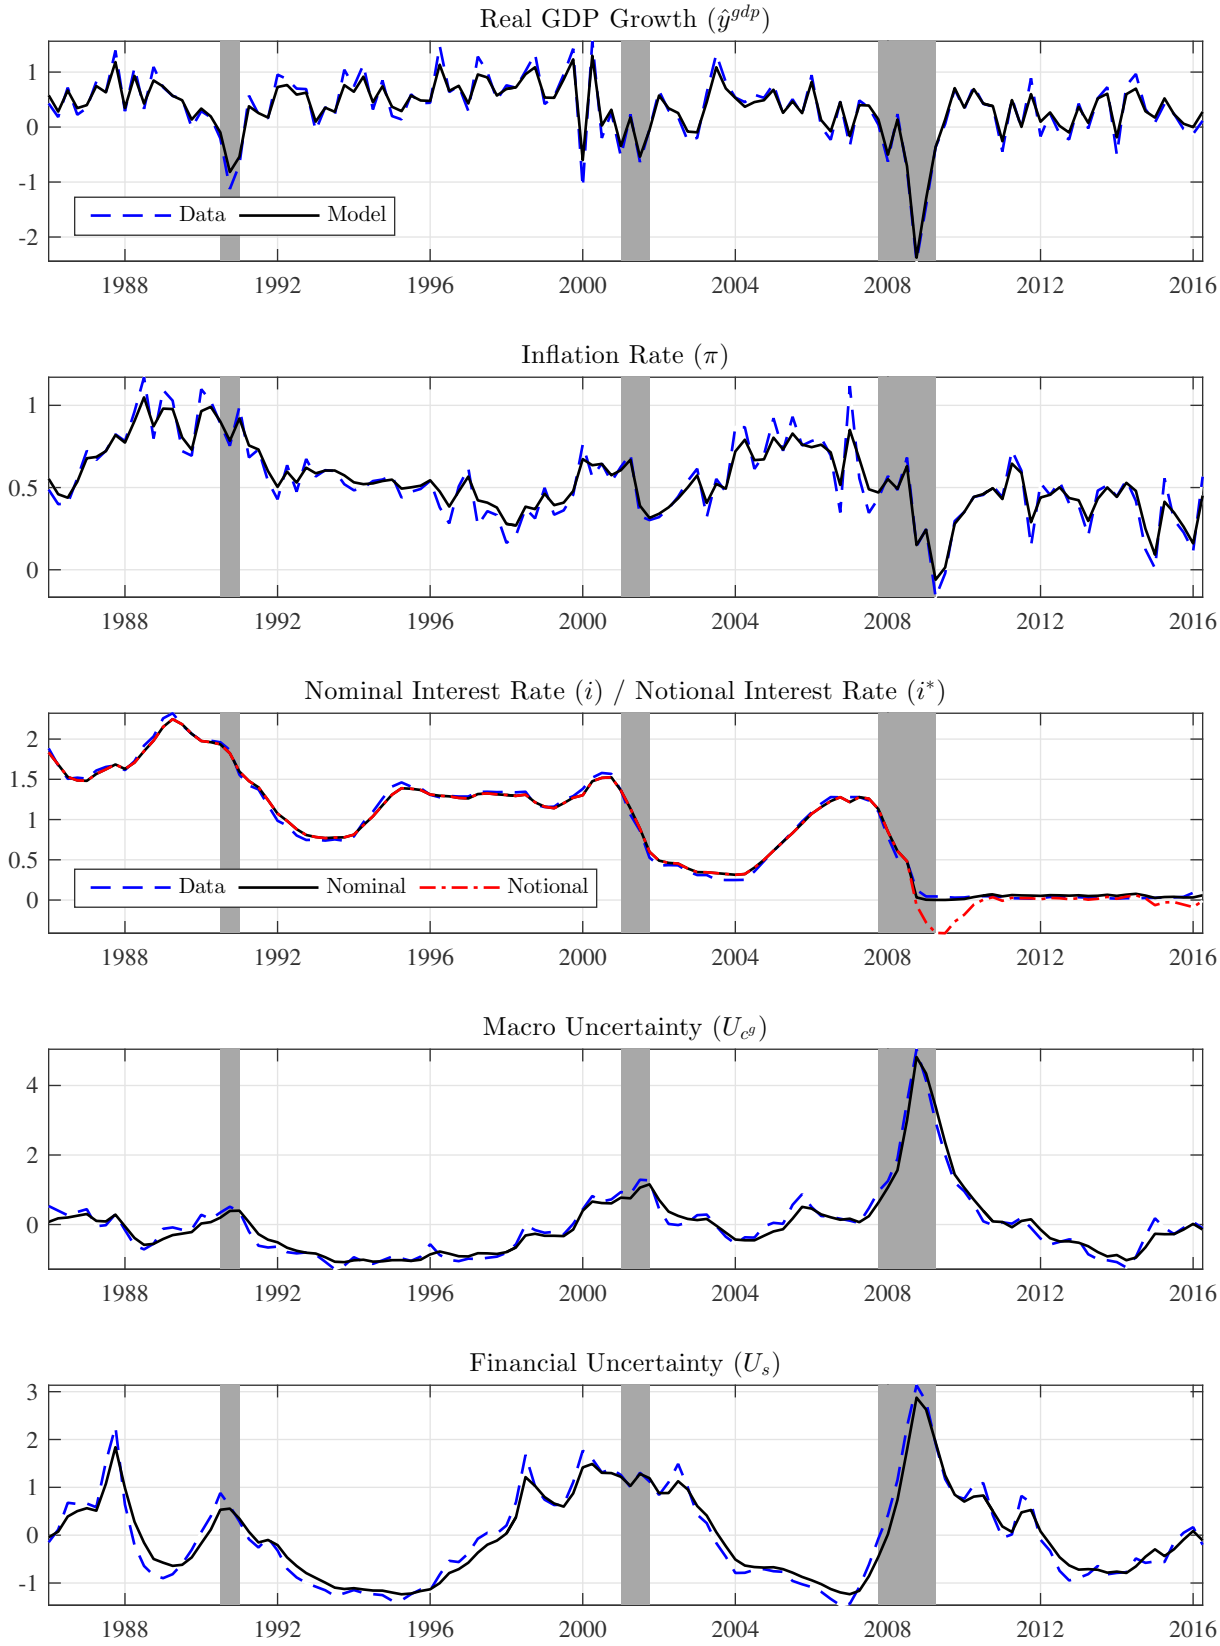

Figure 10: Time paths of the data (dashed line) and the median