

A New Way to Quantify the Effect of Uncertainty*

Alexander W. Richter

Nathaniel A. Throckmorton

First Draft: May 4, 2017

This Draft: February 14, 2018

ABSTRACT

This paper develops a new way to quantify the effect of uncertainty and other higher-order moments. First, we estimate a nonlinear model using Bayesian methods with data on uncertainty, in addition to common macro time series. This key step disciplines the model and allows us to generate data-driven policy functions for any higher-order moment. Second, we use the Euler equation to analytically decompose consumption into several terms—expected consumption, the *ex-ante* real interest rate and the *ex-ante* variance and skewness of future consumption, technology growth, and inflation—and then use the policy functions to filter the data and generate a time series for the effect of each term. We apply our method to a familiar New Keynesian model with a zero lower bound constraint on the nominal interest rate and two stochastic volatility shocks, but it is adaptable to any dynamic model. Over a 1-quarter horizon, uncertainty has a very small effect on consumption, similar to the volatility shocks in our model. Over horizons that remove the influence of expected consumption, the effect of uncertainty is an order of magnitude larger. Other higher-order moments have much smaller effects.

Keywords: Endogenous Uncertainty; Stochastic Volatility; Particle Filter; Zero Lower Bound

JEL Classifications: C11; D81; E32; E58

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

1 INTRODUCTION

There is widespread agreement that uncertainty decreases economic activity. The debate rests on whether the effect is quantitatively significant, which is difficult to determine for two reasons. One, uncertainty is unobserved, so there is disagreement on what constitutes a good measure. Until recently, the literature has relied on proxies for uncertainty, such as realized or implied volatility, indexes based on keywords in print or online media, and survey-based forecast dispersion, which are often weakly correlated with each other and loosely connected with the definition of uncertainty. Two, uncertainty is endogenous. Not only can uncertainty affect economic activity, as intuition suggests, what is happening in the economy can also affect uncertainty. A few of the mechanisms emphasized in the literature include financial frictions and constraints that create an adverse feedback loop between net worth and asset prices [Brunnermeier and Sannikov (2014)], incomplete information that endogenously gives rise to pessimism during recessions [Fajgelbaum et al. (2017); Saijo (2017); Van Nieuwerburgh and Veldkamp (2006)], and a zero lower bound (ZLB) constraint on the short-term nominal interest rate that restricts a central bank’s ability to stabilize the economy [Plante et al. (2018)]. As a consequence, it is difficult to quantify the causal effect of uncertainty.

There are two common ways to examine the effects of uncertainty in the literature. Empirically, the literature typically adds a measure of uncertainty to a set of variables in a vector autoregression (VAR) and then computes impulse responses using a recursive identification scheme.¹ While that approach is easy to implement, the results depend on where uncertainty is ordered in the list of variables. If the uncertainty series is ordered first, then the state of the economy has no contemporaneous effect on the impulse responses. If it is ordered last, then the uncertainty shock has no economic effect on impact. Therefore, the modeler must specify whether uncertainty is exogenous or endogenous. The challenge is even greater when trying to account for multiple sources of uncertainty. In a theoretical model, the most common way to determine the effects of uncertainty is with impulse responses to stochastic volatility shocks (i.e., time-varying variances). While that is also a fine approach, it ignores the fact that all dynamic models possess intrinsic sources of uncertainty.

This paper develops a new way to quantify the effect of uncertainty and other higher-order moments. First, we estimate a nonlinear model with Bayesian methods using the uncertainty series from Jurado et al. (2015) and Ludvigson et al. (2017), in addition to common macro time series. We chose those series over other popular measures of uncertainty because they are constructed with a rich set of macro and financial variables and are based on the *ex-ante* variance of a given variable—the same statistic we use to measure uncertainty—which focuses on predictability instead of *ex-post* variability. This key step disciplines the parameters, shocks, and predictions of the model and allows us to generate data-driven policy functions for any higher-order moment. Second, we use the consumption Euler equation to analytically decompose current consumption into several terms—expected future consumption, the *ex-ante* real interest rate and the *ex-ante* variance and skewness of future consumption, technology growth, and inflation—and then use the policy functions to filter the data at each posterior draw and generate a time series for the effect of each term.

Our approach improves on previous methods in four important ways. One, it directly links the

¹See Alexopoulos and Cohen (2009), Bachmann et al. (2013), Basu and Bundick (2017), Bekaert et al. (2013), Bloom (2009), Jurado et al. (2015), and Leduc and Liu (2016). An 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 changes in real activity. They find financial uncertainty causes sharp declines in real activity, but not macro uncertainty. Negative shocks to real activity increase macro uncertainty but have little effect on financial uncertainty.

measures of uncertainty in our model—second moments—to equivalent measures in the data with likelihood based methods, whereas previous work relied on first moments such as real activity and interest rates. Two, it captures the effect of exogenous sources of uncertainty as well as the uncertainty that naturally arises in dynamic macro models, which we refer to as endogenous uncertainty. Three, it quantifies not only the effects of uncertainty but also other higher-order moments, such as the skewness of consumption and the covariance between consumption and inflation, which have received less attention in the literature. Four, it accounts for the *ex-ante* effects of uncertainty over horizons beyond one quarter by recursively decomposing expected consumption into expected future real interest rates and higher-order moments. To summarize, our decomposition quantifies the overall effect of uncertainty and other higher-order moments in each period by accounting for the combination of first and second moment shocks that best explain both macro and uncertainty data.²

We apply our new methodology to a familiar New Keynesian model with an occasionally binding ZLB constraint on the nominal interest rate and two stochastic volatility shocks—one to the risk premium on a 1-period nominal bond and the other to the growth rate of technology. The Fed cut its policy rate to its ZLB in December 2008 and uncertainty about real GDP increased at the same time. Plante et al. (2018) find that uncertainty about future real GDP growth increased because the ZLB restricted the Fed’s ability to stabilize the economy. Thus, the ZLB endogenously generates uncertainty that depends on how severely the Fed is constrained. The stochastic volatility shocks generate time-varying exogenous uncertainty. The risk premium volatility shock introduces financial or demand uncertainty, while the technology growth volatility shock adds supply uncertainty.

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

Over a 1-quarter horizon, we find consumption uncertainty reduced consumption by less than 0.01% in every quarter, similar to the volatility shocks in our model, because expected consumption hides the influence of higher-order moments in future periods. Over horizons long enough to eliminate expected consumption, consumption uncertainty on average reduced consumption by 0.06% and the peak effect was 0.15% during the Great Recession, of which roughly one-third was due to the ZLB constraint. Inflation uncertainty and both consumption and inflation skewness had much smaller impacts on consumption. When we extend our baseline model without capital so households can invest, the average effect of consumption uncertainty increases to -0.08% and the peak effect rises to -0.22% , but the differences from the baseline model are statistically insignificant.

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

Using the capital Euler equation, we find rental rate uncertainty had a small effect on consumption, but uncertainty about Tobin's q had roughly half as large of an effect as consumption uncertainty.

We conduct two exercises to uncover the drivers of our results. Both are easily applied to other models. One, we decompose uncertainty into its endogenous and exogenous sources using counterfactual simulations. The uncertainty that naturally arises in the economy due to first moment shocks accounted for about 95% of total consumption uncertainty. However, nearly all of the variation in uncertainty was driven by the exogenous sources. One exception is when the Fed was constrained. In 2009Q1, about 8.5% of the increase in uncertainty was due to endogenous uncertainty while 38% was due to the endogenous amplification of second moment shocks. Two, we determine the relative importance of each parameter in our model for the results of our Euler equation decomposition using posterior predictive analysis. While price adjustment costs play an important role as others have emphasized, we find the coefficient of relative risk aversion and the monetary response to inflation had the largest impact on the transmission of uncertainty among the deep parameters.

We conclude our analysis by comparing impulse responses to a financial uncertainty shock in our nonlinear model to the same shock in a linear VAR using a recursive identification scheme, since that is the most common way to identify the effects of uncertainty in the literature. Using data simulated from the nonlinear model, we find the VAR generates a quantitatively similar response of consumption growth to a financial uncertainty shock in samples with and without ZLB events. While these results paint VAR methods in a positive light, we point out several reasons for caution.

