Federal Reserve
Bank of Dallas

Complementarity and Macroeconomic Uncertainty

Tyler Atkinson, Michael Plante, Alexander W. Richter
and Nathaniel A. Throckmorton

Complementarity and Macroeconomic Uncertainty[*]

Tyler Atkinson[†], Michael Plante[‡], Alexander W. Richter[§] and Nathaniel A. Throckmorton[‡]

March 24, 2020

Abstract

Macroeconomic uncertainty—the conditional volatility of the unforecastable component of a future value of a time series—shows considerable variation in the data. A typical assumption in business cycle models is that production is Cobb-Douglas. Under that assumption, this paper shows there is usually little, if any, endogenous variation in output uncertainty, and first moment shocks have similar effects in all states of the economy. When the model departs from Cobb-Douglas production and assumes capital and labor are gross complements, first-moment shocks have state-dependent effects and can cause meaningful variation in uncertainty compared to the data. Estimating several variants of a nonlinear real business cycle model reveals the data strongly prefers a model with high complementarity between capital and labor inputs.

Keywords: State-Dependent; Time-Varying Volatility; CES Production; Nonlinear Estimation

JEL Classifications: C15; D81; E32; E37

1 INTRODUCTION

Macroeconomic uncertainty—the conditional volatility of the unforecastable component of a future value of a time series—shows considerable variation in the data. The literature typically models changes in uncertainty in business cycle models using shocks to the variance of an exogenous variable, such as productivity, while holding its conditional mean fixed.[1] This approach produces estimates for the responses of real activity to an exogenous increase in uncertainty, but it is silent on the effect first moment shocks have on uncertainty. Equivalently, it cannot tell us anything about situations where uncertainty endogenously varies over time in response to the state of the economy.

A common assumption in business cycle models is that production is Cobb-Douglas. Under that assumption, we analytically and numerically show there is typically little, if any, endogenous variation in output uncertainty because the model is close to log-linear. Importantly, the model's ability to generate endogenous fluctuations in uncertainty is limited because log-linearity implies: (1) first moment shocks have nearly the same effect on output in all states of the economy and (2) conditional distributions of future real variables are similar across different states of the economy.

The model can have significant nonlinearity when output is produced using a constant elasticity of substitution (CES) production function. In that case, first moment shocks have state-dependent impacts on output, resulting in conditional distributions that vary across states and endogenous fluctuations in uncertainty. When capital and labor are gross complements, we find first-moment capital and labor productivity shocks generate fluctuations in uncertainty consistent with the data.

Complementarity can endogenously generate fluctuations in uncertainty because it influences how output responds to first moment shocks in different states of the economy, with the effective capital intensity ratio playing a key role. For example, consider a positive labor productivity shock. When the capital intensity ratio is low, the change in labor demand will be relatively small, as will the change in output, compared to situations where the capital-to-labor ratio is high. The distributions of future real variables inherit the state-dependent nature of the responses, generating time-varying endogenous uncertainty. On the other hand, in the Cobb-Douglas case, the responses of labor and output to such a shock are essentially the same regardless of the state of the economy.

An important question is whether the mechanism is empirically relevant in a more fully fledged business cycle model. To that end, we estimate several nonlinear models with real frictions by matching uncertainty, output, consumption, and investment moments in the data. We consider

---

[1]In a small open-economy model, Fernández-Villaverde et al. (2011) examine volatility shocks to a country-specific interest rate spread. They find a 1 standard deviation shock lowers output 0.15%-0.2% in Argentina and Ecuador and 0.01%-0.02% in Brazil and Venezuela. Most papers rely on closed-economy New Keynesian models. Mumtaz and Zanetti (2013) examine monetary policy volatility shocks in a model without capital. They find doubling the volatility reduces output growth by only 0.03%. Born and Pfeifer (2014) introduce variable capital utilization and investment adjustment costs. They show a simultaneous 2 standard deviation increase in uncertainty about government spending, monetary policy, and capital and labor taxes reduces output by only 0.065%. In contrast, Fernández-Villaverde et al. (2015) find a volatility shock to only capital taxes lowers output by 0.1% and the effect is larger when the ZLB binds.