filtered series from the model (solid line).

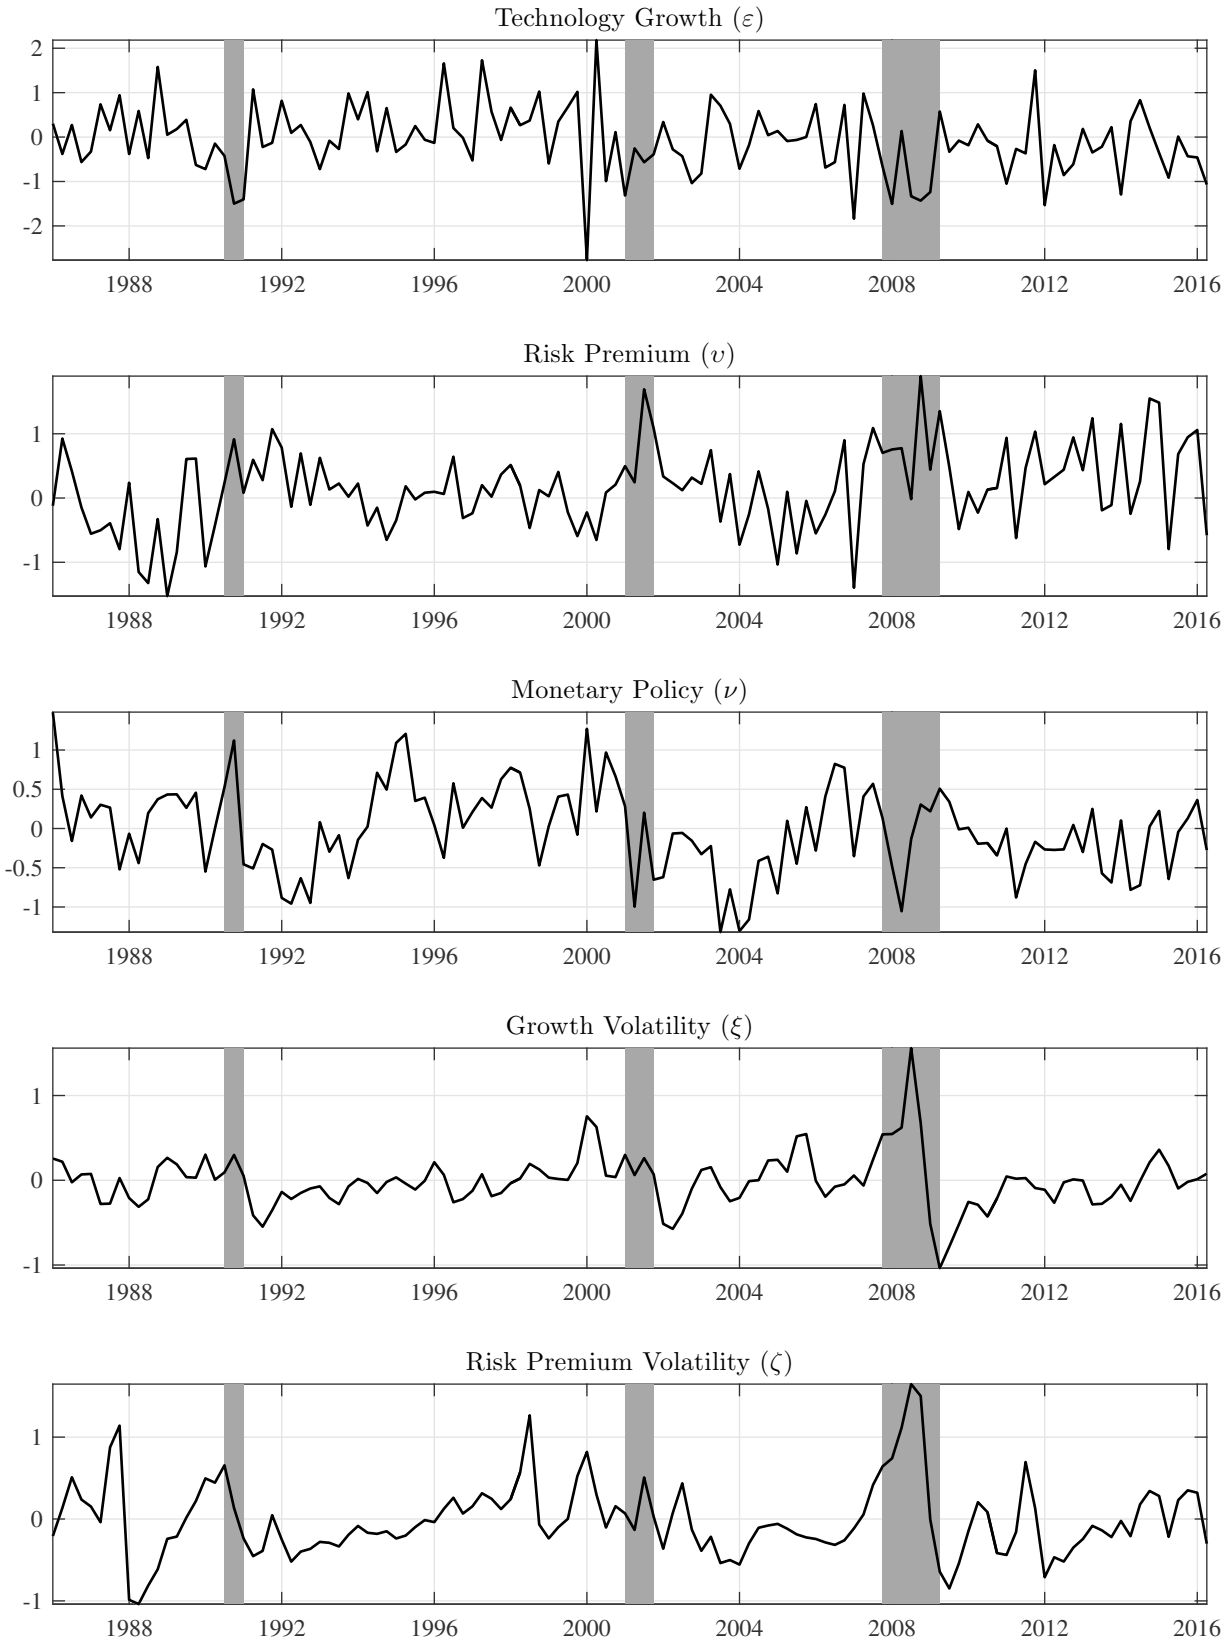

Figure 11: Median paths of the estimated shocks normalized by their respective posterior mean standard deviation.

