Complementarity and Macroeconomic Uncertainty*

Tyler Atkinson Michael Plante

Alexander W. Richter Nathaniel A. Throckmorton

March 24, 2021

ABSTRACT

Macroeconomic uncertainty regularly fluctuates in the data. Theory suggests complementarity between capital and labor inputs in production can generate time-varying endogenous uncertainty because the concavity in the production function influences how output responds to productivity shocks in different states of the economy. This paper examines whether complementarity is a quantitatively significant source of time-varying endogenous uncertainty by estimating a nonlinear real business cycle model with a constant elasticity of substitution production function and exogenous volatility shocks. When matching labor share and uncertainty moments, we find at most 16% of the volatility of uncertainty is endogenous. An estimated model without exogenous volatility shocks can endogenously generate all of the empirical variation in uncertainty, but only at the expense of significantly overstating the volatility of the labor share.

Keywords: State-Dependence; Stochastic Volatility; CES Production; Endogenous Uncertainty
JEL Classifications: C12; C13; D81; E32; E37

*Atkinson, Plante, and Richter, Research Department, Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (tyler.atkinson@dal.frb.org; michael.plante@dal.frb.org; alex.richter@dal.frb.org); Throckmorton, Department of Economics, William & Mary, P.O. Box 8795, Williamsburg, VA 23187 (nat@wm.edu). We thank Loukas Karabarbounis and two anonymous referees for very useful comments that improved the paper. We also thank Jarod Coulter for excellent research assistance. This work was supported by computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas. The views expressed in this paper are our own and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 INTRODUCTION

Macroeconomic uncertainty regularly fluctuates in the data. The literature typically accounts for these fluctuations in business cycle models using shocks to the variance of an exogenous variable, such as productivity, while holding its conditional mean fixed. That produces estimates for the responses of output to an *exogenous* increase in uncertainty, but it is silent about whether uncertainty *endogenously* fluctuates over time in response to first-moment shocks and the state of the economy.

Straub and Ulbricht (2019) show that variance-preserving shifts in the productivity distribution across firms cause endogenous fluctuations in the cross-sectional dispersion of output when capital and labor are gross complements in production because each firm's output becomes a concave function in productivity.[1] We build on their work in two ways. First, we show that complementarity also generates endogenous fluctuations in aggregate uncertainty—the expected forecast error volatility of future aggregate output. Complementarity can endogenously generate fluctuations in aggregate uncertainty because the concavity in the production function influences how output responds to first moment shocks in different states of the economy. For example, a positive labor productivity shock generates a larger change in output when the capital-to-labor ratio is high compared to when the capital-labor ratio is low because complementarity increases the marginal product of labor when capital is abundant. The differences in the responses imply that forecasts for output and hence the level of uncertainty also depend on the initial state, creating time-varying endogenous uncertainty.

Second, we examine whether the endogenous variation in aggregate uncertainty is quantitatively significant. To conduct our analysis, we estimate a real business cycle model with a constant elasticity of substitution (CES) production function, exogenous volatility shocks to productivity, and real frictions in the form of habit persistence in consumption and investment adjustment costs. We estimate the nonlinear model using a global solution method and a simulated method of moments that targets uncertainty, real activity, and labor share moments. The model generates movements in aggregate uncertainty through exogenous volatility shocks and endogenous fluctuations due to complementarity in the production function. Allowing for exogenous and endogenous movements in uncertainty lets the data determine the relative importance of the two mechanisms.

We begin by looking at the predictions of the model without exogenous volatility shocks to focus on the effects of complementarity. We find that the model without exogenous volatility can successfully match either the uncertainty moments or the labor share moments, but not both. When the labor share moments are excluded from the set of empirical targets, the model is able to endogenously generate all of the empirical variation in aggregate uncertainty, confirming the theory that gross complementarity can generate significant variation in aggregate uncertainty. However, the estimated CES is only 0.14, well below almost all estimates in the literature, and the volatility

[1] Ilut et al. (2018) show concave hiring rules generate countercyclical volatility in response to productivity shocks.

of the labor share is about six times higher than it is in the data. Alternatively, when we exclude the uncertainty moments, the labor share moments pin down the estimated value of the CES, which increases to 0.49. However, the higher estimate leads to far less concavity in the production function, so the model is able to generate only about 20% of the volatility of aggregate uncertainty in the data.

We then turn to our baseline model with exogenous volatility shocks, which offers a competing mechanism for the variation in uncertainty. The model jointly matches the volatility of uncertainty and labor share dynamics, and the CES estimate is the same as when we match labor share moments without exogenous volatility shocks. However, a forecast error variance decomposition reveals that endogenous uncertainty explains at most 16% of the variation in aggregate uncertainty. A lower CES would increase the strength of the endogenous uncertainty channel, but it would also cause the model to overstate the empirical volatility of the labor share. Despite the theoretical appeal of complementarity, these results show that it plays a limited role in generating the fluctuations in macro uncertainty. However, the model underpredicts the cyclicality of uncertainty, which suggests there is a potentially important role for other sources of time-varying endogenous uncertainty.

We also find the dynamics of real activity under a Cobb-Douglas production function are similar to those implied by our estimated CES production function. Under a Cobb-Douglas production function with constant returns to scale, the labor share is constant, which motivates using a more general CES production function. However, we find the empirical volatility of the labor share leads to a CES estimate that is not small enough to generate meaningful differences in the dynamics implied by the two production functions. This suggests a Cobb-Douglas production function provides a reasonable approximation of business cycle dynamics, even when accounting for nonlinearities.

Related Literature A large literature has estimated the elasticity of substitution between capital and labor. Klump et al. (2012) surveys reduced-form estimates of the CES. Across 17 studies, the average CES is 0.58, well below the unit elasticity implied by Cobb-Douglas production. Oberfield and Raval (2021), who account for substitution across plants, firms, and manufacturing industries, estimate a CES for the U.S. manufacturing sector between 0.5 and 0.7. We obtain a similar CES estimate of 0.49 in our baseline model using a different framework and identification strategy. We pin down the CES in a one-sector business cycle model using quarterly variation in the labor share and macro uncertainty. When we do not target the labor share, the CES falls to 0.14. This estimate is consistent with Cantore et al. (2015), who estimate a New Keynesian model with a CES production function but without targeting labor share dynamics. While these estimates generate more endogenous variation in uncertainty, they cause the model to over-predict the volatility of the labor share.

Several papers have also estimated a CES above unity (Eden and Gaggl, 2018; Hubmer, 2020; Karabarbounis and Neiman, 2014). This work identifies the CES by focusing on long-run trends, such as the decline in the labor share or the relative price of investment goods. Our estimation procedure focus on business cycle dynamics and therefore matches the short-run properties of the

de-trended labor share.[2] However, the two sets of estimates are not necessarily inconsistent, as it is eminently plausible that the long-run elasticity is higher than the short-run elasticity we estimate.

Several papers study alternative sources of endogenous uncertainty. One segment emphasizes the role of a financial sector, where the severity and duration of financial crises are stochastic. Most papers focus on crises that result from financial frictions, collateral constraints, or the zero lower bound constraint on the nominal interest rate (Brunnermeier and Sannikov, 2014; He and Krishnamurthy, 2019; Mendoza, 2010; Plante et al., 2018), while others incorporate the role of firm default (Arellano et al., 2019; Gourio, 2014; Navarro, 2014). A separate segment of the literature examines the implications of incomplete information. Some 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, 2020; Straub and Ulbricht, 2015). In these models, adverse shocks under asymmetric learning reduce economic activity and make it harder for agents to learn about the economy, amplifying equilibrium dynamics. Finally, recent papers have shown that search and matching frictions can generate uncertainty (Bernstein et al., 2020; Ilut et al., 2018). Consistent with many of these mechanisms, uncertainty in our model occurs at business cycle turning points. One major benefit of our mechanism is that it is easy to incorporate into any model.

There are also papers that study the effects of exogenous volatility shocks. For example, this literature has examined volatility shocks to technology (Bloom, 2009; Leduc and Liu, 2016), fiscal policy (Born and Pfeifer, 2014; Fernández-Villaverde et al., 2015), monetary policy (Mumtaz and Zanetti, 2013), and the real interest rate (Fernández-Villaverde et al., 2011). We build on this literature by examining the role of exogenous volatility shocks in driving aggregate uncertainty. Our model accounts for exogenous and endogenous sources of uncertainty, and our estimation disciplines the stochastic processes by matching the volatilities of both real activity and aggregate uncertainty. This allows the data to decide which source is driving the empirical variation in uncertainty.

Our paper also makes two technical contributions to the literature. First, we estimate our model using a global solution method, which is crucial to calculate aggregate uncertainty and account for state-dependance. Cantore et al. (2015) estimate a similar model but use a linear solution method and focus on issues besides uncertainty. Similarly, the exogenous uncertainty literature typically applies third-order perturbation methods. Second, we are the first to discipline uncertainty dynamics in a business cycle model by directly linking to uncertainty dynamics in the data. We rely on the real uncertainty series in Ludvigson et al. (2020) over other popular measures of uncertainty (e.g., realized volatility, indexes based on keywords in print or online media, survey-based forecast dispersion) since it is based on the same statistic we use to measure uncertainty in our structural model.

The paper proceeds as follows. Section 2 explains the underlying nonlinearity in a simplified

[2]Business cycle models typically do not separate short- and long-run elasticities of substitution. A notable exception is Leon-Ledesma and Satchi (2019), who embed a technological choice problem into a real business cycle model.

setting. Section 3 presents our quantitative model. Section 4 re-examines the source of the nonlinearity in our quantitative model. Section 5 presents our estimation results, and Section 6 concludes.

2 UNDERLYING MECHANISM

This section first defines our measure of uncertainty and then shows complementarity can generate time-varying endogenous uncertainty even with independent inputs. The benefit of this example is that it shows it is possible to incorporate our proposed mechanism into any equilibrium framework.

2.1 UNCERTAINTY DEFINITION

Following Plante et al. (2018), uncertainty is defined as the expected volatility of a variable in the model. The same definition is used in empirical work (e.g., Jurado et al. 2015; Ludvigson et al. 2020). It is possible to calculate uncertainty over any horizon, but we focus on the 1-quarter ahead forecast error. For variable x, aggregate uncertainty is given by

$$\mathcal{U}^x_{t,t+1} \equiv \sqrt{E_t[(x_{t+1} - E_t[x_{t+1}])^2]}, \tag{1}$$

where E_t is the mathematical expectation operator conditional on information available at time t. A key aspect of this definition is that it removes the predictable component, $E_t[x_{t+1}]$, from a 1-period ahead forecast of x. While x can represent any variable in the model, the discussion in this paper is centered on log output ($\hat{y}_t$) and log output growth ($\hat{y}^g_t = \hat{y}_{t+1} - \hat{y}_t$). The first useful result is that the uncertainty surrounding these two endogenous variables are equal. This result occurs because $\hat{y}_t$ is known at time t and therefore cancels out when removing the predictable component.

2.2 ENVIRONMENT

A common assumption is that output is produced according to a Cobb-Douglas production function with constant returns to scale. In this case, log output, $\hat{y}_t$, is given by

$$\hat{y}_t = \alpha \hat{x}_{1,t} + (1-\alpha)\hat{x}_{2,t},$$

where hats denote logs, $\hat{x}_{i,t}$, $i \in \{1, 2\}$, are random variables that evolve according to

$$\hat{x}_{i,t} = (1-\rho_i)\log \bar{x}_i + \rho_i \hat{x}_{i,t-1} + \nu_i \varepsilon_{i,t},\ 0 \leq \rho_i < 1,\ \varepsilon_i \sim \mathbb{N}(0,1), \tag{2}$$

and bars denote steady-states. Uncertainty about log output 1-period ahead is constant and given by

$$\mathcal{U}^{\hat{y}}_{t,t+1} = \sqrt{\alpha^2 \nu_1^2 + (1-\alpha)^2 \nu_2^2}. \tag{3}$$

This expression shows the conditional volatility of log output is a weighted average of the variance of each shock, with weights equal to the squared share of each variable in the production function.

While this example is stylized, it illustrates the important role of the production function and the assumptions underlying the random variables. There are four key assumptions behind the

results: (1) Log-linearity of the production function; (2) Constant weights on the variances; (3) Log-linearity of the stochastic processes; (4) Constant conditional variances of the shocks (ε_i). Deviations from any of these assumptions would create time-varying uncertainty about log output. Relaxing the last assumption by making the variance of an exogenous variable stochastic is the most common way to introduce time-varying uncertainty. Since it does not depend on any equilibrium feature of the model, we refer to it as time-varying *exogenous* uncertainty. This paper focuses on the effects of relaxing the first assumption, which creates time-varying *endogenous* uncertainty.

An alternative way to think about the importance of assumptions (1)-(4) is to note that they imply the response of log output to an unexpected shock is *not* state-dependent. Mathematically,

$$\partial\hat{y}_t/\partial\varepsilon_{1,t} = \alpha\nu_1 \quad \text{and} \quad \partial\hat{y}_t/\partial\varepsilon_{2,t} = (1-\alpha)\nu_2.$$

The impact of a shock to log output is the same in every period and for any state of the economy.