four specifications: the typical model with Cobb-Douglas production and a labor productivity shock and a model with CES production with a labor productivity shock, a capital productivity shock, and both shocks. The Cobb-Douglas specification performs poorly because it is not able to generate nearly enough volatility in uncertainty while simultaneously matching the volatilities and autocorrelations of real activity in the data. All three CES specifications provide a significant improvement in fit and the model with both productivity shocks overwhelmingly passes a test of over-identifying restrictions. In each of these cases, the estimated CES is well below one, indicating that the data prefers a large degree of complementarity between capital and labor inputs.

The paper closest to ours is the recent article by Straub and Ulbricht (2019). They show firm-specific productivity shocks create endogenous fluctuations in the cross-sectional dispersion of output when deviating from Cobb-Douglas production. We build on their work in three ways. First, we examine aggregate uncertainty by computing the conditional volatility of future real activity, rather than the variance in output across firms. Second, we examine the effects of several first-moment shocks, including shocks to capital and labor productivity, the markup of price over marginal cost, government spending, household preferences and monetary policy, in models with and without sticky prices. Third, we conduct a more thorough empirical assessment of the merits of gross complementarity. Specifically, we use a richer model that includes habit persistence in consumption and investment adjustment costs and formally assess its empirical fit using a simulated method of moments procedure that matches key moments related to uncertainty and real activity.

Several other papers consider alternative mechanisms that endogenously generate uncertainty in structural models. 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, collateral constraints, or the zero lower bound constraint on the short-term nominal interest rate (Brunnermeier and Sannikov, 2014; He and Krishnamurthy, 2019; Mendoza, 2010; Plante et al., 2018), while others in this area 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 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, 2020; Straub and Ulbricht, 2015). In these models, adverse shocks under asymmetric learning reduce economic activity and make it harder for households and firms to learn about the economy, which amplifies first moment shocks.

The CES literature provides strong motivation for our proposed mechanism. Klump et al. (2012) surveys empirical estimates of the CES over the last several decades. Across the 17 reported studies, the average CES is 0.58, well below a unitary elasticity implied by Cobb-Douglas production. More recently, Gechert et al. (2019) examine over 3,000 estimates across 121 studies. They estimate a CES of 0.3 after correcting for publication bias and controlling for other aspects

of quality. Antràs (2004) finds that when the econometric specification is adjusted to account for biased technological change, the estimate of the CES is much lower. Consistent with those results, Cantore et al. (2015) estimates a CES between 0.15 and 0.18 using a medium-scale New Keynesian model that accounts for biased technological changes. These values are consistent with our estimates from a real business cycle model that includes both capital and labor productivity shocks. Another benefit of our mechanism is that it affects uncertainty in all periods and is not unique to a specific economic downturn. This is important because uncertainty regularly fluctuates in the data.

Overall, our results demonstrate that an empirically plausible degree of complementarity causes first-moments shocks to generate significant variation in uncertainty. This suggests any study that fails to account for these effects will lead to an over-statement of the causal effects of uncertainty.

The paper proceeds as follows. Section 2 analytically shows why complementarity generates endogenous uncertainty. Section 3 quantifies the amount of time-varying uncertainty in models with productivity shocks. Section 4 estimates several versions of a real business cycle model with both Cobb-Douglas and CES production functions. Section 5 shows numerical results with several other common shocks. Section 6 extends the model to include sticky prices. Section 7 concludes.

2 ANALYTICAL RESULTS

This section develops some understanding for how uncertainty behaves in dynamic models using simplified settings that permit analytical solutions. Following Plante et al. (2018), macroeconomic uncertainty in our model is defined as the expected volatility of a variable in the model. The same definition is also used in empirical work on uncertainty (e.g., Jurado et al. 2015). It is possible to calculate uncertainty over any horizon, but we concentrate on the one-quarter ahead forecast error.