	Real GDP Growth ($\hat{y}_t^{gdp}$)		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		
	$(\hat{y}_t^{gdp}, \hat{y}_{t-1}^{gdp})$	(π_t, π_{t-1})	(i_t, i_{t-1})	$(\hat{y}_t^{gdp}, \pi_t)$	$(\hat{y}_t^{gdp}, i_t)$	(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.

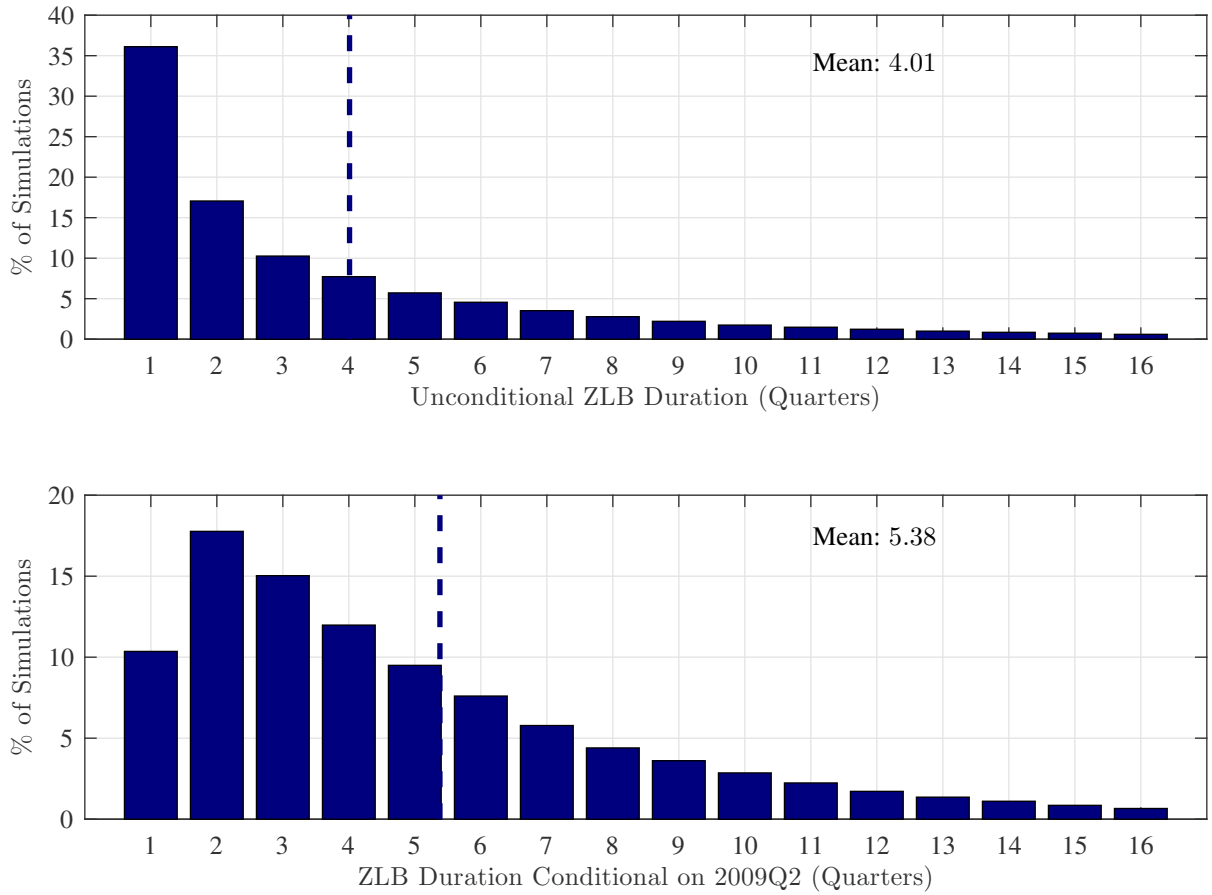

Figure 12: Distribution of ZLB events in each model. The vertical dashed line represents the expected ZLB duration.