2.3 UNCERTAINTY APPROXIMATIONS Taylor expansions are useful for approximating moments of a variable when the variable is a nonlinear function of other random variables. For example, a second-order approximation of the variance of some generic function $f(x)$ is given by

$$\text{Var}[f(x)] \approx (f'(E[x]))^2\nu_x^2, \tag{4}$$

where $f'(E[x])$ is the first derivative of $f(x)$ evaluated at the unconditional mean of x and ν_x^2 is the unconditional variance of x. Extensions to the multivariate case are straightforward as long as the random variables are uncorrelated. In that case, the variance is approximated by a weighted average of the unconditional variances of the random variables, with the weights equal to the first partial derivatives of the function evaluated at the unconditional means of the random variables.

Approximating $\hat{y}$ around the conditional mean of $\hat{x}_1$ and $\hat{x}_2$ in the Cobb-Douglas case implies

$$\text{Var}_t[\hat{y}_{t+1}] \approx \alpha^2\nu_1^2 + (1-\alpha^2)\nu_2^2.$$

This approximation is exact due to the log-linearity of the production function and its two inputs. Taking a square root of the conditional variance produces our measure of uncertainty defined in (1).

Now consider a more general CES production function given by

$$\hat{y}_t = \tfrac{\sigma}{\sigma-1}\ln\left(\alpha\exp(\tfrac{\sigma-1}{\sigma}\hat{x}_{1,t}) + (1-\alpha)\exp(\tfrac{\sigma-1}{\sigma}\hat{x}_{2,t})\right),$$

where σ is the elasticity of substitution. The approximation for the conditional variance implies

$$\text{Var}_t[\hat{y}_{t+1}] \approx (f_1(E_t\hat{x}_{1,t+1}, E_t\hat{x}_{2,t+1}))^2\nu_1^2 + (f_2(E_t\hat{x}_{1,t+1}, E_t\hat{x}_{2,t+1}))^2\nu_2^2,$$

where $f_i(\cdot)$ is the partial derivative of $\hat{y}$ with respect to $\hat{x}_i$. As in the Cobb-Douglas case, the condi-

tional variance for log output is a weighted average of the conditional variances of the two random variables. However, the weights are now time-varying due to the state-dependent effects of the shocks. This feature generates time-varying endogenous uncertainty under CES production.[3]

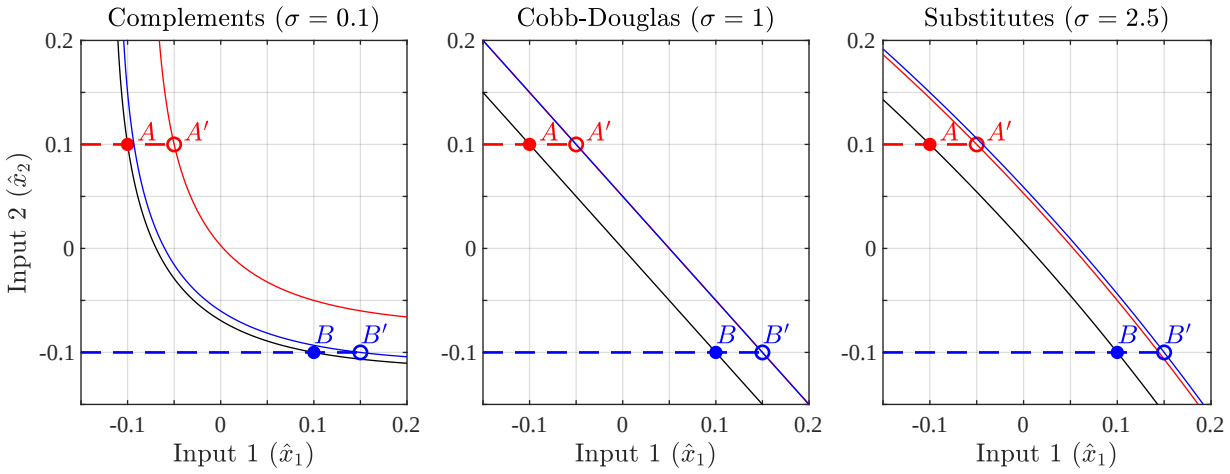

Figure 1: Isoquants. Points A' and B' are shifted right by 0.05 units of output and α is set to 0.5.

2.4 ISOQUANTS Figure 1 plots isoquants to help visualize how complementarity generates state-dependent responses of output and time-varying uncertainty. The inner solid line traces out combinations of inputs that produce a given level of output. Points A and B are initial conditions for the same level of output but with a high level of one input relative to the other. Points A' and B' are both shifted to the right by 0.05 units of output, reflecting an increase in $\hat{x}_1$ with no change in $\hat{x}_2$.

In the left panel where the inputs are strong complements ($\sigma = 0.1$), the shift results in a large increase in output from point A to A', where $\hat{x}_1$ was initially low relative to $\hat{x}_2$ and the isoquants are steep. This reflects the effect of complementarity, in that a high level of one input increases the marginal product of the other. Alternatively, output is largely unchanged from point B to B', as $\hat{x}_1$ is already elevated relative to $\hat{x}_2$ and the isoquants are relatively flat. The middle panel shows the same exercise for Cobb-Douglas production. Since the isoquants are linear, increases in $\hat{x}_1$ raise output by the same amount regardless of $\hat{x}_2$. In the right panel, where the inputs are substitutes ($\sigma = 2.5$), the isoquants are concave rather than convex, meaning an elevated level of one input relative to the other increases its marginal product—the opposite of what happens under complementarity.

The partial derivative of $\hat{y}$ with respect to $\hat{x}_1$ is $1/(1 + \frac{1-\alpha}{\alpha} \exp(\frac{\sigma-1}{\sigma}(\hat{x}_2 - \hat{x}_1))$, so the response of output depends on the elasticity of substitution and the log difference in inputs (or the ratio of the inputs when written in levels). The nonlinearity stems from the concavity in the production function. Differences in the responses of output imply that forecasts for output and hence the level

[3] Appendix A provides more information about the approximations. See Benaroya et al. (2005) for further details.

of output uncertainty also depends on the initial state of the economy. Under Cobb-Douglas production, output uncertainty is time-invariant because the response of output is not state-dependent.[4]

2.5 EXOGENOUS VOLATILITY SHOCKS

The discussion so far has focused on the implications of relaxing the log-linearity of the production function (Assumption 1). As previously mentioned, it is possible to generate time-varying *exogenous* uncertainty by relaxing the assumption of constant conditional variances of the shocks (Assumption 4). To consider the impact, generalize (2) to

$$\hat{x}_{i,t} = (1-\rho_i)\log\bar{x}_i + \rho_i\hat{x}_{i,t-1} + \nu_{i,t}\varepsilon_{i,t},\ 0 \leq \rho_i < 1,\ \varepsilon_i \sim \mathbb{N}(0,1),$$

where $\nu_{i,t}$ is uncorrelated with $\varepsilon_{i,t}$ and evolves according to a stationary autoregressive process.

In this case, it is straightforward to show that uncertainty in the Cobb-Douglas case becomes

$$\mathcal{U}^{\hat{y}}_{t,t+1} = \sqrt{\alpha^2 E_t[\nu^2_{1,t+1}] + (1-\alpha)^2 E_t[\nu^2_{2,t+1}]}, \tag{5}$$

which is time-varying. In the CES case our approximation for the conditional variance is given by

$$\mathrm{Var}_t[\hat{y}_{t+1}] \approx (f_1(E_t\hat{x}_{1,t+1}, E_t\hat{x}_{2,t+1}))^2 E_t[\nu^2_{1,t+1}] + (f_2(E_t\hat{x}_{1,t+1}, E_t\hat{x}_{2,t+1}))^2 E_t[\nu^2_{2,t+1}].$$

The conditional volatility of output varies across time not only because the effects of the shocks are state-dependent, but also because the conditional variances of ν_i fluctuate over time. Thus, it is possible to decompose the variation in aggregate uncertainty into endogenous and exogenous sources.

3 BUSINESS CYCLE MODEL

To conduct our quantitative analysis we use a real business cycle model with complementary inputs in production, real frictions, and a common shock to the volatility of capital and labor productivity.

Firms The representative firm produces output according to a CES production function given by

$$y_t = \begin{cases} y_0\left(\alpha(z_{k,t}k_{t-1}/k_0)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)(\bar{g}^t z_{n,t}n_t/n_0)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, & \sigma \neq 1, \\ y_0(z_{k,t}k_{t-1}/k_0)^{\alpha}(\bar{g}^t z_{n,t}n_t/n_0)^{1-\alpha}, & \sigma = 1, \end{cases} \tag{6}$$

where k is the capital stock, n is the labor supply, and $\bar{g}$ is the average growth rate of labor-augmenting productivity. z_i, $i \in \{k, n\}$, is a stationary productivity shock that evolves according to

$$\hat{z}_{i,t} = \rho_{zi}\hat{z}_{i,t-1} + \nu_{zi}\exp(\hat{\nu}_t)\varepsilon_{zi,t},\ 0 \leq \rho_{zi} < 1,\ \varepsilon_{zi} \sim \mathbb{N}(0,1), \tag{7}$$

$$\hat{\nu}_t = \rho_\nu\hat{\nu}_{t-1} + \nu_\nu\varepsilon_{\nu,t},\ 0 \leq \rho_\nu < 1,\ \varepsilon_\nu \sim \mathbb{N}(0,1). \tag{8}$$

[4] When the inputs adjust proportionally, such as under Leontief production with flexible inputs, the nonlinearity and state-dependence vanishes. However, as long as there is a non-unitary elasticity of substitution and the log difference in inputs fluctuates over time, the curvature in the production function will generate time-varying endogenous uncertainty.

Following the literature, we include normalizing constants, y_0, n_0 and k_0, to ensure that α is the cost-share of capital under both a Cobb-Douglas ($\sigma = 1$) and CES ($\sigma \neq 1$) production function.[5]

The firm chooses $\{n_t, k_{t-1}\}$ to maximize $y_t - w_t n_t - r_t^k k_{t-1}$ subject to the production function, where w is the wage rate and r^k is the rental rate on capital. The optimality conditions are given by

$$r_t^k = \alpha(y_0 z_{k,t}/k_0)^{(\sigma-1)/\sigma}(y_t/k_{t-1})^{1/\sigma}, \tag{9}$$

$$w_t = (1-\alpha)(y_0 \bar{g}^t z_{n,t}/n_0)^{(\sigma-1)/\sigma}(y_t/n_t)^{1/\sigma}. \tag{10}$$

The labor share, $s_{n,t} \equiv w_t n_t / y_t$, is time-varying when $\sigma \neq 1$, but always equals $1-\alpha$ when $\sigma = 1$.

Households A representative household chooses $\{c_t, n_t, k_t, x_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \beta^t [\frac{(c_t - h c_{t-1})^{1-\gamma}}{1-\gamma} - \chi \frac{n_t^{1+\eta}}{1+\eta}]$, where c is consumption, x is investment, χ determines the steady-state labor supply, γ is the coefficient of relative risk aversion, $1/\eta$ is the Frisch elasticity of labor supply, and h is the degree of external habit persistence. These choices are constrained by

$$c_t + x_t = w_t n_t + r_t^k k_{t-1},$$

$$k_t = (1-\delta)k_{t-1} + x_t(1 - \varphi^x (x_t^{gap} - 1)^2/2),$$

where $x_t^{gap} = x_t/(\bar{g} x_{t-1})$ is investment growth, φ^x scales the size of the investment adjustment cost, and δ is the depreciation rate on capial. The household's optimality conditions are given by

$$w_t = \chi n_t^{\eta}(c_t - h c_{t-1})^{\gamma}, \tag{11}$$

$$1 = q_t[1 - \varphi^x (x_t^{gap} - 1)(3x_t^{gap} - 1)/2] + \bar{g}\varphi^x E_t[m_{t+1} q_{t+1} (x_{t+1}^{gap})^2 (x_{t+1}^{gap} - 1)], \tag{12}$$

$$q_t = E_t[m_{t+1}(r_{t+1}^k + (1-\delta)q_{t+1})], \tag{13}$$

where q is Tobin's q and $m_{t+1} = \beta((c_t - h c_{t-1})/(c_{t+1} - h c_t))^{\gamma}$ is the stochastic discount factor.

Equilibrium The resource constraint is given by $c_t + x_t = y_t$. Due to the trend in labor productivity, we detrend the model by defining $\tilde{x}_t \equiv x_t/\bar{g}^t$. Appendix B provides the equilibrium system.

Competitive equilibrium consists of sequences of quantities $\{\tilde{y}_t, \tilde{k}_t, \tilde{c}_t, n_t, \tilde{x}_t, m_{t+1}, x_t^{gap}\}_{t=0}^{\infty}$, prices $\{r_t^k, \tilde{w}_t, q_t\}_{t=0}^{\infty}$, and exogenous variables $\{z_{k,t}, z_{n,t}, \nu_t\}_{t=0}^{\infty}$ that satisfy the detrended equilibrium system, given initial conditions $\{\tilde{c}_{-1}, \tilde{k}_{-1}, \tilde{x}_{-1}, z_{k,0}, z_{n,0}, \nu_0\}$ and shocks $\{\varepsilon_{zk,t}, \varepsilon_{zn,t}, \varepsilon_{\nu,t}\}_{t=1}^{\infty}$.