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

2 RELATED LITERATURE

The literature finds mixed results on the effects of uncertainty, depending on the type of uncertainty and methodology. There are three general approaches to modeling uncertainty: shocks to the cross-sectional standard deviation of firm-level productivity, shocks to the volatility of aggregate exogenous variables, and endogenously-driven sources. When investment is partially irreversible, Bloom (2009) and Bloom et al. (2016) find simultaneous shocks to the volatility of aggregate and firm-level productivity generate roughly a 2% decline in output, though shocks of that nature occur infrequently. In contrast, Bachmann and Bayer (2013) find firm-level uncertainty shocks are a small source of business cycle fluctuations, contributing 2%-5% of the variance of output. Chugh (2016) calibrates a financial accelerator model without irreversible investment using firm-level data and obtains a similar result as Bachmann and Bayer (2013). Christiano et al. (2014) estimate a financial accelerator model with aggregate data and find uncertainty shocks account for 60% of the fluctuations in output. Gilchrist et al. (2014) combine irreversible investment and the financial accelerator mechanism to show that both features are important for the transmission of uncertainty.

Research that studies the effects of aggregate volatility shocks has considered many different exogenous sources of uncertainty. In a small open-economy real business cycle model, Fernández-Villaverde et al. (2011) examine volatility shocks to a country-specific interest rate spread. They find a one standard deviation shock reduces output 0.15%-0.2% in Argentina and Ecuador but only

0.01%-0.02% in Brazil and Venezuela. Other papers use closed-economy New Keynesian models. For example, Mumtaz and Zanetti (2013) focus on monetary policy volatility shocks in a model without capital. They find doubling the volatility reduces output growth by only 0.03%, about five times less than their VAR result. Born and Pfeifer (2014) introduce variable capital utilization and investment adjustment costs. They show a simultaneous two standard deviation increase in uncertainty about government spending, monetary policy, and capital and labor taxes reduces output by only 0.065%. In contrast, Fernández-Villaverde et al. (2015) find a volatility shock to only capital taxes reduces output by 0.1% and the effects are larger when the ZLB binds. Leduc and Liu (2016) include search frictions and habit formation and find a one standard deviation increase in technology volatility increases unemployment by about 2.6%, consistent with their VAR evidence. In a textbook model with recursive preferences, Basu and Bundick (2017) find a one standard deviation preference volatility shock—a proxy for demand uncertainty—reduces output by 0.2%. However, de Groot et al. (2017) show the way the shock enters their preferences creates an asymptote in the parameter space that amplifies the response of output. Without the asymptote, preference volatility shocks have very little real effects. The exogenous volatility shocks in our model also have a small impact, but 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 emphasize the importance of accounting for the *ex-ante* effects of uncertainty over horizons beyond one quarter.

As an alternative to exogenous uncertainty shocks, several papers propose models that endogenously generate uncertainty. One segment emphasizes the role of a financial sector under complete information, where the severity and duration of financial crises are stochastic. Most papers focus on crises that result from financial frictions and collateral constraints [Brunnermeier and Sannikov (2014); He and Krishnamurthy (2014); Mendoza (2010)], while a few papers incorporate the role of firm default [Arellano et al. (2016); Gourio (2014); Navarro (2014)]. Another segment examines the implications of incomplete information. Some of the papers feature learning with aggregate shocks [Fajgelbaum et al. (2017); Saijo (2017); Van Nieuwerburgh and Veldkamp (2006)], while others focus on firm-specific shocks [Ilut and Saijo (2016); Straub and Ulbricht (2015)]. In these models, an adverse shock under asymmetric learning lowers economic activity and makes it harder for households to learn about the economy, which amplifies the effects of first moment shocks. Our paper bridges the gap between the stochastic volatility and endogenous uncertainty literatures by providing a flexible methodology that is easily applied to models with both types of uncertainty.

Our paper is also related to the cost of business cycles literature. Lucas (1987) examines the welfare cost of “instability” by calculating the fraction of consumption goods a household would give up each period to live in a world without volatility. With constant relative risk aversion preferences, He finds the welfare cost of the consumption volatility in post-World War II data ranges from 0.008% (log utility) to 0.17% (risk aversion, $\gamma = 20$). The conclusion is that the cost of instability is insignificant. Several papers have examined these estimates in different settings. For example, Tallarini (2000) finds much higher welfare costs with non-expected utility in a model that matches asset prices and volatilities of macro aggregates. Otrok (2001), however, finds the welfare costs are similar to Lucas’s estimate in a model disciplined to match the persistence in the data. Lester et al. (2014) calculate welfare for several types of preferences and parameter specifications. They find Lucas’s estimates overstate the cost of business cycles and in some parts of the parameter space volatility increases welfare. We build on this literature by calculating welfare at each point in our sample using an estimated model that matches both macro and uncertainty data. We find the welfare costs of first and second moment shocks are well within the range Lucas reported. We view

this important exercise as complementary to our Euler equation decomposition. However, one major advantage of our decomposition is that it tells us which moments (e.g., uncertainty, skewness, covariance) are most important, while still being able to decompose the effects of specific shocks.

3 NEW KEYNESIAN MODEL AND UNCERTAINTY MEASURES

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

3.1 FIRMS The production sector consists of a continuum of monopolistically competitive intermediate goods firms and a final goods firm. Intermediate firm $f \in [0, 1]$ produces a differentiated good, $y_t(f)$, according to $y_t(f) = z_t n_t(f)$, where $n_t(f)$ 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

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

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

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

$$\tilde{y}_t = n_t, \quad (3)$$

$$mc_t = \tilde{w}_t, \quad (4)$$

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

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

3.2 HOUSEHOLDS The representative household chooses $\{c_t, n_t, b_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \beta^t [((c_t/z_t)^{1-\gamma} - 1)/(1-\gamma) - \chi n_t^{1+\eta}/(1+\eta)]$, where γ is the coefficient of relative risk aversion, $\chi > 0$ is a preference parameter that determines the steady state labor supply, $1/\eta$ is the Frisch elasticity of labor supply, c is consumption, n is labor hours, b is the real value of a privately-issued 1-period nominal bond that is in zero net supply, and E_0 is the mathematical expectation operator conditional on information in period 0. Following An and Schorfheide (2007),

households receive utility from consumption relative to the level of technology, which is a proxy for the habit stock. That assumption allows us to use additively separable preferences and parameterize the degree of risk aversion while maintaining a balanced growth path. The household's choices are constrained by $c_t + b_t/(i_t s_t) = w_t n_t + b_{t-1}/\pi_t + d_t$, where π is the gross inflation rate, w is the real wage rate, i is the gross nominal interest rate set by the central bank, and d is a real dividend received from owning the intermediate goods firms. Following Smets and Wouters (2007) and Gust et al. (2017), s is a shock to the risk premium on the nominal bond and it evolves according to

$$s_t = (1 - \rho_s) + \rho_s s_{t-1} + \sigma_{v,t} v_t, \quad 0 \leq \rho_s < 1, \quad v \sim \mathbb{N}(0, 1), \quad (6)$$

$$\sigma_{v,t} = \bar{\sigma}_v (\sigma_{v,t-1}/\bar{\sigma}_v)^{\rho_\sigma} \exp(\sigma_\zeta \zeta_t), \quad 0 \leq \rho_{\sigma_v} < 1, \quad \zeta \sim \mathbb{N}(0, 1), \quad (7)$$

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

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

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

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

Equation (9) is the consumption Euler equation that we will use to examine the economic effects of the real interest rate and higher-order moments, including consumption uncertainty and skewness.

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

$$i_t = \max\{1, i_t^n\}, \quad (10)$$

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

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, $\bar{i}$ 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. When the ZLB binds, a more negative net notional rate means the central bank is more constrained.

3.4 COMPETITIVE EQUILIBRIUM The aggregate resource constraint is given by

$$\tilde{c}_t = \tilde{y}_t^{gdp}, \quad (12)$$

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

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, $\{s_t, g_t, \sigma_{\varepsilon,t}, \sigma_{v,t}\}_{t=0}^\infty$, that satisfy the detrended equilibrium system, (1)-(13), given the initial conditions, $\{c_{-1}, i_{-1}^n, s_0, a_0, \nu_0, \sigma_{\varepsilon,0}, \sigma_{v,0}\}$, and the five sequences of shocks, $\{\varepsilon_t, v_t, \nu_t, \xi_t, \zeta_t\}_{t=1}^\infty$.