For any variable x, macroeconomic uncertainty is given by

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

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 ($\hat{y}_t^g = \hat{y}_{t+1} - \hat{y}_t$). 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. Given this result, the following sections will always refer to the uncertainty surrounding log output.

2.1 ATHEORETICAL MODEL   A common modeling assumption is that output is produced with a Cobb-Douglas function, so the elasticity of substitution between the factors of production equals 1. In this case, it is possible to derive analytical results without a fully specified theoretical model.

Suppose log output, $\hat{y}_t$, is produced according to a Cobb-Douglas function given by

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

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

$$\hat{x}_{jt} = (1 - \rho_j) \log \bar{x}_j + \rho_j \hat{x}_{jt-1} + \sigma_j \varepsilon_{jt}, \ \varepsilon_j \sim \mathbb{N}(0, 1),$$

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

$$\mathcal{U}_{t,t+1}^{\hat{y}} = \sqrt{\alpha^2 \sigma_1^2 + (1 - \alpha)^2 \sigma_2^2}.$$

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. It is also straightforward to derive analytical expressions for macroeconomic uncertainty over horizons beyond 1 quarter. For example, uncertainty about log output k periods ahead is given by

$$\mathcal{U}_{t,t+k}^{\hat{y}} = \sqrt{\alpha^2 (1 + \rho_1^2 + \cdots + \rho_1^k)\sigma_1^2 + (1 - \alpha)^2 (1 + \rho_2^2 + \cdots + \rho_2^k)\sigma_2^2}$$

and the unconditional volatility, which equals the uncertainty over an infinite horizon, is given by

$$\mathcal{U}^{\hat{y}} = \sqrt{\alpha^2 \sigma_1^2 / (1 - \rho_1^2) + (1 - \alpha)^2 \sigma_2^2 / (1 - \rho_2^2)}.$$

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_j). Deviations from 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 examines the effects of relaxing the first two assumptions, 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,

$$\partial \hat{y}_t / \partial \varepsilon_{1t} = \alpha \sigma_1 \quad \text{and} \quad \partial \hat{y}_t / \partial \varepsilon_{2t} = (1 - \alpha)\sigma_2.$$

At any point in time and for any state of the world, the impact of an unexpected shock to log output is the same. Therefore, the conditional distribution of log output one period in the future does not change shape over time. This result could change if a subset of the four assumptions was violated.

Uncertainty Approximations Taylor expansions are useful for approximating moments of a variable, such as log output, 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

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

where $f'(E[x])$ is the first derivative of $f(x)$ evaluated at its unconditional mean and σ_x^2 is the unconditional variance of x.[2] 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 component random variables, with the weights equal to the first partial derivatives of the function, all evaluated at their respective unconditional means.

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

$$\text{Var}_t[\hat{y}_{t+1}] \approx \alpha^2 \sigma_1^2 + (1 - \alpha^2)\sigma_2^2.$$

This approximation is exact due to the log-linearity of the production function and its two inputs.

Now consider a CES production function given by

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

where α is the cost-share of capital and σ is the elasticity of substitution. It is not possible to analytically solve for log output uncertainty. The approximation for the conditional variance implies

$$\text{Var}_t[\hat{y}_{t+1}] \approx (f_1(E_t\hat{x}_{1t+1}, E_t\hat{x}_{2t+1}))^2\sigma_1^2 + (f_2(E_t\hat{x}_{1t+1}, E_t\hat{x}_{2t+1}))^2\sigma_2^2,$$

where $f_j(\cdot)$ is the partial derivative of $\hat{y}$ with respect to the conditional mean of $\hat{x}_j$. As in the Cobb-Douglas case, the conditional variance for log output is a weighted average of the conditional variances of the two random variables. However, the weights are now state-dependent and vary over time, reflecting the fact that the impact of the two shocks on log output varies across different states of the economy. This generates a time-varying conditional variance under CES production.