4 REINSPECTING THE MECHANISM

This section builds on the intuition from the production economy in Section 2. We begin by solving our quantitative model under conditions that permit an analytical solution. We then report generalized impulse responses and distributions of future output in different states of the economy.

[5] Leon-Ledesma et al. (2010) show that a normalized CES production function improves identification, in addition to its theoretical benefits. See Klump et al. (2012) for more information about the effects of the normalizing constants.

4.1 ANALYTICAL RESULTS To analytically solve our model, we assume capital fully depreciates each period ($\delta = 1$), turn off the real frictions ($h = \varphi^x = 0$), remove the capital productivity and volatility shocks ($\nu_{zk} = \nu_\nu = 0$), and set $\gamma = 1/\sigma$. In this case, guessing that $\tilde{c}_t = \theta \tilde{y}_t$ implies

$$\tilde{c}_t = (1 - (\alpha\beta)^\sigma (y_0/(\bar{g}k_0))^{\sigma-1})\tilde{y}_t, \tag{14}$$

$$\tilde{k}_t = (\alpha\beta)^\sigma (y_0/(\bar{g}k_0))^{\sigma-1}\tilde{y}_t, \tag{15}$$

$$n_t = \left[\frac{1-\alpha}{\chi}\left(\frac{1}{1-(\alpha\beta)^\sigma (y_0/(\bar{g}k_0))^{\sigma-1}}\right)^{1/\sigma}\left(\frac{y_0 z_{n,t}}{n_0}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{1+\eta\sigma}}, \tag{16}$$

which confirms the consumption-to-output ratio is constant. Labor depends on labor productivity.

Cobb-Douglas Case When $\sigma = 1$, uncertainty about 1-period ahead log output is given by

$$\mathcal{U}^{\hat{y}}_{t,t+1} = (1-\alpha)\sqrt{E_t[(\hat{z}_{n,t+1} - E_t[\hat{z}_{n,t+1}])^2]} = (1-\alpha)\nu_{zn}.$$

Only one term appears because α multiplies the capital stock, which is known at time t. If we included both the capital- and labor-augmenting productivity shocks, uncertainty would equal $\sqrt{\alpha\nu_{zk}^2 + (1-\alpha)\nu_{zn}^2}$, which is the same as we reported in (3) when both inputs were independent.

This example satisfies the assumptions in Section 2.2. The production function and shock process are log-linear, the shock is homoskedastic, and the cost-shares are constant. Furthermore, $\partial\hat{y}_t/\partial\varepsilon_{zn,t} = (1-\alpha)\nu_{zn}$. Thus, a model with Cobb-Douglas production, log utility, and full depreciation cannot generate time-varying output uncertainty or state-dependent impulse responses.

CES Case When $\sigma \neq 1$, log output and log labor in period $t+1$ are given by

$$\hat{y}_{t+1} = \hat{y}_0 + \tfrac{\sigma}{\sigma-1}\ln\left(\alpha\exp(\tfrac{\sigma-1}{\sigma}(\hat{k}_t - \hat{g} - \hat{k}_0)) + (1-\alpha)\exp(\tfrac{\sigma-1}{\sigma}(\hat{z}_{n,t+1} + \hat{n}_{t+1} - \hat{n}_0))\right), \tag{17}$$

$$\hat{n}_{t+1} = \tfrac{\sigma}{1+\eta\sigma}\left(\hat{\kappa}_n + \tfrac{\sigma-1}{\sigma}\hat{z}_{n,t+1}\right), \tag{18}$$

where $\hat{\kappa}_n$ collects the constant terms in the labor policy function. Combining the policy functions provides an equation for log output that is solely a function of labor productivity. In contrast with the Cobb-Douglas case, output is no longer log-linear. As a result, deriving an exact analytical expression for log output uncertainty is no longer possible, so we turn to analytical approximations.

To apply the Taylor approximation in (4), first combine (17) and (18) and reorganize to obtain

$$\hat{y}_{t+1} = \hat{y}_0 + \tfrac{\sigma}{\sigma-1}\ln\left(\alpha\exp(\tfrac{\sigma-1}{\sigma}(\hat{k}_t - \hat{g} - \hat{k}_0)) + (1-\alpha)\exp(\hat{\kappa}_y + \tfrac{(\sigma-1)(1+\eta)}{1+\eta\sigma}\hat{z}_{n,t+1})\right),$$

where $\hat{\kappa}_y$ collects the constant terms. The conditional variance of log output is approximated by

$$\mathrm{Var}_t[\hat{y}_{t+1}] \approx [y_{\hat{z}}(\hat{k}_t, E_t\hat{z}_{n,t+1})]^2\nu_{zn}^2.$$

Differentiation implies the weight, $y_{\hat{z}}$, increases with $\hat{z}_n$ when $\sigma > 1$ and decreases when $\sigma < 1$, so uncertainty is time-varying under CES production and constant in the Cobb-Douglas case ($\sigma = 1$).[6] Importantly, when $\sigma < 1$, output becomes a concave function of capital and labor productivity and the conditional variance depends on their relative levels as explained in the discussion of Figure 1.

Introducing exogenous volatility shocks would change the variance formulas in the same way as Section 2.5, so the variances would become time-varying in the Cobb-Douglas and CES cases.

4.2 NUMERICAL RESULTS We now turn to numerical methods to permit solutions under more general conditions. The literature often linearizes business cycle models. While this approach works well for most applications, it prevents analysis of time-varying uncertainty. Therefore, we solve the nonlinear model using the policy function iteration algorithm described in Richter et al. (2014). The algorithm is based on the theoretical work on monotone operators in Coleman (1991).

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. We then minimize the Euler equation errors on every node in the 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. Once the algorithm converges, we use the nonlinear solution and numerical integration to generate a policy function for log output uncertainty. See Appendix C for a detailed description of the algorithm.

Impulse Responses To generate time-varying endogenous uncertainty, the responses to first-moment shocks must vary across states of the economy. Figure 2 shows the state-dependence by plotting generalized impulse responses to a positive 2 standard deviation labor productivity shock when $\sigma \in \{0.1, 1, 2.5\}$. These are the same values we used in Section 2.4 to illustrate the nonlinear effects of complementarity. Our quantitative analysis, discussed in the next section, estimates the CES. To simplify the state of the economy, we continue to exclude the real frictions and only include a labor productivity shock, though these assumptions do not affect our qualitative results.[7]

Following Koop et al. (1996), the impulse response of variable x_{t+h} over horizon h is given by

$$\mathcal{G}_t(x_{t+h}|\varepsilon_{zn,t+1} = 2, \mathbf{z}_t) = E_t[x_{t+h}|\varepsilon_{zn,t+1} = 2, \mathbf{z}_t] - E_t[x_{t+h}|\mathbf{z}_t],$$

where $\mathbf{z}_t$ is a vector of initial states and 2 is the size of the labor productivity shock. The conditional expectations are computed based on the mean path from 20,000 simulations of the model. Each

[6] Define $\hat{f}_{1,t} = \frac{\sigma-1}{\sigma}(\hat{k}_t - \hat{g} - \hat{k}_0)$ and $\hat{f}_{2,t+1} = \hat{\kappa}_n + \frac{(\sigma-1)(1+\eta)}{1+\eta\sigma}\hat{z}_{n,t+1}$. Then the comparative statics are given by

$$\hat{y}_{\hat{z}}(\hat{k}_t, \hat{z}_{n,t+1}) = \frac{\sigma(1+\eta)}{1+\eta\sigma}\frac{(1-\alpha)\exp(\hat{f}_{2,t+1})}{\alpha\exp(\hat{f}_{1,t})+(1-\alpha)\exp(\hat{f}_{2,t+1})}, \quad \hat{y}_{\hat{z}\hat{z}}(\hat{k}_t, \hat{z}_{n,t+1}) = \frac{\sigma(\sigma-1)(1+\eta)^2}{(1+\eta\sigma)^2}\frac{\alpha(1-\alpha)\exp(\hat{f}_{1,t})\exp(\hat{f}_{2,t+1})}{(\alpha\exp(\hat{f}_{1,t})+(1-\alpha)\exp(\hat{f}_{2,t+1}))^2},$$

where $\hat{y}_{\hat{z}}(\hat{k}_t, \hat{z}_{n,t+1}) > 0$. Therefore, it is easy to see that the sign of $\hat{y}_{\hat{z}\hat{z}}(\hat{k}_t, \hat{z}_{n,t+1})$ depends on the sign of $\sigma - 1$.

[7] The deep parameters are described in Section 5. For this exercise, we temporarily set $\rho_{zn} = 0.95$ and $\nu_{zn} = 0.02$.

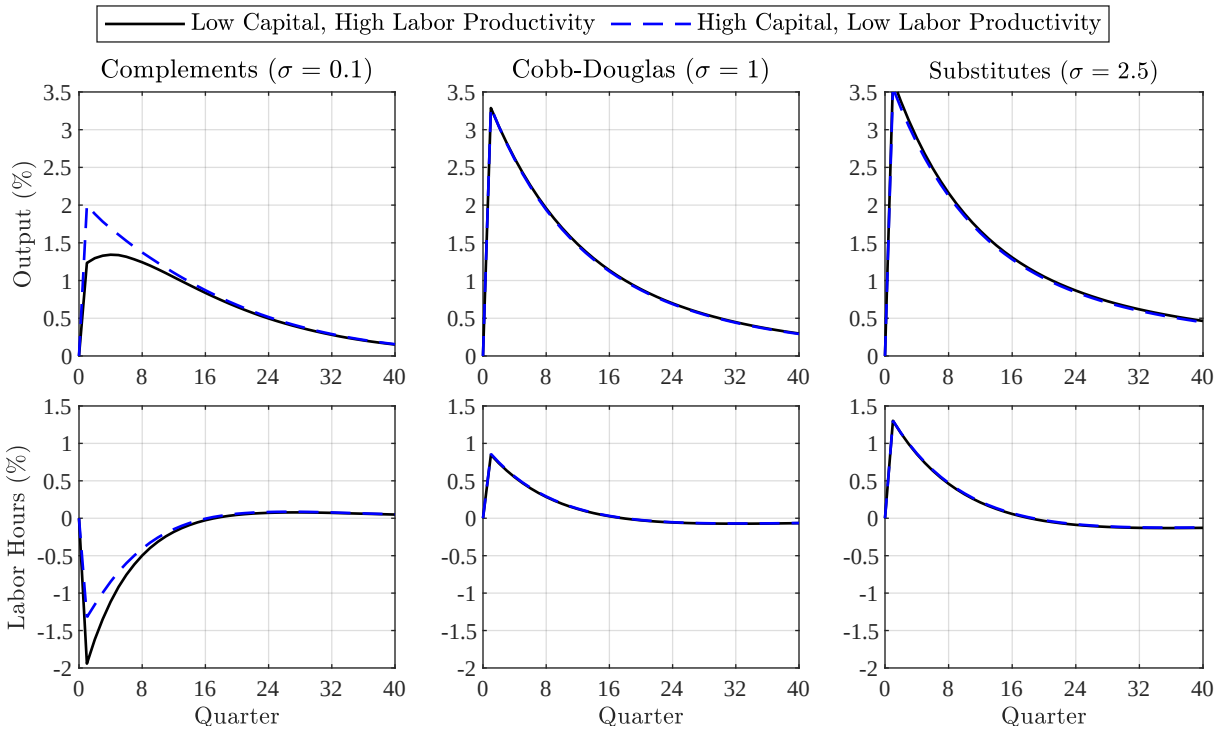

Figure 2: Generalized impulse responses to a 2 standard deviation shock to labor productivity in different states of the economy. To simplify the state, the model only includes labor productivity shocks and there are no real frictions. The capital state is $\pm 2\%$ of steady state. The labor productivity state is $\pm 4.5\%$ of steady state. These values correspond to the $(16, 84)$ percentiles of the ergodic distribution when $\sigma = 0.1$.

line reflects the response from a different initial state. The low (high) capital state is 2% below (above) steady state. The low (high) labor productivity state is 4.5% percent below (above) steady state. These values correspond to the $(16, 84)$ percentiles of the ergodic distribution when $\sigma = 0.1$.

There are several important takeaways from the impulse responses. First, a lower CES dampens the response of output ($\hat{y}$) to a labor productivity shock regardless of the initial state. The response of hours ($\hat{n}$) declines with a lower CES and is negative when $\sigma = 0.1$. This is driven by the complementarity in the production function. Since capital cannot immediately adjust in response to the shock, higher complementarity (a lower σ) reduces the increase in the marginal product of labor from a labor productivity shock and dampens the increase in labor demand. In the extreme case of perfect complementarity (Leontief production), a positive labor productivity shock causes a one-for-one decline in hours and no change in output, since the marginal product of labor is zero.[8]

Second, the impulse responses are state-dependent when there is a high degree of complementarity, as shown in the left panel where $\sigma = 0.1$. The responses of hours and output depend on the initial effective capital-to-labor ratio, $z_{k,t}k_{t-1}/(z_{n,t}n_t)$. When this ratio is low, as it is when k_{t-1} is

[8]Francis and Ramey (2005) show hours decline in response to a technology shock in a model with Leontief production. Cantore et al. (2014) find the sign of the response depends on the type of productivity shock and the CES value.

low and $z_{n,t}$ is high, hours decline more because complementarity reduces the marginal product of labor in these states. In turn, effective labor and hence output increases less. When k_{t-1} is high and $z_{n,t}$ is low, labor productivity shocks lead to a smaller decline in hours and larger increase in output.