3.5 MEASURES OF UNCERTAINTY The stochastic volatility processes, (2) and (7), create exogenous sources of time-varying supply and demand uncertainty. Uncertainty is measured by the

ex-ante standard deviation of future technology growth and the future risk premium, which equal

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

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

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

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

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

which is the same way we measure exogenous uncertainty, except it is calculated with an endogenous variable. Both measures of uncertainty remove the predictable component of the forecasted variable instead of only a constant trend, so they distinguish between uncertainty and conditional volatility. However, the endogenous uncertainty measure not only fluctuates due to exogenous volatility shocks, but also due to events that happen in the economy. For example, when the notional interest rate is negative, the economy is more sensitive to first moment shocks that adversely affect the economy, which increases the endogenous uncertainty about consumption growth. The ZLB constraint also creates uncertainty by amplifying the effect of the exogenous volatility shocks. When the ZLB does not bind, first moment shocks still affect U_{c^g} but the magnitudes are smaller.³

4 NUMERICAL METHODS AND DECOMPOSITION

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

We discretize the state space and then approximate the stochastic volatility processes, (2) and (7), and first moment shocks, ε , v , 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 E](#) for a detailed description of the solution method.

³With a nonlinear solution, it is possible to calculate the endogenous uncertainty about any variable in a model.

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

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

where μ_i and σ_i , $i \in \{MacroU, FinU\}$, are the mean and standard deviation of a given time series.




Figure 1: Measures of uncertainty.

Figure 1 plots the standardized 1-quarter ahead $MacroU$ and $FinU$ series, which inform the parameters in our model and ensure it produces the same fluctuations in uncertainty as the data. The uncertainty series are based on a factor augmented vector autoregression (FAVAR) that accounts for 132 macroeconomic and 147 financial variables. Repeated simulations of the FAVAR are used to obtain estimates of uncertainty for each macro (financial) variable and then averaged to obtain the $MacroU$ ($FinU$) time series. The benefit of these particular series is that they are calculated in the same way as (14), so they distinguish between uncertainty and conditional volatility, and they reflect the uncertainty surrounding a large set of macro (financial) variables. Examples of the variables in the $MacroU$ series include output, employment, housing starts, interest rates, and prices, while the $FinU$ series includes dividends, returns, spreads, and the Fama-French factors.

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

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 nonlinear model and particle filter.

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

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

$\xi \sim \mathcal{N}(0, \Sigma)$ is a vector of MEs, and $\Sigma = \text{diag}([\sigma_{me,y^g}^2, \sigma_{me,\pi}^2, \sigma_{me,i}^2, \sigma_{me,macrou}^2, \sigma_{me,finu}^2])$. It is not practical to estimate the ME variances because they are inversely related to the model likelihood. The Metropolis-Hastings algorithm would prefer lower ME variances, which would decrease the effective sample size in the particle filter and reduce its accuracy. Following Herbst and Schorfheide (2016), we set the ME variance of real GDP growth, the inflation rate, and both uncertainty series to 20% of their variance in the data. However, the ME variance for the policy rate is set to 2% of its variance in the data because the federal funds rate is less noisy and it affects the level of uncertainty predicted by the model near the ZLB. We decided to link consumption growth uncertainty to the macro uncertainty index and risk premium uncertainty to the financial uncertainty index because Ludvigson et al. (2017) find financial uncertainty is an exogenous impulse that causes recessions, whereas macro uncertainty endogenously responds to other shocks that affect the business cycle. In our baseline model, consumption 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 correlation 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 second moment shocks as well as endogenous dynamics. One way to quantify the effect of uncertainty is by decomposing the consumption with the Euler equation, (9). A third-order approximation around the balanced growth path implies

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

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 , $\hat{r}_t \equiv \hat{i}_t + \hat{s}_t - E_t \hat{\pi}_{t+1} - E_t \hat{g}_{t+1}$ is the *ex-ante* real interest rate, and a hat denotes log deviation from the balanced growth path. 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 they had almost no effect on consumption in our sample. The variance, skewness, and covariance terms quantify the effect of the uncertainty, upside and downside risk, and the pairwise linear relationships between consumption, inflation, technology growth next period. Higher risk aversion means households are less willing to intertemporally substitute consumption goods, which makes them less sensitive to the real interest rate and more sensitive to the variance and skewness of consumption next period. Much of our analysis will focus on the variance of consumption because it will have the exact same effect on current consumption regardless of which Euler equation is used for the decomposition.

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

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

where $q \geq 1$ is the forecast horizon. The sum of each variance term over q quarters captures the effect of a given type of uncertainty, conditional on expected consumption in quarter q . When q becomes sufficiently large, the conditional expectation drops out of the decomposition, so we are able to determine the unconditional effects of each higher-order moment. Over a 1-quarter horizon, expected consumption closely tracks current consumption, which hides the effect of higher-order moments in future quarters. By decomposing expected future consumption, we can show how the uncertainty, skewness, and covariance terms affect consumption over horizons beyond one quarter.⁴

⁴After iterating, we obtain $E_t[\text{cov}_{t+j}(x_{t+j+1}, y_{t+j+1})] = \text{cov}_t(x_{t+j}, y_{t+j}) - \text{cov}_t(E_{t+j}[x_{t+j+1}], E_{t+j}[y_{t+j+1}])$ by the law of total covariance. In our derivation, we ignore the second term because its effects are quantitatively small.

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

5 ESTIMATED EFFECTS OF UNCERTAINTY

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

Parameter	Prior			Posterior			
	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_v})	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
Growth Shock SD ($\bar{\sigma}_\varepsilon$)	IGam	0.0075	0.0075	0.00371	0.00054	0.00288	0.00463
Risk Shock SD ($\bar{\sigma}_v$)	IGam	0.0025	0.0025	0.00139	0.00022	0.00107	0.00177
Growth 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.

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

The last four columns display the posterior means, standard deviations, and 90% credible sets for the estimated parameters. Low frequency movements in the macro and financial uncertainty time series coupled with sharp increases in both series during the Great Recession generate highly persistent stochastic volatility processes with large shock standard deviations. For example, a two standard deviation supply uncertainty shock causes a 25.1% increase in the volatility of technology growth with a half-life of about 15.9 quarters. The monetary policy parameters imply a high degree of interest rate smoothing and strong responses to real GDP growth and inflation, which are necessary for the model to explain the long ZLB period. The mean estimates of the annualized technology growth and inflation rates are 1.77% and 2.62%, which are slightly higher than the values in the data since they are unconditional and under-represent the effects of the ZLB period. The mean coefficient of relative risk aversion is consistent with An and Schorfheide (2007). The Rotemberg price adjustment cost parameter implies a slope of the Phillips curve of about 0.035, which is in line with other estimates in the literature. Overall, the priors and posterior means are consistent with Gust et al. (2017), who estimate a similar model with an occasionally binding ZLB constraint but without stochastic volatility. [Appendix G](#) provides additional estimation diagnostics, including the filtered states and shocks, the ZLB distributions, and various unconditional moments.

5.2 IMPULSE RESPONSES We begin our analysis by showing impulse responses to first and second moment shocks to illustrate the underlying dynamics in the model. [Figure 2](#) plots the responses to a 2 standard deviation positive risk premium, risk premium volatility, growth, and growth volatility shock. The parameters are set to their posterior means and the simulations are initialized at two different states. Our benchmark simulation is initialized at the stochastic steady state and reflective of any state of the economy where there is virtually no expectation of hitting the ZLB. We compare those responses to the responses when the notional rate is negative by initializing the simulation at the filtered state vector corresponding to 2009Q2. The effect of mean reversion is removed from the responses by plotting the percentage point difference (percent change 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 consumption, 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.




Figure 2: Impulse responses to a 2 standard deviation positive shock at and away from the ZLB. The steady-state simulation (solid line) is initialized at the stochastic steady state. The other simulation (dashed line) is initialized at the filtered state corresponding to 2009Q2 so the ZLB binds. The vertical 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.

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 cause bigger changes 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 in a few ways. One, growth volatility directly affects consumption volatility. Therefore, uncertainty increases more than it does in response to a risk premium volatility shock. Two, the response of uncertainty is similar in both initial states. Three, the increase in uncertainty away from the ZLB is much larger than the increase from a risk premium volatility shock. Therefore, growth volatility shocks play a larger role in explaining the fluctuations in uncertainty when the ZLB does not bind.




Figure 3: Sources of uncertainty in our nonlinear model.

5.3 SOURCES OF UNCERTAINTY The impulse responses show uncertainty can arise in our nonlinear model due to exogenous volatility shocks or first moment shocks that interact with the economy. Figure 3a decomposes consumption growth uncertainty into its exogenous and endogenous sources using counterfactual simulations conditional on the mean parameterization of our model.

To isolate the contribution of technology growth uncertainty, we turn off the risk premium volatility shocks. Similarly, we zero out the technology growth volatility shocks to identify the amount of risk premium uncertainty. We then turn off both volatility shocks to determine the amount of endogenous uncertainty. We also show the endogenous amplification of the exogenous volatility shocks when the Fed was most constrained using the solution to the unconstrained nonlinear model.