Third, complementarity implies business cycle turning points driven by labor productivity shocks exhibit the most extreme levels of uncertainty. When capital is high and a large negative labor productivity shock occurs at the onset of a recession, the response to the other shocks is elevated, increasing uncertainty. When capital is low and a large positive productivity shock arrives at the start of an expansion, the response to the other shocks is weak, and future outcomes relatively certain. This generates a negative correlation between uncertainty and output growth in the model.

For conciseness, the responses to a capital productivity shock are shown in Appendix F. However, it is useful to note that the responses of hours and output strengthen with a lower CES because a capital productivity shock has the opposite effect on the marginal product of labor. The state-dependency also differs from a labor productivity shock. Initial periods with low k_{t-1} and $z_{k,t}$ have the largest responses of output, while periods with high k_{t-1} and $z_{k,t}$ have the smallest responses. This result is expected given the impact of the shock on the effective capital-to-labor ratio. Also, the highest levels of uncertainty occur at low levels of output, rather than business cycle turning points.

Under Cobb-Douglas production, there is no state dependence, confirming the intuition from the analytical results and Section 2. When $\sigma = 2.5$, the state dependence is in the opposite direction but negligible even though we consider a CES well above one. Overall, the state dependence is relatively weak even when σ is near zero and the economy is in the extreme states shown in Figure 2.

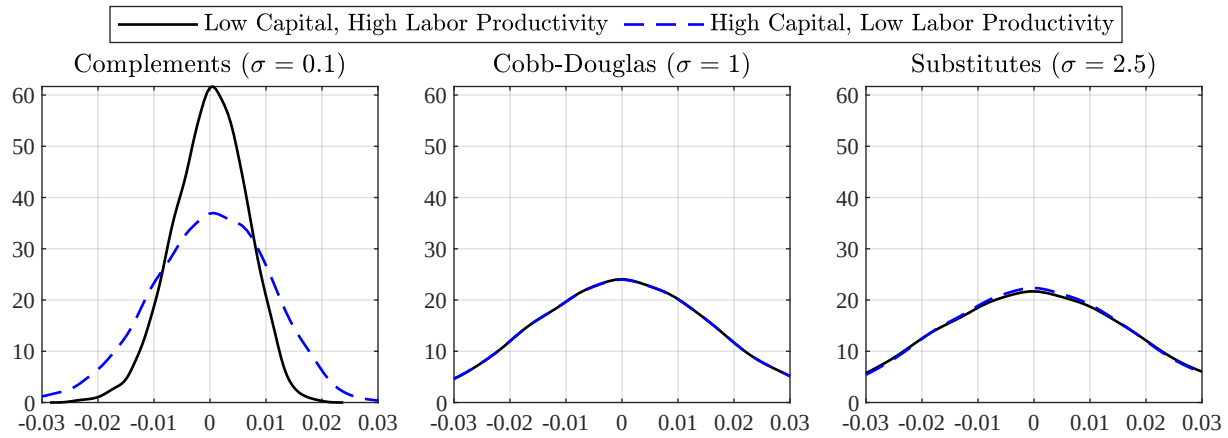

Figure 3: Conditional distributions for log output one period in the future in different states of the economy.

Conditional Output Distributions Another intuitive way to visualize how the degree of complementarity affects output uncertainty is to look at the conditional distributions for future log output. To generate these distributions, we first conduct $20{,}000$ 1-period simulations of the model at each state used to generate the impulse responses. We then use the simulated values from each

specification to construct a kernel density estimator of the distribution of next period's log output.

Figure 3 plots the conditional distribution for output in the same setting as the impulse responses. Consistent with earlier results, the distributions are invariant to the state of the economy when $\sigma \geq 1$. However, when $\sigma = 0.1$, the distributions differ depending upon the initial state. In situations where the initial capital-to-labor ratio is low, the distributions are tighter around the conditional mean of the distribution, whereas the distributions are wider in situations when the ratio is initially high. This reflects that in initial states with a depressed level of capital relative to labor productivity, changes in labor productivity next period will have a smaller impact on output. A narrower (wider) distribution generates a lower (higher) level of uncertainty. Therefore, the larger the differences in the conditional distributions, the greater the endogenous variation in uncertainty.

5 ESTIMATION RESULTS

This section presents our main results. We first outline our estimation strategy and then report the parameter estimates and targeted moments under several specifications of our quantitative model.

Data The nonlinear model is estimated with quarterly data from 1964-2019. Five parameters are set externally in line with our data sample and the literature. The discount factor, $\beta = 0.9959$, equals the inverse of the average real interest rate, which corresponds to the ratio of the average federal funds rate to the average GDP deflator inflation rate. The trend in the model, $\bar{g} = 1.0039$, matches the average growth rate of per capita output. The capital depreciation rate, $\delta = 0.0247$, equals the average rate on private fixed assets and consumer durable goods. The steady-state labor share of income $\bar{s}_n = 1 - \alpha$, since $n_0 = \bar{n} = 1/3$ and $y_0 = \bar{y} = 1$. We calibrate $\alpha = 0.3969$ using labor share data for the total economy from the Bureau of Labor Statistics. Finally, the Frisch elasticity of labor supply, $1/\eta = 0.5$, is set to the intensive margin estimate in Chetty et al. (2012).

The rest of the parameters are set to target twelve moments in the data: the standard deviations and first-order autocorrelations of output, consumption, and investment growth, and the standard deviations, first-order autocorrelations, and cyclicality of uncertainty and the labor share of income. The uncertainty data is based on the real uncertainty series from Ludvigson et al. (2020), shown in Figure 4. This series is a sub-index of the macro uncertainty series from Jurado et al. (2015) that accounts for 73 real activity variables (e.g., measures of output, income, housing, consumption, orders, and inventories). Most of their time series are transformed into growth rates and standard normalized. Simulations of a factor augmented vector autoregression are used to obtain estimates of uncertainty for each real variable and then averaged to generate the real uncertainty series. The benefit of this series is that it is based on the same definition of uncertainty as this paper, so it distinguishes between uncertainty and *ex-post* volatility. To make the units from our model comparable to the real uncertainty series, we define $SD(\mathcal{U}^{\hat{y}}) \equiv SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$, where $\hat{y}^g_t$ is

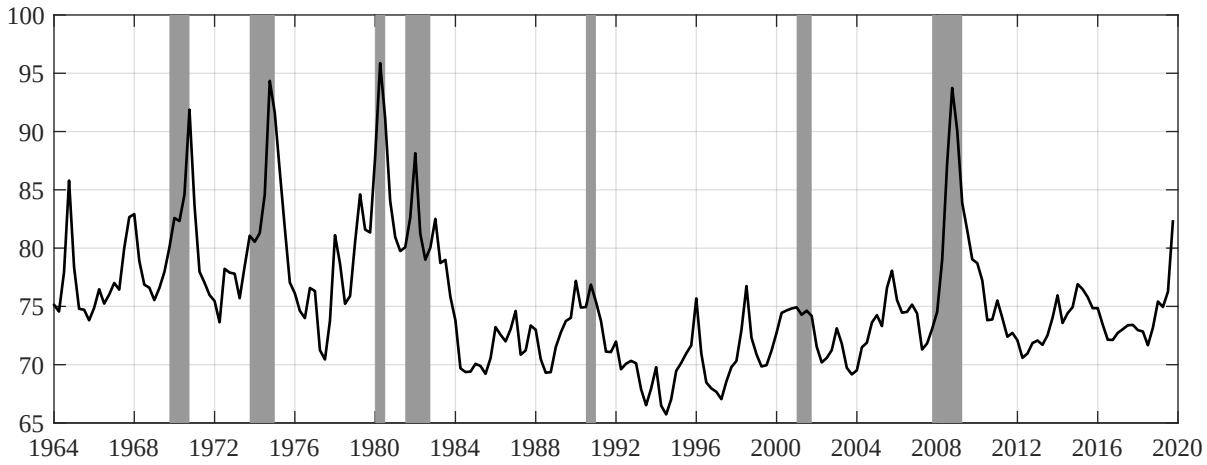

Figure 4: Real uncertainty series from Ludvigson et al. (2020). Shaded regions denote NBER recessions.

the quarter-over-quarter log-difference in output. Consumption includes expenditures on services and nondurables. Investment is composed of durable consumption and private fixed investment. Appendix D provides detailed descriptions of all our data sources and how they were transformed.

Methodology The empirical targets are stored in $\hat{\Psi}_T^D$ and estimated with a two-step Generalized Method of Moments (GMM) estimator, where $T = 224$ is the sample size. Conditional on the GMM estimates, the nonlinear model is estimated with Simulated Method of Moments (SMM). For parameterization θ and shocks $\mathcal{E}$, we solve the model and simulate it $R = 1{,}000$ times for T periods. The analogues of the targets are the median moments across the R simulations, $\bar{\Psi}_{R,T}^M(\theta, \mathcal{E})$.

The parameter estimates, $\hat{\theta}$, are obtained by minimizing the following loss function:

$$J(\theta, \mathcal{E}) = [\hat{\Psi}_T^D - \bar{\Psi}_{R,T}^M(\theta, \mathcal{E})]'[\hat{\Sigma}_T^D(1 + 1/R)]^{-1}[\hat{\Psi}_T^D - \bar{\Psi}_{R,T}^M(\theta, \mathcal{E})],$$

where $\hat{\Sigma}_T^D$ is the diagonal of the GMM estimate of the variance-covariance matrix. Monte Carlo methods are used to calculate the standard errors on the parameters.[9] We run our SMM algorithm $N_s = 100$ times, each time conditional on a particular sequence of shocks $\mathcal{E}^s$ but holding fixed the empirical targets, $\hat{\Psi}_T^D$, and weighting matrix, $\hat{\Sigma}_T^D$, used in the loss function. Given the set of parameter estimates $\{\hat{\theta}^s\}_{s=1}^{N_s}$, we report the mean, $\bar{\theta} = \sum_{s=1}^{N_s} \hat{\theta}^s / N_s$, and $(5, 95)$ percentiles.[10] While this method is numerically intensive, there are two major benefits to obtaining standard errors by Monte Carlo. First, it provides more reliable estimates of the standard errors than using the asymptotic variance of the estimator, which is the more common method used in the literature. Second, it is an effective way to determine whether the parameters are identified and check for multiple modes.

[9] Ruge-Murcia (2012) applies SMM to several nonlinear business cycle models and finds that asymptotic standard errors tend to overstate the variability of the estimates. This underscores the importance of using Monte Carlo methods.

[10] The practice of re-estimating with different sequences of shocks follows the recommendation of Fabio Canova (see http://apps.eui.eu/Personal/Canova/Teachingmaterial/Smm_eui2014.pdf, slide 16).

Parameter		No Exogenous Volatility Shocks			Baseline
		Labor Share Moments Untargeted	Uncertainty Moments Untargeted	All Moments Targeted	All Moments Targeted
Volatility Shock AC	ρ_ν	–	–	–	0.902 (0.891, 0.912)
Volatility Shock SD	ν_ν	–	–	–	0.028 (0.026, 0.030)
Labor Productivity AC	ρ_{zn}	0.954 (0.944, 0.964)	0.750 (0.654, 0.820)	0.953 (0.946, 0.959)	0.765 (0.701, 0.809)
Labor Productivity SD	ν_{zn}	0.025 (0.022, 0.028)	0.036 (0.030, 0.048)	0.027 (0.026, 0.028)	0.036 (0.032, 0.044)
Capital Productivity AC	ρ_{zk}	0.952 (0.913, 0.986)	0.550 (0.312, 0.765)	0.291 (0.251, 0.353)	0.388 (0.250, 0.643)
Capital Productivity SD	ν_{zk}	0.008 (0.007, 0.009)	0.009 (0.008, 0.010)	0.022 (0.020, 0.024)	0.009 (0.008, 0.009)
Investment Adjustment Cost	φ_x	2.36 (1.65, 3.25)	6.80 (5.45, 9.73)	5.10 (4.55, 5.73)	6.76 (5.79, 8.65)
External Habit Persistence	h	0.84 (0.78, 0.88)	0.94 (0.93, 0.96)	0.95 (0.95, 0.96)	0.95 (0.94, 0.96)
Elasticity of Substitution	σ	0.14 (0.11, 0.18)	0.49 (0.47, 0.51)	0.65 (0.63, 0.66)	0.49 (0.48, 0.50)
Loss Function	J	2.06 (1.70, 2.52)	6.40 (6.27, 6.48)	46.62 (45.52, 48.13)	14.73 (14.37, 15.05)

Table 1: Average and $(5, 95)$ percentiles of the parameter estimates and model fit.

Empirical Fit Table 1 shows the parameter estimates and Table 2 shows the model-implied moments under various specifications. We first report estimates from the model without exogenous volatility shocks to concentrate on the effects of complementarity. The first column ("Labor Share Moments Untargeted") shows the results when we exclude the standard deviation, autocorrelation, and cyclicality of the labor share from the set of targeted moments. In this case, the estimated model is able to endogenously generate all of the empirical volatility of uncertainty (5.51 in the model vs. 5.39 in the data), while almost perfectly matching all of the other empirical targets ($J = 2.06$). These results confirm the theory that complementarity can be a significant source of time-varying endogenous uncertainty. However, achieving these results requires a very low CES, which causes the model to overstate the volatility of the labor share (9.76 in the model vs. 1.55 in the data).

The second column ("Uncertainty Moments Untargeted") shows the results when we exclude the standard deviation, autocorrelation, and cyclicality of uncertainty from the set of targeted moments. These results illustrate that the labor share moments are crucial to identify the CES. The estimated CES increases to 0.49 to match the volatility of the labor share in the data. While the CES estimate is still well below the unitary elasticity of a Cobb-Douglas production function, the model is able to generate only about 20% of the empirical volatility of uncertainty (1.03 in the model vs. 5.39 in the data). The fit of the model ($J = 6.40$) is also slightly worse than the first specification.

In the third column, we target all twelve moments in our model without exogenous volatility shocks. The CES estimate increases to 0.65 and the model is able to match all of the targeted labor

	Data		No Exogenous Volatility Shocks			Baseline
Moment	Mean	SE	Labor Share Moments Untargeted	Uncertainty Moments Untargeted	All Moments Targeted	All Moments Targeted
$SD(\mathcal{U}^{\hat{y}})$	5.39	0.56	5.51	1.03^u	2.54	5.07
$AC(\mathcal{U}^{\hat{y}})$	0.87	0.04	0.88	0.51^u	0.86	0.86
$Corr(\hat{y}^g, \mathcal{U}^{\hat{y}})$	-0.43	0.09	-0.43	-0.79^u	-0.37	-0.19
$SD(\tilde{s}_n)$	1.55	0.12	9.76^u	1.56	1.58	1.57
$AC(\tilde{s}_n)$	0.84	0.04	0.93^u	0.84	0.82	0.84
$Corr(\hat{y}^g, \tilde{s}_n)$	-0.25	0.07	-0.28^u	-0.25	-0.19	-0.27
$SD(\hat{y}^g)$	0.80	0.07	0.85	0.92	0.96	0.91
$SD(\hat{c}^g)$	0.51	0.05	0.51	0.50	0.38	0.46
$SD(\hat{x}^g)$	2.03	0.20	1.84	1.73	2.14	1.78
$AC(\hat{y}^g)$	0.29	0.09	0.26	0.36	0.44	0.37
$AC(\hat{c}^g)$	0.39	0.07	0.38	0.39	0.42	0.39
$AC(\hat{x}^g)$	0.42	0.09	0.38	0.36	0.47	0.37

Table 2: Data and model-implied moments. The first section shows the uncertainty moments, the second section shows the labor share moments, and the last section shows the real activity moments. A superscript u denotes an untargeted moment. A tilde denotes a detrended variable. The trend in the data is based on a Hamilton (2018) filter with an 8-quarter window. The model-implied trend is equal to the simulated mean.

share moments. In spite of the higher CES estimate, the model is able to endogenously generate twice as much volatility in uncertainty as the second column (2.54, or about 40% of the data), but that improvement comes at the expense of overstating several real activity moments, including the volatility and autocorrelation of output growth. The fit of the model is much worse than other the specifications ($J = 46.62$), though that is mostly because it understates the volatility of uncertainty.

The last column shows the estimates from our baseline model with exogenous volatility shocks. The parameter estimates are similar to those in the second column. The model provides a good fit of most of the twelve targeted moments ($J = 14.73$). In particular, exogenous volatility shocks allow the model to simultaneously match the volatility of uncertainty and labor market dynamics. However, the cyclicality of uncertainty now falls below its empirical counterpart (-0.19 in the model vs. -0.43 in the data), which suggests there is a role for other sources of endogenous uncertainty.

To summarize, complementarity by itself can generate empirically consistent endogenous variation in uncertainty, but it cannot do so while simultaneously matching labor share dynamics. Moreover, if the model includes exogenous volatility shocks to allow for a horse-race between the endogenous and exogenous sources of time-varying uncertainty, the estimated parameters are similar to estimates we obtain when the uncertainty moments are excluded from the set of empirical targets. This suggests endogenous uncertainty is playing a relatively minor role in matching the data.

Uncertainty Decomposition Exogenous volatility shocks introduce an exogenous source of time-varying uncertainty, while complementarity endogenously generates fluctuations in uncertainty.

To determine how much of the variance in aggregate uncertainty is driven by complementarity, we compute a generalized forecast error variance decomposition (GFEVD). The method applied to linear models is easily generalized to a nonlinear model by replacing a linear impulse response function (IRF) with a generalized impulse response function (GIRF), which accounts for the initial state and shock size. Recall from Section 4.2 that the GIRF of variable x_{t+h} over horizon h is given by

$$\mathcal{G}_{i,t}(x_{t+h}|\varepsilon_{i,t+1} = \xi, \mathbf{z}_t) = E_t[x_{t+h}|\varepsilon_{i,t+1} = \xi, \mathbf{z}_t] - E_t[x_{t+h}|\mathbf{z}_t],$$

where ξ is the size of shock $i \in \{zn, zk, \nu\}$ and $\mathbf{z}_t$ is a vector of initial states. Following Lanne and Nyberg (2016), the GFEVD of variable x_{t+h} into component i over time horizon h is then given by

$$\lambda_{i,t}(x_{t+h}|\mathbf{z}_t) = \int_{-\infty}^{\infty} \frac{\sum_{j=1}^{h} \left(\mathcal{G}(x_{t+j}|\varepsilon_{i,t+1} = \xi, \mathbf{z}_t)\right)^2}{\sum_{i=1}^{m}\sum_{j=1}^{h} \left(\mathcal{G}(x_{t+j}|\varepsilon_{i,t+1} = \xi, \mathbf{z}_t)\right)^2} f(\xi)d\xi,$$

where $f(\cdot)$ is the probability density function of ε_i. This method requires us to integrate across the shock since the response of an endogenous variable is not necessarily a linear function of the shock size in a nonlinear model. We use Gauss Hermite quadrature with 10 points to discretize the shock.

	$h = 1$	$h = 8$	$h = 16$	Long-run
Exogenous Uncertainy	84.2	93.5	94.1	94.2
Endogenous Uncertainy	15.8	6.5	5.9	5.8

Table 3: Forecast error variance decomposition of aggregate uncertainty ($\mathcal{U}^{\hat{y}}$) over different horizons (h).

Table 3 shows the decomposition of aggregate uncertainty, $SD(\mathcal{U}^{\hat{y}})$, into its exogenous, λ_ν, and endogenous, $\lambda_{zn} + \lambda_{zk}$, sources over various horizons (h). When $h = 1$, only 15.8% of the variation in aggregate uncertainty is endogenous. That percentage steadily declines to 5.8% in the long-run ($h \geq 20$), because the exogenous volatility shocks are more persistent ($\rho_\nu = 0.9$) than capital and labor productivity ($\rho_{zk} = 0.39$; $\rho_{zn} = 0.77$) and complementarity is not strong enough to significantly affect the internal propagation in the model under our baseline parameter estimates.

Impulse Responses A useful way to see the role of complementarity is with the GIRFs that underlie the variance decomposition in Table 3. Figure 5 plots the responses of output, aggregate uncertainty, the effective capital-to-labor ratio, and the labor share in our baseline model to 1 standard deviation positive shocks to capital and labor productivity. We show the results using three different values of the CES: (1) Our estimate when we exclude exogenous volatility shocks and the labor share targets ($\sigma = 0.14$), (2) our baseline estimate ($\sigma = 0.49$), and (3) the value under Cobb-Douglas production ($\sigma = 1$). All of the other parameters are set to the baseline estimates in Table 1.

For all three CES values, an increase in labor productivity leads to higher output by raising the effective labor supply ($z_{n,t}n_t$). With labor relatively more abundant and the capital stock fixed on

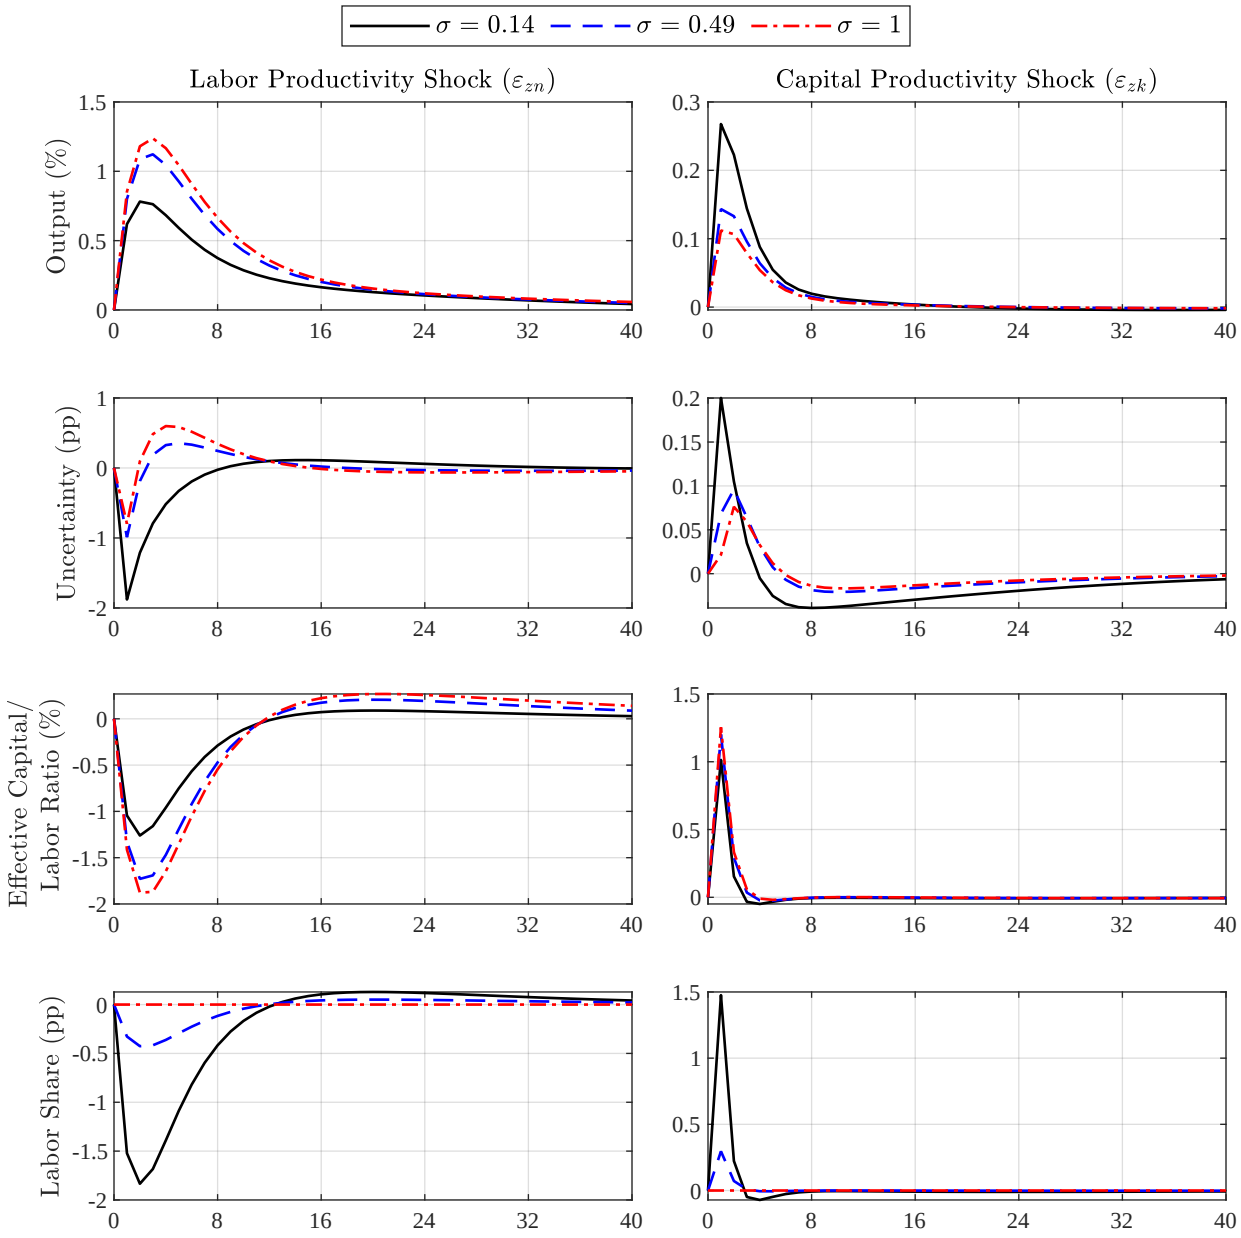

Figure 5: Generalized impulse responses to a 1 standard deviation positive shock. Output is shown in percent deviations from steady state. Uncertainty is normalized by $SD(\hat{y}^g)$ and shown as a percentage point (pp) difference from steady state. Except for the CES, all of the parameters are set to their baseline estimates.

impact, the effective capital-to-labor ratio ($z_{k,t}k_{t-1}/z_{n,t}n_t$) falls. As demonstrated in Figure 2 and Figure 3, starting from a lower effective capital-to-labor ratio makes output less sensitive to future labor productivity shocks (lower uncertainty) but more sensitive to capital productivity shocks (higher uncertainty). Under the baseline parameter estimates, the volatility of labor productivity is much higher than the volatility of capital productivity. Therefore, uncertainty falls in response to a labor productivity shock. The intuition works in reverse for the capital productivity shock. It

raises the effective capital-to-labor ratio, making output more sensitive to future labor productivity shocks and less sensitive to future capital productivity shocks and, on net, increasing uncertainty.[11]

The magnitudes of the impulse responses are affected by the degree of complementarity. With a lower CES, higher labor productivity leads to a smaller increase in the marginal product of labor, which translates into a smaller increase in output. A lower CES also leads to a smaller decline in the effective capital-to-labor ratio. As inputs become more complementary, their relative values become less volatile, which by itself would dampen the uncertainty response. However, the production function also becomes more concave with a lower CES. Under our baseline calibration, the latter effect dominates, so the magnitude of the uncertainty response strengthens as the CES declines. Finally, a lower CES significantly increases the response of the labor share to either shock.