On average about 95% of consumption growth uncertainty is due to first moment shocks that interact with the economy—what we refer to as endogenous uncertainty. However, the vast majority of the changes in uncertainty are driven by second moment shocks. Growth volatility shocks are the key driver in most periods, but risk premium volatility shocks play an important role in certain parts of our sample. Typically, endogenous uncertainty is fairly constant, but it increases when the policy rate is near or at its ZLB, which occurs in the mid 2000s and from 2009 to the end of the sample. The sharp increase in uncertainty in 2009, however, primarily occurred due to the endogenous amplification of the exogenous volatility shocks, rather than through first moment shocks. The markers in 2009Q1 show the counterfactual increase in uncertainty that would have occurred if the Fed was not constrained. Those results indicate that about 8.5% $((0.48 - 0.43)/(1.03 - 0.43))$ of the increase in uncertainty in 2009Q1 was due to endogenous uncertainty and about 38% $((1.03 - 0.82)/(1.03 - 0.48))$ was due to the endogenous amplification of second moment shocks.

Despite some nonlinear interactions between the exogenous volatility shocks and the ZLB, we are able to approximate the relative contribution of each volatility shock over time, similar to a variance decomposition in a linear model. The dark bars in [figure 3b](#) represent the technology growth counterfactual relative to the endogenous uncertainty counterfactual (circles minus diamonds) and the light bars represent the risk premium counterfactual relative to the endogenous uncertainty counterfactual (triangles minus diamonds), which is approximately equal to consumption growth uncertainty relative to the endogenous uncertainty counterfactual (solid minus diamonds). The results reiterate that technology growth uncertainty is typically the biggest contributor to consumption growth uncertainty, but the two sources of exogenous uncertainty typically move together. There are two notable exceptions. One, the model predicts that risk premium uncertainty precedes the 2001 recession. Two, technology growth uncertainty increases before the rise in risk premium uncertainty during the Great Recession, but the effects of risk premium uncertainty linger while the impact of technology growth uncertainty is negligible for a few years after the Great Recession. During the Great Recession, technology growth and risk premium volatility shocks have nearly equal roles. By the end of the sample, consumption growth uncertainty declined to its lowest point.

5.4 EULER EQUATION DECOMPOSITION We now turn to our main results. [Figure 4](#) shows the filtered time series of the terms in the Euler equation decomposition in (16) over different forecast horizons. The values on the vertical axes are the 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) from the higher-order terms (right panel) so the effects of expected consumption and the real interest rate do not drown out the effects of either the uncertainty or skewness terms. We also plot current consumption in the top left panel.

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 about 0.1%. The peak effect was -0.37% during the Great Recession, but that effect quickly declined as the economy rebounded. The higher-order terms show uncertainty about consumption in the next quarter had time-varying adverse effects on current consumption. Un-




Figure 4: 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.

certainty had by far its largest effect during the Great Recession, since the ZLB constraint made the economy more sensitive to adverse shocks and increased the expected volatility of future consumption. However, the quantitative effects were small throughout our sample. Even during the Great Recession, the peak increase in uncertainty reduced consumption by less than 0.01%. Consumption skewness and both inflation uncertainty and inflation skewness also had very little effect on current consumption. Interestingly, the effects of uncertainty over a one-quarter horizon have a similar magnitude as the impact effects of the two exogenous volatility shocks shown in [figure 2](#). However, both sets of results significantly understate the impacts of uncertainty because they hide the effects that future real interest rates and higher-order moments have on expected consumption.

The middle left panel shows how expected 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, explained a larger fraction of the changes in consumption. We 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%.

The middle right panel shows the effect of consumption uncertainty over the horizons shown in the left panel, but 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 decompose the influence of expected future consumption. Over a 24-quarter horizon, consumption uncertainty on average decreases current consumption by about 0.06% and the largest effect was about 0.15% in 2009Q1, which accounted for 16.6% of the total decline in that quarter.

The other higher order moments are shown in the bottom left panel. During the Great Recession, the peak effects of technology growth uncertainty, inflation uncertainty, and consumption skewness over a 24-quarter horizon were -0.023% , -0.005% , and -0.001% , respectively, and the average effects were much smaller. We do not show the effect of inflation skewness because it is 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 since it prevents the Fed from responding to adverse shocks through conventional channels. Evidently, those effects on consumption are small when controlling for other terms. The covariance between inflation and consumption reduced consumption by about 0.02% on average, the second largest effect behind consumption uncertainty.

The bottom right panel shows the effect of consumption uncertainty over a 24-quarter horizon along with two of the counterfactuals shown in [figure 3a](#). First, we plot the effect of consumption uncertainty after removing the influence of the ZLB using the solution to the unconstrained nonlinear model. The differences from the baseline path show how much the ZLB increased the adverse effects of uncertainty. In most quarters, the differences are small because there is a low probability of going to and staying at the ZLB. Larger differences between the two paths occurred from 2008Q4 to 2009Q4, when the notional rate was well below zero and there was a strong expectation of staying at the ZLB. For example, in 2009Q1 consumption uncertainty reduced consumption by

about 0.06 percentage points more than it would have if the Fed was not constrained. Second, we restore the ZLB constraint but zero out both of the exogenous volatility shocks. Since first moment shocks are the main source of most of the uncertainty in the economy, they are also the primary source of the adverse effects of consumption uncertainty. At its peak, endogenous uncertainty only increased the adverse effects of consumption uncertainty by about 0.01 percentage points, whereas exogenous volatility shocks played a much larger role during the last two recessions. For example, the volatility shocks without the ZLB contributed about 0.04 percentage points to the decline in consumption growth in 2009Q1 and their amplification contributed another 0.06 percentage points.

The bottom right panel also plots the total effect of uncertainty—the sum of consumption, technology growth, and inflation uncertainty—on consumption over a 24-quarter horizon. On average, total uncertainty reduced current consumption by about 0.07% with a maximum decline of 0.17%.




Figure 5: 68% credible sets of the filtered effects on current consumption. The vertical bars denote NBER recessions. The vertical axis is the contribution to the percentage point deviation of detrended consumption from its steady state.

The results in [figure 4](#) are based on the mean parameterization of our model. Although numerically intensive, it is also possible to generate time series for the terms in the decomposition for each of the 1,000 draws from our posterior distribution. [Figure 5](#) shows the 1 standard deviation (16%-84%) credible sets for the effects of the three types uncertainty over a 24-quarter horizon. The left panel shows the effects of consumption uncertainty while the right panel shows the effects of technology growth and inflation uncertainty. In a typical quarter, the effect of consumption uncertainty ranges from -0.01% to $+0.015\%$ of the median effect. The effects, however, are more asymmetric during recessions. For example, during the peak of the Great Recession there was a 68% chance consumption uncertainty decreased current consumption by at least 0.12% and it could have decreased it by as much as 0.25%. While the effects of technology growth and inflation uncertainty are always considerably smaller than consumption uncertainty, it is possible technology growth uncertainty decreased current consumption by nearly 0.04% during the peak of the Great Recession.

5.5 KEY PARAMETERS In addition to quantifying the contribution of each source of uncertainty, we also determine the relative importance of each parameter for our Euler equation decomposition by conducting posterior predictive analysis with the draws from the joint posterior density, $\{\hat{\theta}_i\}_{i=1}^{1000}$.

We focus on the effect of consumption uncertainty over a 24-quarter horizon, which is given by

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

where θ is a vector of estimated parameters and $\bar{v}_t$ and $\bar{z}_t$ are the median filtered states and shocks conditional on the posterior mean parameterization. Fixing the states and shocks allows us to isolate the role of each parameter. We already calculate $h(\theta_i, t)$ for all i to generate the credible sets in [figure 5](#), which represent the actual effects of consumption uncertainty given each posterior draw. Define $\bar{\theta}_{i,\ell}$ as the i th posterior draw conditional on the posterior mean of parameter ℓ . As a counterfactual, we first calculate $h(\bar{\theta}_{i,\ell}, t)$ using the procedure described at the end of [section 4](#). We then calculate the root mean-squared deviation (RMSD) from that counterfactual, given by

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

For our analysis, we calculate the RMSD for each of the 16 estimated parameters across the entire sample. A higher RMSD implies a given parameter has a bigger impact on the effect of consumption uncertainty on current consumption. [Figure 6](#) plots time series of the RMSD for the nine most consequential parameters. The risk premium persistence (ρ_s) and shock standard deviation (σ_v) have the largest average RMSDs. Of the deep parameters, the coefficient of relative risk aversion (γ) and the monetary response to inflation (ϕ_π) are the most consequential. There is also considerable variation in the importance of the parameters across time. For example, during recessions the RMSD of each parameter increases, but the process parameters (right panel) become relatively more important than the deep parameters. Outside recessions, the deep parameters (left panel) are relatively more important, though the average RMSD of each parameter is much lower.