In spite of the nonlinearities from complementarity, the output and uncertainty responses under our baseline CES estimate are similar to the responses under Cobb-Douglas production. Larger differences only occur when the CES is near zero and the model significantly overstates the empirical volatility of the labor share. This suggests a Cobb-Douglas production function provides a reasonable approximation of business cycle dynamics, even when accounting for higher-order moments.

6 CONCLUSION

Macroeconomic uncertainty regularly fluctuates in the data. The business cycle literature typically accounts for these fluctuations with exogenous volatility shocks, which introduce an exogenous source of time-varying uncertainty. This paper focuses on endogenous movements in uncertainty that arise due to the state-dependent effects of first-moment shocks. Theory suggests that complementarity between capital and labor inputs in production is a potentially important source of time-varying endogenous uncertainty because output becomes a concave function of productivity. We examine the quantitative significance of this channel by estimating several variants of a nonlinear real business cycle model with a CES production function and exogenous volatility shocks.

We find an estimated model without exogenous volatility shocks can endogenously generate all of the empirical variation in uncertainty, but only at the expense of significantly overstating the volatility of the labor share. When we match labor share dynamics, our estimate of the CES increases, decreasing the model's ability to endogenously generate time-varying uncertainty. In our baseline model that matches labor share and uncertainty dynamics, at most 16% of the variation in aggregate uncertainty is endogenous. Furthermore, the responses of output and uncertainty to a productivity shock are similar to the responses under Cobb-Douglas production. Despite the theoretical appeal of complementarity, these results indicate that its effects are quantitatively small.

There are three main avenues to build on our results. First, it would be interesting to add a competing source of variation in the labor share (e.g., price markup). Currently, variation in the labor

[11] Appendix F compares the responses to a capital and labor productivity shock under different shock sizes and signs.

share and endogenous uncertainty are both driven by the CES. An additional source of variation in the labor share would decouple these movements. Second, one could examine the quantitative significance of other sources of endogenous uncertainty. Business cycle models that endogenously generate a large share of the empirical fluctuations in aggregate uncertainty would likely help the model match other higher-order moments due to the state-dependent effects of first-moment shocks. Three, one could develop richer models that are able to simultaneously explain the empirical volatility of macroeconomic and financial uncertainty. Any of these avenues would advance not only the uncertainty literature, but also the quantitative performance of business cycle models.

REFERENCES

ARELLANO, C., Y. BAI, AND P. J. KEHOE (2019): "Financial Frictions and Fluctuations in Volatility," *Journal of Political Economy*, 127, 2049–2103, https://doi.org/10.1086/701792.

BENAROYA, H., S. M. HAN, AND M. NAGURKA (2005): *Probability models in engineering and science*, CRC press.

BERNSTEIN, J., A. W. RICHTER, AND N. A. THROCKMORTON (2020): "Cyclical Net Entry and Exit," Federal Reserve Bank of Dallas Working Paper 2018, https://doi.org/10.24149/wp2018r1.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685, https://doi.org/10.3982/ECTA6248.

BORN, B. AND J. PFEIFER (2014): "Policy Risk and the Business Cycle," *Journal of Monetary Economics*, 68, 68–85, https://doi.org/10.1016/j.jmoneco.2014.07.012.

BRUNNERMEIER, M. K. AND Y. SANNIKOV (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104, 379–421, https://doi.org/10.1257/aer.104.2.379.

CANTORE, C., M. LEON-LEDESMA, P. MCADAM, AND A. WILLMAN (2014): "Shocking Stuff: Technology, Hours, And Factor Substitution," *Journal of the European Economic Association*, 12, 108–128, https://doi.org/10.1111/jeea.12038.

CANTORE, C., P. LEVINE, J. PEARLMAN, AND B. YANG (2015): "CES technology and business cycle fluctuations," *Journal of Economic Dynamics and Control*, 61, 133–151, 10.1016/j.jedc.2015.09.00.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2012): "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," in *NBER Macroeconomics Annual 2012, Volume 27*, ed. by D. Acemoglu, J. Parker, and M. Woodford, MIT Press, Cambridge, 1–56.

COLEMAN, II, W. J. (1991): "Equilibrium in a Production Economy with an Income Tax," *Econometrica*, 59, 1091–1104, https://doi.org/10.2307/2938175.

EDEN, M. AND P. GAGGL (2018): "On the Welfare Implications of Automation," *Review of Economic Dynamics*, 29, 15–43, 10.1016/j.red.2017.12.003.

FAJGELBAUM, P., M. TASCHEREAU-DUMOUCHEL, AND E. SCHAAL (2017): "Uncertainty Traps," *The Quarterly Journal of Economics*, 132, 1641–1692, https://doi.org/10.1093/qje/qjx02.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, K. KUESTER, AND J. F. RUBIO-RAMÍREZ (2015): "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, 105, 3352–84, https://doi.org/10.1257/aer.20121236.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–61, https://doi.org/10.1257/aer.101.6.2530.

FRANCIS, N. AND V. A. RAMEY (2005): "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 52, 1379–1399, https://doi.org/10.1016/j.jmoneco.2004.08.009.

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

HAMILTON, J. D. (2018): "Why You Should Never Use the Hodrick-Prescott Filter," *Review of Economics and Statistics*, 100, 831–843, https://doi.org/10.1162/rest_a_00706.

HE, Z. AND A. KRISHNAMURTHY (2019): "A Macroeconomic Framework for Quantifying Systemic Risk," *American Economic Journal: Macroeconomics*, 11, 1–37, 10.1257/mac.20180011.

HUBMER, J. (2020): "The Race Between Preferences and Technology," Working Paper, University of Pennslyvania.

ILUT, C., M. KEHRIG, AND M. SCHNEIDER (2018): "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News," *Journal of Political Economy*, 126, 2011–2071, https://doi.org/10.1086/699189.

ILUT, C. AND H. SAIJO (2020): "Learning, confidence, and business cycles," *Journal of Monetary Economics*, forthcoming, https://doi.org/10.1016/j.jmoneco.2020.01.010.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216, https://www.doi.org/10.1257/aer.20131193.

KARABARBOUNIS, L. AND B. NEIMAN (2014): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129, 61–103, https://doi.org/10.1093/qje/qjt032.

KLUMP, R., P. MCADAM, AND A. WILLMAN (2012): "The normalized CES production function: theory and empirics," *Journal of Economic Surveys*, 26, 769–799, https://doi.org/10.1111/j.1467-6419.2012.00730.x.

KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): "Impulse Response Analysis in Nonlinear Multivariate Models," *Journal of Econometrics*, 74, 119–147, https://doi.org/10.1016/0304-4076(95)01753-4.

KOPECKY, K. AND R. SUEN (2010): "Finite State Markov-chain Approximations to Highly Persistent Processes," *Review of Economic Dynamics*, 13, 701–714, https://doi.org/10.1016/j.red.2010.02.002.

LANNE, M. AND H. NYBERG (2016): "Generalized Forecast Error Variance Decomposition for Linear and Nonlinear Multivariate Models," *Oxford Bulletin of Economics and Statistics*, 78, 595–603, https://doi.org/10.1111/obes.12125.

LEDUC, S. AND Z. LIU (2016): "Uncertainty Shocks are Aggregate Demand Shocks," *Journal of Monetary Economics*, 82, 20–35, https://doi.org/10.1016/j.jmoneco.2016.07.002.

LEON-LEDESMA, M. A., P. MCADAM, AND A. WILLMAN (2010): "Identifying the Elasticity of Substitution with Biased Technical Change," *American Economic Review*, 100, 1330–1357, https://doi.org/10.1257/aer.100.4.1330.

LEON-LEDESMA, M. A. AND M. SATCHI (2019): "Appropriate Technology and Balanced Growth," *Review of Economic Studies*, 86, 807–835, https://doi.org/10.1093/restud/rdy002.

LUDVIGSON, S. C., S. MA, AND S. NG (2020): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Journal: Marcoeconomics*, forthcoming.

MENDOZA, E. G. (2010): "Sudden Stops, Financial Crises, and Leverage," *American Economic Review*, 100, 1941–1966, https://doi.org/10.1257/aer.100.5.1941.

MUMTAZ, H. AND F. ZANETTI (2013): "The Impact of the Volatility of Monetary Policy Shocks," *Journal of Money, Credit and Banking*, 45, 535–558, https://doi.org/10.1111/jmcb.12015.

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

NEWEY, W. K. AND K. D. WEST (1987): "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708, https://www.doi.org/10.2307/1913610.

OBERFIELD, E. AND D. RAVAL (2021): "Micro Data and Macro Technology," *Econometrica*, 89, 703–732, https://doi.org/10.3982/ECTA12807.

PLANTE, M., A. W. RICHTER, AND N. A. THROCKMORTON (2018): "The Zero Lower Bound and Endogenous Uncertainty," *Economic Journal*, 128, 1730–1757, https://doi.org/10.1111/ecoj.12445.

RICHTER, A. W., N. A. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476, https://doi.org/10.1007/s10614-013-9399-2.

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.

RUGE-MURCIA, F. (2012): "Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles," *Journal of Economic Dynamics and Control*, 36, 914–938, http:https://doi.org/10.1016/j.jedc.2012.01.00.

SAIJO, H. (2017): "The Uncertainty Multiplier and Business Cycles," *Journal of Economic Dynamics and Control*, 78, 1–25, https://doi.org/10.1016/j.jedc.2017.02.008.

SIMS, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20, https://doi.org/10.1023/A:1020517101123.

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

——— (2019): "Endogenous second moments: A unified approach to fluctuations in risk, dispersion, and uncertainty," *Journal of Economic Theory*, 183, 625–660, https://doi.org/10.1016/j.jet.2019.07.007.

VAN NIEUWERBURGH, S. AND L. VELDKAMP (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53, 753–772, https://doi.org/10.1016/j.jmoneco.2005.02.003.

A UNCERTAINTY APPROXIMATIONS

Suppose $y = g(x)$ is a function of a random variable x with mean $\bar{x}$, unconditional variance $\nu^2_{x,u}$, and conditional (one-step ahead) variance $\nu^2_{x,c}$. A second-order Taylor expansion around $\bar{x}$ implies

$$y \approx g(\bar{x}) + g'(\bar{x})(x - \bar{x}) + g''(\bar{x})(x - \bar{x})^2/2,$$

so the unconditional expectation is given by

$$E[y] \approx g(\bar{x}) + g''(\bar{x})\nu^2_{x,u}/2.$$

By definition $\text{Var}[y] = E[(y - \bar{y})^2] = E[y^2] - (E[y])^2$. Therefore, using the same approximation

$$\text{Var}[y] \approx [g(\bar{x})]^2 + ([g'(\bar{x})]^2 + g(\bar{x})g''(\bar{x}))\nu^2_{x,u} - (g(\bar{x}) + g''(\bar{x})\nu^2_{x,u}/2)^2.$$

Expanding the approximation and dropping the term involving $\nu_{x,u}^4$ implies $\text{Var}[y] \approx [g'(\bar{x})]^2 \nu_{x,u}^2$.

An approximation of the time-t conditional variance of y_{t+1}, $\text{Var}_t[y_{t+1}]$, follows the same derivation, except it is approximated around the conditional mean of x_{t+1}, $E_t[x_{t+1}]$. Therefore, $\text{Var}_t[y_{t+1}] \approx [g'(E_t[x_{t+1}])]^2 \nu_{x,c}^2$. Now suppose x follows an $AR(1)$ process with standard-normal shock ε_x where the volatility of the shock, ν_x, is an independent stochastic process. The conditional variance of x is now stochastic and given by $E_t[(x_{t+1} - E_t[x_{t+1}])^2] = E_t[(\nu_{x,t+1}\epsilon_{x,t+1})^2] = E_t[\nu_{x,t+1}^2]$, because ε_x and ν_x are independent. Therefore, $\text{Var}_t[y_{t+1}] \approx [g'(E_t[x_{t+1}])]^2 E_t[\nu_{x,t+1}^2]$.

B Detrended Equilibrium System

Our detrended baseline model includes the exogenous processes in (7)-(8) and

$$
\begin{aligned}
r_t^k &= \alpha\big(y_0 z_{k,t}/k_0\big)^{\frac{\sigma-1}{\sigma}}\big(\bar{g}\tilde{y}_t/\tilde{k}_{t-1}\big)^{\frac{1}{\sigma}},\\
\tilde{y}_t &= y_0\left[\alpha\left(z_{k,t}\tilde{k}_{t-1}/(\bar{g}k_0)\right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)\Big(z_{n,t}n_t/n_0\Big)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},\\
\tilde{w}_t &= (1-\alpha)\left(y_0 z_{n,t}/n_0\right)^{\frac{\sigma-1}{\sigma}}\left(\tilde{y}_t/n_t\right)^{\frac{1}{\sigma}},\\
\tilde{w}_t &= \chi n_t^{\eta}\tilde{\lambda}_t,\\
\tilde{\lambda}_t &= \tilde{c}_t - (h/\bar{g})\tilde{c}_{t-1},\\
\tilde{c}_t + \tilde{x}_t &= \tilde{y}_t,\\
\tilde{k}_t &= (1-\delta)\tilde{k}_{t-1}/\bar{g} + \tilde{x}_t(1-\varphi^x(x_t^{gap}-1)^2/2),\\
1 &= q_t[1-\varphi^x(x_t^{gap}-1)(3x_t^g-1)/2] + \beta\varphi^x E_t[(\tilde{\lambda}_t/\tilde{\lambda}_{t+1})q_{t+1}(x_{t+1}^{gap})^2(x_{t+1}^{gap}-1)],\\
q_t &= (\beta/\bar{g})E_t[(\tilde{\lambda}_t/\tilde{\lambda}_{t+1})(r_{t+1}^k + (1-\delta)q_{t+1})],\\
x_t^{gap} &= \tilde{x}_t/\tilde{x}_{t-1}.
\end{aligned}
$$

C Nonlinear Solution Method

We begin by compactly writing the detrended nonlinear equilibrium system as

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

where f is a vector-valued function, $\mathbf{s}_t$ are the variables, $\varepsilon_t = [\varepsilon_{zn,t}, \varepsilon_{zk,t}, \nu_t]$ are the shocks, $\mathbf{z}_t \equiv [\tilde{k}_{t-1}, \tilde{c}_{t-1}, \tilde{x}_{t-1}, z_{n,t}, z_{k,t}, \nu_t]$ are the initial state variables, and ϑ are the parameters.

There are many ways to discretize the exogenous volatility shock, ν. We use the Markov chain in Rouwenhorst (1995), which Kopecky and Suen (2010) show outperforms other methods for approximating autoregressive processes. The bounds on $\tilde{k}_{t-1}$, $\tilde{c}_{t-1}$, $z_{n,t}$, and $z_{k,t}$, are set to $\pm 10\%$ of their deterministic steady state values, while $\tilde{x}_{t-1}$ is set to $\pm 25\%$ of its deterministic steady state. These values were chosen so the grids contain at least 99% of the simulated values for each state.

We discretize the endogenous state variables ($\tilde{c}_{t-1}$, $\tilde{k}_{t-1}$, and $\tilde{x}_{t-1}$) into 9 evenly-spaced points and the exogenous variables into 7 evenly-spaced points. There are $D = 250{,}047$ nodes in the state space, and the realization of $\mathbf{z}_t$ on node d is denoted $\mathbf{z}_t(d)$. The Rouwenhorst method provides integration nodes $\nu_{t+1}(m)$ that are the same as the state variable. However, the processes for z_n and z_k do not have a standard autoregressive form because of the exogenous volatility shocks. Thus, the first moment shocks, $[\varepsilon_{zn,t}, \varepsilon_{zk,t}]$, are discretized separately from the volatility process. The policy functions are interpolated at realizations of $z_{n,t+1}(m)$ and $z_{k,t+1}(m)$ that can occur between nodes in the state space. We use the same number of interpolation nodes as the state variables, $(7, 7, 7)$, so $M = 343$. The Rouwenhorst method provides weights, $\phi(m)$, for $m \in \{1, \ldots, M\}$.

The vector of policy functions is denoted $\mathbf{pf}_t \equiv [n_t(\mathbf{z}_t), q_t(\mathbf{z}_t)]$ and the realization on node d is denoted $\mathbf{pf}_t(d)$. Our choice of policy functions, while not unique, simplifies solving for the variables in the nonlinear system of equations given $\mathbf{z}_t$. The following steps outline our algorithm:

1. Use Sims's (2002) `gensys` algorithm to solve the log-linear model. Then map the solution for the policy functions to the discretized state space. This provides an initial conjecture.

2. On iteration $j \in \{1, 2, \ldots\}$ and each node $d \in \{1, \ldots, D\}$, use Chris Sims's `csolve` to find $\mathbf{pf}_t(d)$ to satisfy $E[f(\cdot)|\mathbf{z}_t(d), \vartheta] \approx 0$. Guess $\mathbf{pf}_t(d) = \mathbf{pf}_{j-1}(d)$. Then apply the following:
   (a) Solve for all variables dated at time t, given $\mathbf{pf}_t(d)$ and $\mathbf{z}_t(d)$.
   (b) Linearly interpolate the policy functions, $\mathbf{pf}_{j-1}$, at the updated state variables, $\mathbf{z}_{t+1}(m)$, to obtain $\mathbf{pf}_{t+1}(m)$ on every integration node, $m \in \{1, \ldots, M\}$.
   (c) Given $\{\mathbf{pf}_{t+1}(m)\}_{m=1}^{M}$, solve for the other elements of $\mathbf{s}_{t+1}(m)$ and compute

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

   When the nonlinear solver converges, set $\mathbf{pf}_j(d) = \mathbf{pf}_t(d)$.

3. Repeat Step 2 until $\text{maxdist}_j < 10^{-6}$, where $\text{maxdist}_j \equiv \max\{|\mathbf{pf}_j - \mathbf{pf}_{j-1}|\}$. When that criterion is satisfied, the algorithm has converged to an approximate nonlinear solution.

D DATA SOURCES

We use the following quarterly time-series from 1964-2019 provided by Haver Analytics:

1. **Civilian Noninstitutional Population: 16 Years and Over**,
   Not Seasonally Adjusted, Thousands (LN16N@USECON)

2. **Gross Domestic Product: Implicit Price Deflator**,
   Seasonally Adjusted, 2012=100 (DGDP@USNA)

3. **Personal Consumption Expenditures: Nondurable Goods**, Seasonally Adjusted, Billions of Dollars (CN@USECON)
4. **Personal Consumption Expenditures: Services**, Seasonally Adjusted, Billions of Dollars (CS@USECON)
5. **Personal Consumption Expenditures: Durable Goods**, Seasonally Adjusted, Billions of Dollars (CD@USECON)
6. **Private Fixed Investment**, Seasonally Adjusted, Billions of Dollars (F@USECON)
7. **Gross Domestic Product**, Seasonally Adjusted, Billions of Dollars, (GDP@USECON)
8. **Labor Share**, Total Economy, All Employed Persons (LXEBL@USECON)
9. **Net Stock: Private Fixed Assets**, Billions of Dollars (EPT@CAPSTOCK)
10. **Net Stock: Consumer Durable Goods**, Billions of Dollars (EDT@CAPSTOCK)
11. **Depreciation: Private Fixed Assets**, Billions of Dollars (KPT@CAPSTOCK)
12. **Depreciation: Consumer Durable Goods**, Billions of Dollars (KDT@CAPSTOCK)
13. **Effective Federal Funds Rate**, Percent per Annum (FFED@USECON)

We also use the Real Uncertainty Index developed in Ludvigson et al. (2020), which is regularly updated on Sydney Ludvigson's personal website. We use a quarterly average of monthly values based on a 1-quarter forecast horizon ($h = 3$). The data was retrieved on September 7, 2020.

We applied the following transformations to the above data sources:

1. **Per Capita Real Output Growth**:

$$\Delta\hat{y}_t = 100\left(\log\left(\frac{GDP_t}{DGDP_t+LN16N_t}\right) - \log\left(\frac{GDP_{t-1}}{DGDP_{t-1}+LN16N_{t-1}}\right)\right).$$

2. **Per Capita Real Consumption Growth**:

$$\Delta\hat{c}_t = 100\left(\log\left(\frac{CN_t+CS_t}{DGDP_t+LN16N_t}\right) - \log\left(\frac{CN_{t-1}+CS_{t-1}}{DGDP_{t-1}+LN16N_{t-1}}\right)\right).$$

3. **Per Capita Real Investment Growth**:

$$\Delta\hat{x}_t = 100\left(\log\left(\frac{F_t+CD_t}{DGDP_t+LN16N_t}\right) - \log\left(\frac{F_{t-1}+CD_{t-1}}{DGDP_{t-1}+LN16N_{t-1}}\right)\right).$$

4. **Subjective Discount Factor**:

$$\beta = \tfrac{1}{T}\textstyle\sum_{t=1}^{T}(DGDP_t/DGDP_{t-1})/(1+FFED/100)^{1/4}.$$

5. **Capital Depreciation Rate**:

$$\delta = (1 + \tfrac{1}{T}\textstyle\sum_{t=1}^{T}(KPT_t + KDT_t)/(EPT_{t-1} + EDT_{t-1}))^{1/12} - 1.$$

E ESTIMATION METHOD

The estimation procedure has two stages. The first stage estimates moments in the data using a 2-step Generalized Method of Moments (GMM) estimator with a Newey and West (1987) weighting matrix with 5 lags. The second stage is a Simulated Method of Moments (SMM) procedure that searches for a parameter vector that minimizes the distance between the GMM estimates in the data and short-sample predictions of the model, weighted by the diagonal of the GMM estimate of the variance-covariance matrix. The second stage is repeated for many different draws of shocks to obtain standard errors for the parameter estimates. The following steps outline the algorithm:

1. Use GMM to estimate the moments, $\hat{\Psi}_T^D$, and the diagonal of the covariance matrix, $\hat{\Sigma}_T^D$.

2. Use SMM to estimate the detrended linear model. Given a random seed, h, draw a $B+T$ period sequence for each shock in the model, where B is a $1{,}000$ period burn-in and T is the sample size of the quarterly time series. Denote the shock matrix by $\mathcal{E}^s = [\varepsilon_{zk}^s, \varepsilon_{zn}^s, \varepsilon_{\nu}^s]_{t=1}^{B+T})$.

   For shock sequence $s \in \{1, \ldots, N_s\}$, run the following steps:

   (a) Specify a guess, $\hat{\theta}_0$, for the N_p estimated parameters and the covariance matrix, $\Sigma_P^{s,0}$. For all $i \in \{1, \ldots, N_m\}$, apply the following steps:

      i. Draw $\hat{\theta}_i$ from a multivariate normal distribution centered at some mean parameter vector, $\bar{\theta}$, with a diagonal covariance matrix, Σ_0.

      ii. Solve the linear model with Sims's (2002) gensys algorithm given $\hat{\theta}_i$. Repeat the previous step if the solution does not exist or is not unique.

      iii. Given $\mathcal{E}^s(r)$, simulate the quarterly model R times for $B+T$ periods. We draw initial states from the ergodic distribution by burning off the first B periods. For each repetition r, calculate the moments based on T quarters, $\Psi_T^M(\hat{\theta}_i, \mathcal{E}^s(r))$.

      iv. Calculate the median moments across the R simulations,

$$\bar{\Psi}_{R,T}^M(\hat{\theta}_i, \mathcal{E}^s) = \text{median}\{\Psi_T^M(\hat{\theta}_i, \mathcal{E}^s(r))\}_{r=1}^R,$$

      and evaluate the loss function:,

$$J_i^s = [\hat{\Psi}_T^D - \bar{\Psi}_{R,T}^M(\hat{\theta}_i, \mathcal{E}^s)]'[\hat{\Sigma}_T^D(1 + 1/R)]^{-1}[\hat{\Psi}_T^D - \bar{\Psi}_{R,T}^M(\hat{\theta}_i, \mathcal{E}^s)].$$

   (b) Find the parameter draw $\hat{\theta}_0$ that corresponds to $\min\{J_i^s\}_{i=1}^{N_d}$, and calculate $\Sigma_P^{s,0}$.

i. Find the N_{best} draws with the lowest J_i^s. Stack the remaining draws in a $N_{best} \times N_p$ matrix, $\hat{\Theta}^s$, and define $\tilde{\Theta}^s = \hat{\Theta}^s - \mathbf{1}_{N_{best} \times 1} \sum_{i=N_{best}}^{N_d} \hat{\theta}_i^s / N_{best}$.

ii. Calculate $\Sigma_{P,0} = (\tilde{\Theta}^s)' \tilde{\Theta}^s / N_{best}$.

(c) Minimize J with simulated annealing. For $i \in \{0, \ldots, N_d\}$, repeat the following steps:

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

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

We set c_0 to target an average acceptance rate of 50% across seeds.

ii. Under Step 2a, repeats Steps ii-iv.

iii. Accept or reject the candidate draw according to

$$(\hat{\theta}_i^s, J_i^s) = \begin{cases} (\hat{\theta}_i^{cand}, J_i^{s,cand}) & \text{if } i = 0, \\ (\hat{\theta}_i^{cand}, J_i^{s,cand}) & \text{if } \min(1, \exp(J_{i-1}^s - J_i^{s,cand})/c_1) > \hat{u}, \\ (\hat{\theta}_{i-1}, J_{i-1}^s) & \text{otherwise,} \end{cases}$$

where c_1 is the temperature and $\hat{u}$ is a draw from a uniform distribution.