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

The results depend on how each parameter affects expected future volatility. The dashed diagonal line is the linear trend. When the parameters governing the risk premium (ρ_s , ρ_{σ_v} , $\bar{\sigma}_v$, and σ_ζ) are above their posterior means (shown by the dashed vertical line) consumption uncertainty has a much larger adverse effect because the variance of the exogenous process and hence expected volatility increase (i.e., the trend line slopes downward). For example, the posterior mean persistence of the risk premium, ρ_s , is 0.911. When that value is two standard deviations higher (0.932), consumption uncertainty reduces consumption by 0.13 percentage points more than at its posterior mean. Similarly, a higher price adjustment cost parameter causes uncertainty to have a larger effect because stickier prices make households more sensitive to changes in the nominal interest rate. Larger values of the other parameters reduce the effect of uncertainty. An increase in the coefficient of relative risk aversion makes households less willing to substitute across time, which makes consumption less volatile. Thus, consumption uncertainty has a smaller adverse effect, even though households react more strongly to expected volatility. A higher monetary response to inflation has a similar effect because it also reduces expected future volatility. Interest rate smoothing




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




Figure 7: Scatter plot of the effect of consumption uncertainty on consumption over a 24-quarter horizon in 2008Q4 at each posterior draw, conditional on fixing the indicated parameter at its posterior mean. The values on the vertical axis are shown in deviations from the effect at the posterior mean parameterization. The horizontal axis denotes the parameter values corresponding to the number of standard deviations away from the posterior mean. The dashed vertical line denotes the posterior mean and the dashed diagonal line is the linear trend based on ordinary least squares.

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

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

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

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

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

where

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

is the expected present-value of the household's utility from consumption and disutility from labor conditional on their information set at time t , $\hat{\Omega}_t$, which contains the median filtered state variables and the posterior mean parameters. Also, $\tilde{c}^\vartheta$ and n^ϑ are the optimal choices of detrended consumption and labor conditional on model $\vartheta \in \{H, L\}$. We denote the higher (lower) volatility economy with H (L), where the expected path of consumption is lower (higher) due to precautionary saving. Therefore, λ_t is the fraction of consumption goods in the low volatility economy that would compensate the household for the lower consumption path in the higher volatility economy. If $\lambda_t > 0$, then households are better off in the lower volatility economy. [Appendix B](#) provides the derivation.

We approximate W_c and W_n by integrating across 1,000 Monte Carlo simulations of 10,000 quarters initialized at the median filtered state corresponding to each point in the sample. We compute W_c and W_n for four different models: (1) the baseline model, (2) the model without the ZLB imposed, (3) the model without volatility shocks, and (4) the model without any shocks (equivalent to its deterministic steady state). Since our compensating variation calculation is conditional on the state of the economy and the amount of uncertainty is conditional on those states, we generate a time-series of estimated compensating variations conditional on the median filtered state, the most likely state preferred by the data at each point in time given the posterior mean parameterization.

[Figure 8](#) shows four estimated compensating variations: the effect of all shocks (models (1) and (4), x markers); the effect of only the stochastic volatility shocks (models (1) and (3), circles markers); the effect of only the first-moment shocks (models (3) and (4), triangle markers); and




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

the effect of only the ZLB constraint (models (1) and (2), diamond markers). The compensating variation is shown as the percent of consumption goods in the lower volatility or no ZLB models.

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

6 THE EFFECT OF CAPITAL ACCUMULATION

In our baseline model without capital, 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 capital. In the data, investment is more volatile than real GDP, especially during recessions, so it is important to add capital to the model since it allows output, consumption, and investment to have different, potentially time-varying, volatilities.

The final goods firm's problem is unchanged. Intermediate firm $f \in [0, 1]$ 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 symmetric equilibrium, the production function and optimality conditions are

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

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

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

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

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

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

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

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

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

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

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

$$\tilde{c}_t + \tilde{x}_t = \tilde{y}_t^{gdp}. \quad (24)$$

Once again, we redefined variables that grow along the balanced growth path in terms of technology. A competitive equilibrium includes infinite sequences of quantities, $\{\tilde{c}_t, \tilde{y}_t, \tilde{y}_t^{gdp}, n_t, \tilde{x}_t, \tilde{k}_t\}_{t=0}^\infty$, 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, $\{s_t, g_t, \sigma_{\varepsilon,t}, \sigma_{v,t}\}_{t=0}^\infty$, that satisfy the detrended equilibrium system, (1), (2), (5)-(11), (13), and (18)-(24), given the initial conditions, $\{c_{-1}, i_{-1}^n, x_{-1}, k_{-1}, s_0, a_0, \nu_0, \sigma_{\varepsilon,0}, \sigma_{v,0}\}$, and the five sequences of shocks, $\{\varepsilon_t, v_t, \nu_t, \xi_t, \zeta_t\}_{t=1}^\infty$.

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. (2017). Although there are some differences between our model and the one in Gust et al. (2017) (e.g., their model includes sticky wages and variable capital 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.

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

Figure 9 shows the influence of the different types of uncertainty. The left panel plots the effects of consumption uncertainty over a 24-quarter horizon in the models with and without capital. In the capital model, consumption uncertainty on average decreases current consumption by 0.08%, which is only 0.02 percentage points more than in our baseline model without capital. The




Figure 9: 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.

difference is more pronounced when the ZLB first binds. In 2008Q4 consumption uncertainty lowered consumption by 0.22% compared with only 0.15% in the baseline model, but that discrepancy quickly dissipated. Furthermore, the median effect in the capital model is typically in the left tail of the credible set implied by the baseline model. The right panel shows the impact of technology growth and inflation uncertainty in the model with capital—the two other types of uncertainty in the bond Euler equation. Both terms have nearly identical effects to those in the baseline model.⁵

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

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

which we can once again iterate forward to eliminate the influence of expected future consumption. Several terms enter the same way as our previous decomposition. For example, the *ex-ante* variance of consumption and technology growth appear in (21) and (25), so they have the exact same effect on consumption. The rental rate and Tobin's q variance terms replace the inflation variance term.

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

⁵Consumption and inflation skewness as well as the covariance terms also have very similar effects in both models.

7 COMPARISON WITH THE TRADITIONAL VAR APPROACH

Due to the nonlinearities introduced by stochastic volatility and the ZLB constraint, we are interested in whether a *linear* VAR, commonly employed in the literature on uncertainty, can recover the dynamic relationship between uncertainty and real activity predicted by our baseline *nonlinear* model. We focus on the impulse response of consumption growth to a change in financial uncertainty, $U_{s,t} = \sqrt{E_t[\sigma_{v,t+1}^2]}$. Since financial uncertainty is exogenous in our structural model, it is easy for us to compare its effects to those in a VAR model. The shocks in the VAR are identified recursively and the variables—financial uncertainty, consumption growth, inflation, wage growth, the risk premium, and the interest rate—are ordered from first to last as in Christiano et al. (2005).⁶




Figure 10: Impulse responses of consumption growth to a 2 standard deviation increase in financial uncertainty. The solid line is the median response and the shaded regions denote the one standard deviation (16%-84%) credible sets.

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

⁶We obtain very similar results using bivariate VARs with uncertainty ordered first and consumption growth second.

The simulated VARs in the next three subplots are estimated with data from short-sample simulations of the baseline model conditional on the posterior mean parameterization. The solid lines represent the median response and the shaded regions represent the 16%-84% credible sets. The first simulated VAR is estimated using artificial data without any ZLB events (i.e., $i^* > 0$ always). The response of consumption growth to a financial uncertainty shock is close to zero across the whole horizon, which is very similar to the prediction of our baseline model initialized at steady state. The next two simulated VARs are estimated with artificial data where the notional rate falls below -0.4% or -1.5% for at least one quarter, so the responses represent averages across quarters when the ZLB does and does not bind. Given these initial states, the median impact effects of financial uncertainty on consumption growth are -0.06% and -0.10% , respectively. Although the responses are not significantly different from zero, the median impact of financial uncertainty shocks identified by the VAR decreases as the quantity and severity of ZLB events increase in the simulated data and it is quantitatively similar to our structural model. Therefore, the *linear* VAR does a good job capturing the endogenous amplification of exogenous financial uncertainty shocks.

Finally, we estimate the same VAR with actual U.S. data. The second to last subplot excludes the Great Recession and subsequent ZLB period (1986Q1-2007Q4), while the last subplot is based on the sample used to estimate our baseline model (1986Q1-2016Q2). Qualitatively, the results are similar to the predictions of our structural model—the effect of uncertainty is more pronounced when the ZLB period is included in the sample, both on impact and over time. However, there are two notable differences from our structural model. One, in the full sample the response of consumption growth is significantly negative over the entire horizon, instead of only on impact. Two, the impact effect is significantly negative in the truncated sample that omits the ZLB period.

The differences between the responses in the simulated and actual VARs have two potential explanations. One, there are features of the economy besides the ZLB constraint that amplify the impact effect of financial uncertainty shocks. Potential examples include borrowing constraints, irreversible investment, and limited information. Any of these features would increase the impact effect of financial uncertainty in the pre-ZLB period. Two, financial uncertainty is not purely exogenous and instead fluctuates due to events that are happening in the economy (i.e., first moment shocks that increase the *ex-ante* variance of the risk premium). In this case, with uncertainty ordered first, the VAR would over-predict the effect of financial uncertainty because it would assume all changes in uncertainty are exogenous. In reality, both explanations likely play important roles. Furthermore, the challenges associated with classifying uncertainty as endogenous or exogenous become even more severe when accounting for multiple types of uncertainty in the VAR model.

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

8 CONCLUSION

The literature has primarily examined the economic effects of uncertainty through stochastic volatility shocks in dynamic macro models and impulse responses to various proxies for uncertainty in VAR models. We develop a new way to quantify the effects of uncertainty that accounts for both

exogenous and endogenous sources of uncertainty. First, we estimate a nonlinear New Keynesian model, linking measures of uncertainty in the data to equivalent measures in the model. Second, we use the Euler equation to decompose current consumption into first, second, and third moments and then filter the data at the posterior mean to create a time series for each term in the decomposition.

Over a 1-quarter horizon, consumption uncertainty reduced consumption by less than 0.01% in all periods, similar to the volatility shocks in our model. Over horizons that remove the influence of expected future consumption, the effect of consumption uncertainty was an order of magnitude larger. On average, it reduced consumption by about 0.06% and the peak effect was 0.15% during the Great Recession. Adding capital accumulation to the baseline model had a fairly small impact—the average effect increases to -0.08% and the peak effect becomes -0.22% but those differences were not statistically significant. Other higher-order terms, such as inflation uncertainty, rental rate uncertainty, and the various skewness terms had smaller effects in both models.

A major benefit of our method is its flexibility. It can be used to examine the economic effects of any type of uncertainty in a broad class of models. While some models are too costly to estimate, it is usually possible to calculate the *ex-ante* variance or skewness surrounding any endogenous variable in the model and then link it to an empirical measure while filtering the data. For example, one could compute the uncertainty surrounding the exchange rate in an open economy model or any financial variable in a model with a banking sector. Similarly, one could generate time series for the terms in the Euler equation in almost any macro model to determine the effects of various higher-order moments on consumption. Using a familiar New Keynesian model, our decomposition reveals that consumption uncertainty had a modest effect. Future research could introduce additional sources of endogenous uncertainty, such as borrowing constraints, search frictions, limited information, irreversible investment, or heterogeneity to examine whether a larger effect emerges.

REFERENCES

- ALEXOPOULOS, M. AND J. COHEN (2009): “Uncertain Times, Uncertain Measures,” University of Toronto Working Paper 352.
- AN, S. AND F. SCHORFHEIDE (2007): “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26, 113–172.
- ARELLANO, C., Y. BAI, AND P. J. KEHOE (2016): “Financial Frictions and Fluctuations in Volatility,” NBER Working Paper 2016.
- BACHMANN, R. AND C. BAYER (2013): “‘Wait-and-See’ Business Cycles?” *Journal of Monetary Economics*, 60, 704–719.
- BACHMANN, R., S. ELSTNER, AND E. SIMS (2013): “Uncertainty and Economic Activity: Evidence from Business Survey Data,” *American Economic Journal: Macroeconomics*, 5, 217–49.
- BASU, S. AND B. BUNDICK (2015): “Endogenous Volatility at the Zero Lower Bound: Implications for Stabilization Policy,” FRB Kansas City Working Paper 15-1.
- (2017): “Uncertainty Shocks in a Model of Effective Demand,” *Econometrica*, 85, 937–958.
- BEKAERT, G., M. HOEROVA, AND M. LO DUCA (2013): “Risk, Uncertainty and Monetary Policy,” *Journal of Monetary Economics*, 60, 771–788.
- BERNDT, A., H. LUSTIG, AND ŞEVİN YELTEKIN (2012): “How Does the US Government Finance Fiscal Shocks?” *American Economic Journal: Macroeconomics*, 4, 69–104.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.

- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2016): “Really Uncertain Business Cycles,” Manuscript, Stanford University.
- BORN, B. AND J. PFEIFER (2014): “Policy Risk and the Business Cycle,” *Journal of Monetary Economics*, 68, 68–85.
- BRUNNERMEIER, M. K. AND Y. SANNIKOV (2014): “A Macroeconomic Model with a Financial Sector,” *The American Economic Review*, 104, 379–421.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- CHRISTIANO, L. J., R. MOTTO, AND M. ROSTAGNO (2014): “Risk Shocks,” *American Economic Review*, 104, 27–65.
- CHUGH, S. K. (2016): “Firm Risk and Leverage Based Business Cycles,” *Review of Economic Dynamics*, 20, 111–131.
- CHUNG, H. AND E. M. LEEPER (2007): “What Has Financed Government Debt?” NBER Working Paper 13425.
- COLEMAN, II, W. J. (1991): “Equilibrium in a Production Economy with an Income Tax,” *Econometrica*, 59, 1091–1104.
- DE GROOT, O., A. RICHTER, AND N. THROCKMORTON (2017): “Uncertainty Shocks in a Model of Effective Demand: Comment,” Federal Reserve Bank of Dallas Working Paper 1706.
- 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. SCHAAAL (2017): “Uncertainty Traps,” *The Quarterly Journal of Economics*, 132, 1641–1692.
- FERNALD, J. G. (2012): “A quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- FERNÁNDEZ-VILLAYERDE, 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-VILLAYERDE, 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-VILLAYERDE, J. AND J. F. RUBIO-RAMÍREZ (2007): “Estimating Macroeconomic Models: A Likelihood Approach,” *Review of Economic Studies*, 74, 1059–1087.
- GILCHRIST, S., J. W. SIM, AND E. ZAKRAJŠEK (2014): “Uncertainty, Financial Frictions, and Investment Dynamics,” NBER Working Paper 20038.
- GORDON, N. J., D. J. SALMOND, AND A. F. M. SMITH (1993): “Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation,” *IEEE Proceedings F - Radar and Signal Processing*, 140, 107–113.
- GOURIO, F. (2014): “Financial Distress and Endogenous Uncertainty,” Manuscript, Federal Reserve Bank of Chicago.
- GUERRÓN-QUINTANA, P. A. AND J. M. NASON (2013): “Bayesian Estimation of DSGE Models,” in *Handbook of Research Methods and Applications in Empirical Macroeconomics*, Edward Elgar Publishing, chap. 21, 486–512.
- GUST, C., E. HERBST, D. LÓPEZ-SALIDO, AND M. E. SMITH (2017): “The Empirical Implications of the Interest-Rate Lower Bound,” *American Economic Review*, 107, 1971–2006.