(d) Find the parameter draw $\hat{\theta}_{\min}^s$ that corresponds to $\min\{J_i^s\}_{i=1}^{N_d}$, and update Σ_P^s.

i. Discard the first $N_d/2$ draws. Stack the remaining draws in a $N_d/2 \times N_p$ matrix, $\hat{\Theta}^s$, and define $\tilde{\Theta}^s = \hat{\Theta}^s - \mathbf{1}_{N_d/2 \times 1} \sum_{i=N_d/2}^{N_d} \hat{\theta}_i^s / (N_d/2)$.

ii. Calculate $\Sigma_P^{s,up} = (\tilde{\Theta}^s)' \tilde{\Theta}^s / (N_d/2)$.

(e) Repeat the previous step N_{SMM} times, initializing at draw $\hat{\theta}_0 = \hat{\theta}_{\min}^s$ and covariance matrix $\Sigma_P = \Sigma_P^{s,up}$. Gradually decrease the temperature. Of all the draws, find the lowest J value, denoted J_{guess}^s, and the corresponding draws, θ_{guess}^s.

(f) Minimize the same loss function with MATLAB's `fminsearch` starting at θ_{guess}^s. The minimum is $\hat{\theta}_{\min}^s$ with a loss function value of $J_{\min}^s$. Repeat, each time updating the guess, until $J_{guess}^s - J_{\min}^s < 0.001$. The parameter estimates correspond to $J_{\min}^s$.

The set of SMM parameter estimates $\{\hat{\theta}^s\}_{s=1}^{N_s}$ approximate the joint sampling distribution of the parameters. We report the mean, $\bar{\theta} = \sum_{s=1}^{N_s} \hat{\theta}^s / N_s$, and $(5, 95)$ percentiles. The reported moments are based on the mean parameter estimates, $\bar{\Psi}_T^M = \bar{\Psi}_{R,T}^M(\bar{\theta}, \mathcal{E})$.

We set $N_s = 100$, $R = 1{,}000$, $N_{SMM} = 3$, and $N_J = 1$. N_m, N_d, N_p, and c_1 are all model-specific. The nonlinear solution and estimation algorithms were both programmed in Fortran 95 and executed with Open MPI on the BigTex supercomputer at the Federal Reserve Bank of Dallas.

F ADDITIONAL RESULTS

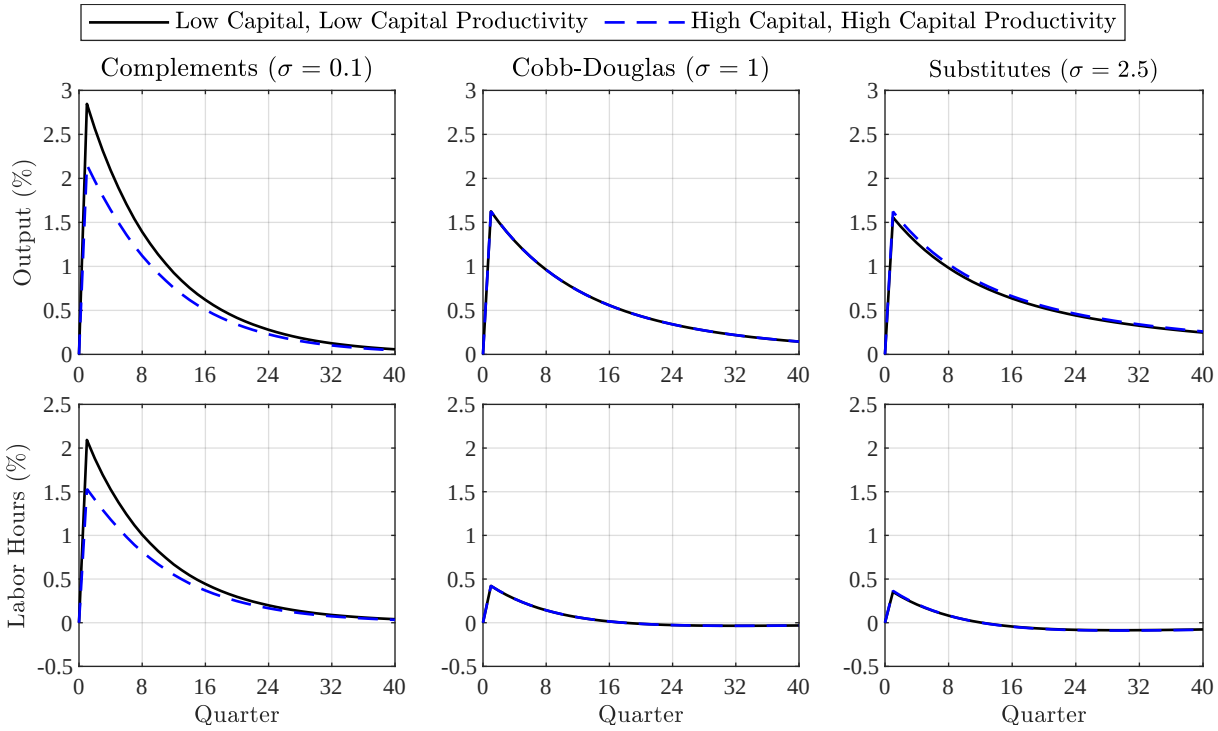

Figure 6: Generalized impulse responses to a 2 standard deviation shock to capital productivity in different states of the economy. To simplify the state, the model only includes capital productivity shocks and there are no real frictions. The capital state is $\pm 2\%$ of steady state. The capital productivity state is $\pm 4.5\%$ of steady state. These values correspond to the $(16, 84)$ percentiles of the ergodic distribution when $\sigma = 0.1$.

Figure 6 shows impulse responses to a capital productivity shock. As in Figure 2, the responses are based on a model without real frictions to highlight the state-dependent effects of the shocks under different degrees of complementarity. In this case, a lower CES strengthens the responses of output and hours because it increases the marginal product of labor. Once again, meaningful state-dependency only occurs with a low CES. The strength of the responses depends on the initial effective capital-to-labor ratio $(z_{k,t}k_{t-1}/(z_{n,t}n_t))$. Periods with low $z_{k,t}$ and k_{t-1} have the largest responses of output and hours, while periods with high $z_{k,t}$ and k_{t-1} have the smallest responses.

Figure 7 shows impulse responses to a labor (top panel) and capital (bottom panel) productivity shock in our baseline model with real frictions and exogenous volatility shocks. The responses are the same as those shown in Figure 5, except we compare different shock sizes and signs to assess the strength of the nonlinearities and asymmetries in the model. Under our baseline parameter estimates, there is very little asymmetry in the model, as the responses to a positive shock are very similar to the responses to a negative shock. The uncertainty response to a 4 standard deviation labor productivity shock is slightly less than double the response to a 2 standard deviation shock, indicating that there is some nonlinearity in uncertainty, but the output responses essentially scale linearly.

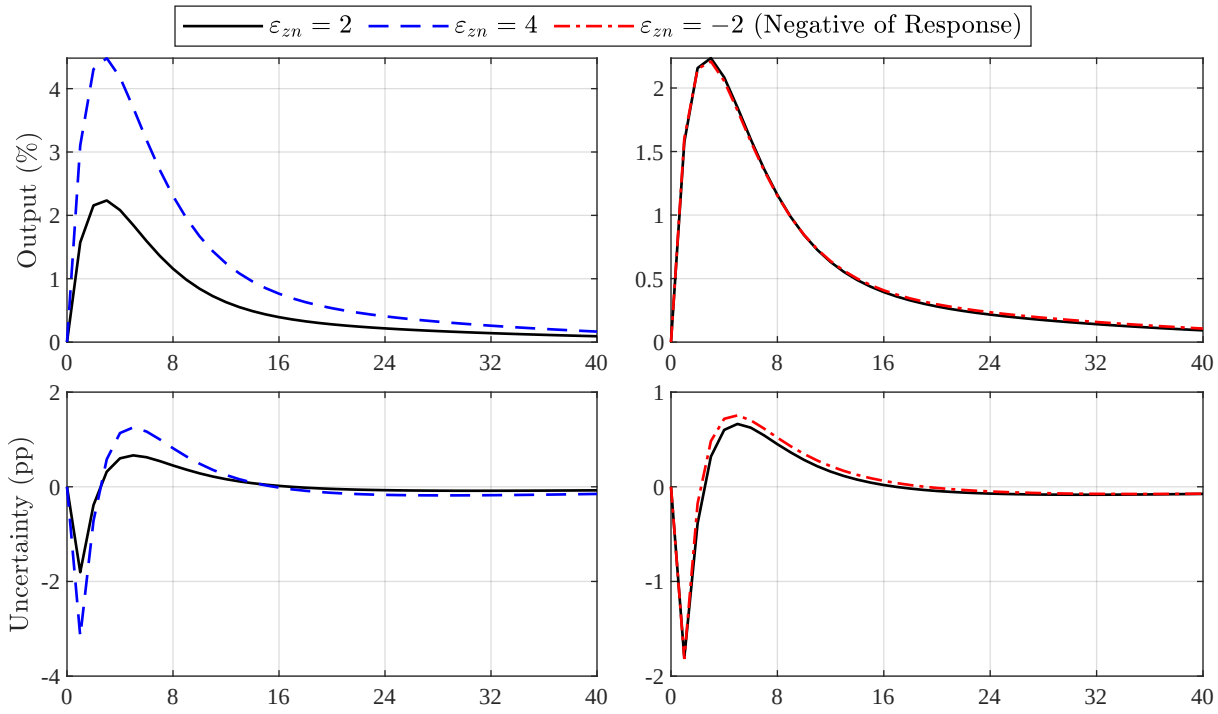

(a) Responses to a labor productivity shock (ε_{zn})

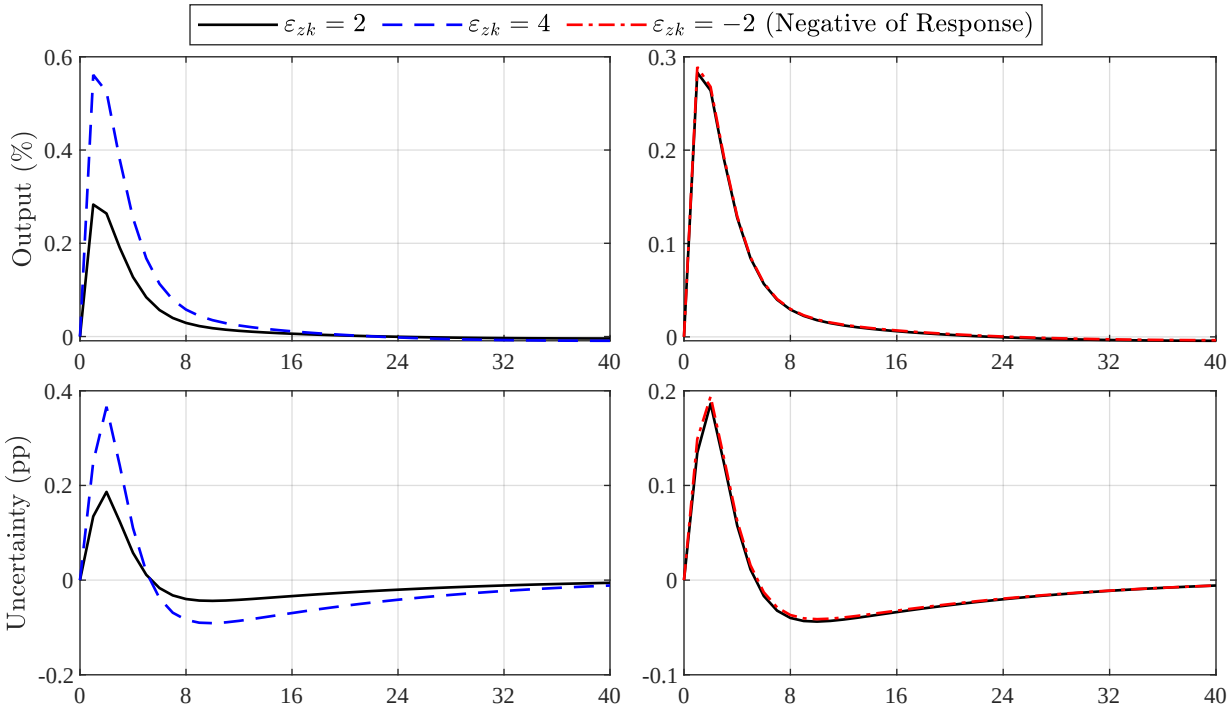

(b) Responses to a capital productivity shock (ε_{zk})

Figure 7: Generalized impulse responses to a shock with various sizes and signs. Output is shown in percent deviations from steady state. Uncertainty is normalized by $SD(\hat{y}^g)$ and shown as a percentage point (pp) difference from steady state. All of the parameters, including the CES, are set to their baseline estimates.