- HALL, G. J. AND T. J. SARGENT (2011): "Interest Rate Risk and Other Determinants of Post-WWII US Government Debt/GDP Dynamics," *American Economic Journal: Macroeconomics*, 3, 192–214.
- HE, Z. AND A. KRISHNAMURTHY (2014): "A Macroeconomic Framework for Quantifying Systemic Risk," NBER Working Paper 19885.
- HERBST, E. P. AND F. SCHORFHEIDE (2016): *Bayesian Estimation of DSGE Models*, Princeton, NJ: Princeton University Press.
- ILUT, C. L. AND H. SAIJO (2016): "Learning, Confidence, and Business Cycles," NBER Working Paper 22958.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.
- KITAGAWA, G. (1996): "Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models," *Journal of Computational and Graphical Statistics*, 5, pp. 1–25.
- KOPECKY, K. AND R. SUEN (2010): "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 13, 701–714.
- LEDUC, S. AND Z. LIU (2016): "Uncertainty Shocks are Aggregate Demand Shocks," *Journal of Monetary Economics*, 82, 20–35.
- LESTER, R., M. PRIES, AND E. SIMS (2014): "Volatility and Welfare," *Journal of Economic Dynamics and Control*, 38, 17–36.
- LUCAS, JR, R. E. (1987): *Models of Business Cycles*, Oxford: Basil Blackwell.
- LUDVIGSON, S. C., S. MA, AND S. NG (2017): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" NBER Working Paper 21803.
- MASON, J. W. AND A. JAYADEV (2014): "'Fisher Dynamics" in US Household Debt, 1929–2011," *American Economic Journal: Macroeconomics*, 6, 214–234.
- MENDOZA, E. G. (2010): "Sudden Stops, Financial Crises, and Leverage," *The American Economic Review*, 100, 1941–1966.
- MUMTAZ, H. AND F. ZANETTI (2013): "The Impact of the Volatility of Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 45, 535–558.
- NAVARRO, G. (2014): "Financial Crises and Endogenous Volatility," Manuscript, New York University.
- OTROK, C. (2001): "On Measuring the Welfare Cost of Business Cycles," *Journal of Monetary Economics*, 47, 61–92.
- PARKER, J. A. AND B. PRESTON (2005): "Precautionary Saving and Consumption Fluctuations," *American Economic Review*, 95, 1119–1143.
- PETERMAN, W. B. (2016): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," *Economic Inquiry*, 54, 100–120.
- PLANTE, M., A. W. RICHTER, AND N. A. THROCKMORTON (2018): "The Zero Lower Bound and Endogenous Uncertainty," *Economic Journal*, forthcoming.
- RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.
- ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.
- ROUWENHORST, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton, NJ: Princeton University Press, 294–330.

- SAIJO, H. (2017): “The Uncertainty Multiplier and Business Cycles,” *Journal of Economic Dynamics and Control*, 78, 1–25.
- SIMS, C. A. (2002): “Solving Linear Rational Expectations Models,” *Computational Economics*, 20, 1–20.
- SMETS, F. AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97, 586–606.
- STEWART, L. AND P. MCCARTY, JR (1992): “Use of Bayesian Belief Networks to Fuse Continuous and Discrete Information for Target Recognition, Tracking, and Situation Assessment,” *Proc. SPIE*, 1699, 177–185.
- STRAUB, L. AND R. ULBRICHT (2015): “Endogenous Uncertainty and Credit Crunches,” Toulouse School of Economics Working Paper 15-604.
- TALLARINI, JR., T. D. (2000): “Risk-Sensitive Real Business Cycles,” *Journal of Monetary Economics*, 45, 507–532.
- TAYLOR, J. B. (1993): “Discretion Versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- VAN NIEUWERBURGH, S. AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53, 753–772.

A DATA SOURCES

We drew from the following data sources to estimate our VAR and New Keynesian 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. **Effective Federal Funds Rate:** Daily. Source: Board of Governors of the Federal Reserve System, Selected Interest Rates, H.15 (FRED ID: FEDFUNDS).
10. **Civilian Noninstitutional Population:** Monthly. Source: U.S. Bureau of Labour Statistics, Current Population Survey (FRED ID: CNP16OV).
11. **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:

12. **Per Capita Real GDP:** $1,000,000 \times \text{Real GDP} / \text{Population}$.
13. **Real PCE, Nondurable Goods:** $\text{Average PCE Nondurables in 2009} \times (\text{PCE Nondurables Quantity Index} / 100)$. Quantity Index FRED ID: DNDGRA3M086SBEA.
14. **Real PCE, Services:** $\text{Average PCE Nondurables in 2009} \times (\text{PCE Services Quantity Index} / 100)$. Quantity Index FRED ID: DSERRA3M086SBEA.
15. **Per Capita Real PCE:** $1,000,000 \times (\text{Real PCE Nondurables} + \text{Real PCE Services}) / \text{Population}$.
16. **Real Wage:** $100 \times \text{Average Hourly Earnings} / \text{Price Index}$.
17. **Real Investment:** $\text{Average FPI in 2009} \times (\text{FPI Quantity Index} / 100)$. Quantity Index FRED ID: A007RA3Q086SBEA.
18. **Per Capita Real Investment:** $1,000,000 \times \text{Real Investment} / \text{Population}$.

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, 15, 6, 16, 8, and 9. The observables used to estimate our nonlinear model without capital include series 12, 6, 7, 1, and 2. When we filter the data using the model with capital, we add series 18 as an observable.

B WELFARE: CONSTANT RELATIVE RISK AVERSION PREFERENCES

The representative household's preferences are given by

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

When $\gamma \neq 1$, the welfare cost, λ , satisfies

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

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

C EULER EQUATION DECOMPOSITION

The bond Euler equation is given by

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

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

$$\begin{aligned} -(\hat{i}_t + \hat{s}_t + \gamma \hat{c}_t) &= \log(E_t[\exp(-\hat{\pi}_{t+1} - \hat{g}_{t+1} - \gamma \hat{c}_{t+1})]) \\ &= \log(E_t[1 - (\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma \hat{c}_{t+1}) + \frac{1}{2}(\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma \hat{c}_{t+1})^2 - \\ &\quad \frac{1}{6}(\hat{\pi}_{t+1} + \hat{g}_{t+1} + \gamma \hat{c}_{t+1})^3 + \dots]) \\ &= \log \left(1 - (E_t[\hat{\pi}_{t+1}] + E_t[\hat{g}_{t+1}] + \gamma E_t[\hat{c}_{t+1}]) + \right. \\ &\quad \frac{1}{2} \left(E_t[\hat{\pi}_{t+1}^2] + E_t[\hat{g}_{t+1}^2] + \gamma^2 E_t[\hat{c}_{t+1}^2] + \right. \\ &\quad \left. 2E_t[\hat{\pi}_{t+1}\hat{g}_{t+1}] + 2\gamma E_t[\hat{g}_{t+1}\hat{c}_{t+1}] + 2\gamma E_t[\hat{\pi}_{t+1}\hat{c}_{t+1}] \right) - \\ &\quad \frac{1}{6} \left(E_t[\hat{\pi}_{t+1}^3] + E_t[\hat{g}_{t+1}^3] + \gamma^3 E_t[\hat{c}_{t+1}^3] + 6\gamma E_t[\hat{\pi}_{t+1}\hat{g}_{t+1}\hat{c}_{t+1}] + \right. \\ &\quad \left. 3E_t[\hat{\pi}_{t+1}^2\hat{g}_{t+1}] + 3\gamma E_t[\hat{\pi}_{t+1}^2\hat{c}_{t+1}] + 3E_t[\hat{g}_{t+1}^2\hat{\pi}_{t+1}] + \right. \\ &\quad \left. 3\gamma E_t[\hat{g}_{t+1}^2\hat{c}_{t+1}] + 3\gamma^2 E_t[\hat{c}_{t+1}^2\hat{\pi}_{t+1}] + 3\gamma^2 E_t[\hat{c}_{t+1}^2\hat{g}_{t+1}] \right) + \dots \Big), \end{aligned}$$

where the second equality follows from the Maclaurin series for $e^x = 1 + x + x^2/2 + x^3/6 + \dots$. Subsequently applying a third-order Maclaurin series to $\log(1 - x) \approx -x - x^2/2 - x^3/3$ implies

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

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

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

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

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

D VECTOR AUTOREGRESSION MODEL

The structural VAR model is given by

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

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

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

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 $j = 1, \dots, p$, $v_t = A_0^{-1} \varepsilon_t$ is a $K \times 1$ vector of shocks that has a multivariate normal distribution with zero mean and variance-covariance matrix Σ , and y is a $K \times 1$ vector of endogenous variables.

The VAR is either estimated with data generated by the baseline model or analogous variables in U.S. data and the variables are ordered as in Christiano et al. (2005). The data consists of the financial uncertainty series in Ludvigson et al. (2017), *FinU*, the per capita real consumption growth rate (nondurables + services), the GDP implicit price deflator inflation rate, the real wage growth rate (production and nonsupervisory employees), the risk premium (BAA corporate bond yield – 10 year treasury yield), and the federal funds rate. A description of our data sources is provided in [Appendix A](#). We estimate the Actual VAR with up to four lags and calculate the Bayesian information criterion (BIC). We rewrite the model as $Y_T = \beta X + U$ and calculate the least squares estimates, $\hat{\beta}$ and $\hat{\Sigma}$. For example, when $p = 4$ the parameters are $\beta = [b_0, B_1, B_2, B_3, B_4]$ and the regressors are $X = [\mathbf{1}, Y'_{T-1}, Y'_{T-2}, Y'_{T-3}, Y'_{T-4}]'$ where $Y_{T-i} = [y_{1-i}, \dots, y_{T-i}]$ and $U = [v_1, \dots, v_T]$. 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 VAR model with one lag, so we focus on that specification.

E SOLUTION METHOD

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

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

There are many ways to discretize the normally-distributed and autoregressive exogenous state variables, ν_t , $\log(\sigma_{\varepsilon,t})$, and $\log(\sigma_{v,t})$. 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, $\{\nu_{t+1}(m), \log(\sigma_{\varepsilon,t+1})(m), \log(\sigma_{v,t+1})(m), \varepsilon_{t+1}(m), v_{t+1}(m)\}_{m=1}^M$.

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 $\tilde{c}_0$ and π_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 $d \in \{1, \dots, D\}$:
 - (a) Solve for $\{\tilde{y}_t, \tilde{y}_t^{gap}, i_t^n, i_t, \tilde{w}_t, mp_t\}$ given $\tilde{c}_t = \tilde{c}_{j-1}(d)$, $\pi_t^{gap} = \pi_{j-1}^{gap}(d)$, and $\mathbf{z}_t(d)$.
 - (b) Linearly interpolate the policy functions, $\tilde{c}_{j-1}$ and π_{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, $m \in \{1, \dots, M\}$.
 - (c) Given $\{\tilde{c}_{t+1}(m), \pi_{t+1}^{gap}(m)\}_{m=1}^M$, solve for the other elements of $\mathbf{v}_{t+1}(m)$ and compute

$$\mathbb{E}[f(\mathbf{v}_{t+1}, \mathbf{v}_t(d)) | \Omega_t(d)] \approx \sum_{m=1}^M \phi(m) f(\mathbf{v}_{t+1}(m), \mathbf{v}_t(d)),$$
 - (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 $\text{maxdist}_j \equiv \max\{|\tilde{c}_j - \tilde{c}_{j-1}|, |\pi_j^{gap} - \pi_{j-1}^{gap}|\}$. When that occurs, the algorithm has converged to an approximate nonlinear solution.




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

Figure 11 shows the distribution of the absolute value of the errors in base 10 logarithms for the consumption Euler equation and the Phillips curve. For example, an error of -3 means there is a mistake of 1 consumption good for every 1,000 goods. The mean Euler equation error is -3.96 and the mean Phillips curve error is -2.32 . By construction, the errors on nodes used in

the solution algorithm are less than the convergence criterion, 10^{-6} . The larger average errors are due to linear interpolation of the policy functions with respect to the (g_t, s_t, mp_{t-1}) states. To measure the errors between the nodes, we created a new grid with a total of $D = 850,500$ nodes by increasing the number of points in the (g_t, s_t, mp_{t-1}) dimensions to $(15, 15, 15)$. We used the same number of points in the $(\nu_t, \log(\sigma_{\varepsilon,t}), \log(\sigma_{v,t}))$ 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.

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

F 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.

F.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, $\theta \equiv \{\gamma, \varphi, \phi_\pi, \phi_y, \bar{g}, \bar{\pi}, \rho_i, \rho_g, \rho_s, \rho_{\sigma_\varepsilon}, \rho_{\sigma_v}, \sigma_\nu, \sigma_\varepsilon, \sigma_v, \sigma_\xi, \sigma_\zeta\}$.
2. The vector of N_x observables consists of per capita real GDP, $RGDP/CNP$, the GDP deflator, DEF , the 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 $t \in \{1, \dots, T\}$. 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 $i \in \{1, \dots, N_m\}$, where $N_m = 5,000$, apply the following steps:

i. Draw $\hat{\theta}_i$ from the joint prior distribution and calculate its density value:

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

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 E](#). 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 F.2](#) to obtain the log-likelihood value for the model, $\log \ell_i^{\text{model}}$, otherwise repeat step 3(a)i.

iv. The posterior log-likelihood is $\log \ell_i^{\text{post}} = \log \ell_i^{\text{prior}} + \log \ell_i^{\text{model}}$

(b) Calculate $\max(\log \ell_1^{\text{post}}, \dots, \log \ell_{N_m}^{\text{post}})$ and find the corresponding parameter vector, $\hat{\theta}_0$.

4. Approximate the covariance matrix for the joint posterior distribution of the parameters, Σ , which is used to draw candidates during the preliminary Metropolis-Hastings step.

(a) Locate the draws with a likelihood in the top decile. Stack the $N_{m,\text{sub}} = (1 - p)N_m$ draws in a $N_{m,\text{sub}} \times N_e$ matrix, $\hat{\Theta}$, and define $\tilde{\Theta} = \hat{\Theta} - \sum_{i=1}^{N_{m,\text{sub}}} \hat{\theta}_{i,j} / N_{m,\text{sub}}$.

(b) Calculate $\Sigma = \tilde{\Theta}'\tilde{\Theta} / N_{m,\text{sub}}$ 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 $i \in \{0, \dots, N_d\}$, where $N_d = 25,000$, perform the following steps:

i. Draw a candidate vector of parameters, $\hat{\theta}_i^{\text{cand}}$, where

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

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

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

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

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

v. Accept or reject the candidate draw according to

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

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 $\log \ell_i^{\text{cand}} = \log \ell_i^{\text{prior}} + \log \ell_i^{\text{model}}$.

- (b) Burn the first $N_b = 5000$ draws and use the remaining sample to calculate the mean draw, $\bar{\theta}^{preMH} = \sum_{i=N_b+1}^{N_{preMH}} \hat{\theta}_i$, 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 p th 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.

F.2 ADAPTED PARTICLE FILTER The following steps outline the filter:

1. Initialize the filter by drawing $\mathbf{e}_{t,p} = \{\nu_{t,p}, \varepsilon_{t,p}, v_{t,p}, \xi_{t,p}, \zeta_{t,p}\}_{t=-24}^0$ for all $p \in \{0, \dots, N_p\}$ 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 $p \in \{1, \dots, N_p\}$ 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 ($\tilde{y}_t^{gdp}$), inflation (π), the policy rate (i), consumption growth uncertainty, and risk premium uncertainty according to $\hat{\mathbf{x}}_t^{model} = [\log(g_t \tilde{y}_t^{gdp} / \tilde{y}_{t-1}^{gdp}), \log(\pi_t), \log(i_t), (U_{cg,t} - \mu_{U_{cg}})/\sigma_{U_{cg}}, (U_{s,t} - \mu_{U_s})/\sigma_{U_s}]$. 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 $H \equiv \text{diag}(\sigma_{me,cg}^2, \sigma_{me,\pi}^2, \sigma_{me,i}^2, \sigma_{me,macrou}^2, \sigma_{me,finu}^2)$ 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, $\mu_{t,p} = \hat{\mathbf{x}}_{t,p}^{model} - \hat{\mathbf{x}}_{t,p}^{data}$. The unnormalized weight on particle p is given by

$$\omega_{t,p} = \frac{p(\mu_t|\mathbf{z}_{t,p})p(\mathbf{z}_{t,p}|\mathbf{z}_{t-1,p})}{g(\mathbf{z}_{t,p}|\mathbf{z}_{t-1,p}, \hat{\mathbf{x}}_{t,p}^{data})} \propto \frac{\exp(-\mu_{t,p}' H^{-1} \mu_{t,p} / 2) \exp(-\mathbf{e}_{t,p}' \mathbf{e}_{t,p} / 2)}{\exp(-(\mathbf{e}_{t,p} - \bar{\mathbf{e}}_t)'(\mathbf{e}_{t,p} - \bar{\mathbf{e}}_t) / 2)}.$$

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 $\ell_t^{model} = \sum_{p=1}^{N_p} \omega_{t,p} / N_p$.

- (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 $t \in \{1, \dots, T\}$. The log-likelihood is then $\log \ell^{model} = \sum_{t=1}^T \log \ell_t^{model}$.

G NEW KEYNESIAN MODEL ESTIMATION DIAGNOSTICS

	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. The vertical dashed line represents the expected ZLB duration.




Figure 13: Trace plots. We obtained 100,000 draws from each posterior distribution and kept every 100th draw.




Figure 14: Prior (solid lines) and posterior kernel (dashed lines) densities of the estimated parameters.




Figure 15: Time paths of the data (dashed line) and the median filtered series from the baseline model (solid line).




Figure 16: Median paths of the estimated shocks normalized by their respective posterior mean standard deviation.