Isoquants Figure 1 plots production function isoquants to help visualize how complementarity generates state-dependency. The inner solid line traces out the combinations of the inputs that are consistent with 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, 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

---

[2]Appendix C provides more information about the approximations. See Benaroya et al. (2005) for further details.
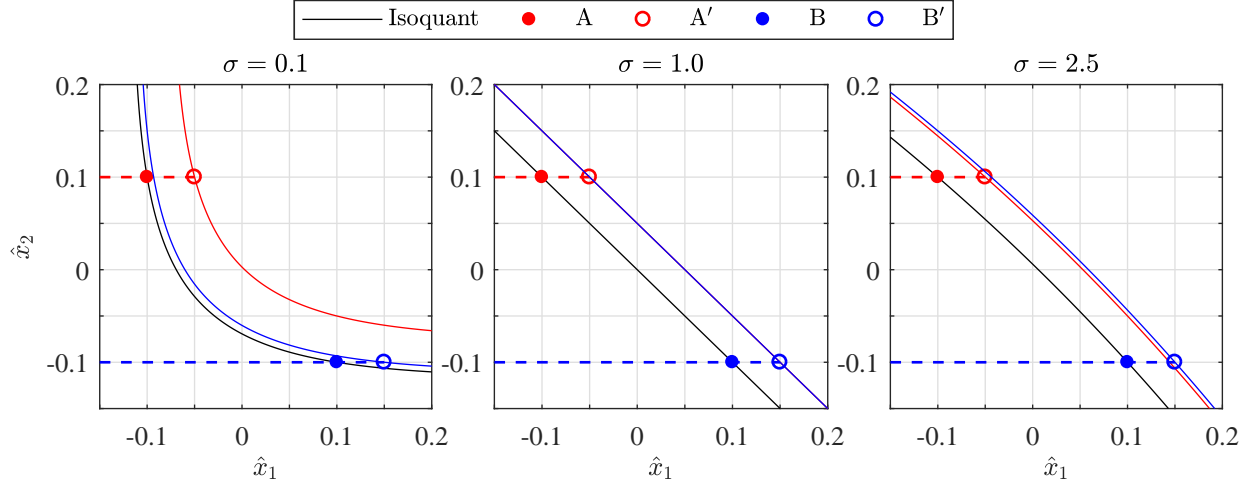
Figure 1: Production function isoquants. Points A′ and B′ are shifted right by 0.05 units of output and α is set to 0.5.

of complementarity, in that a relatively 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 flat. The middle panel shows the same exercise for Cobb-Douglas production. Since the isoquants are linear, increases in an input boost output by the same amount regardless of the initial condition. In the right panel, where the inputs are substitutes ($\sigma = 2.5$), the isoquants are concave rather than convex, meaning an elevated level of an input relative to the other increases its marginal product—the opposite of what happens under complementarity.

The change in $\hat{y}$ relative to $\hat{x}_1$, $1/(1 + \frac{1-\alpha}{\alpha} \exp(\frac{\sigma-1}{\sigma}(\hat{x}_2 - \hat{x}_1)))$, shows the response of output depends on the log difference in the inputs. If x_1 is capital and x_2 is labor, then this difference would correspond to the capital intensity ratio. Changes in one input relative to the other cause variations in uncertainty, whereas that does not necessarily happen due to changes in output. Under Cobb-Douglas production, uncertainty is constant since the responses of output are not state dependent.

2.2 REAL BUSINESS CYCLE MODEL   We now go a step further and introduce a simple model that under some assumptions provides analytical solutions and results regarding uncertainty. Output is produced with capital and labor according to a normalized CES production function given by

$$y_t = \begin{cases} y_0\big(\alpha(z_t^k k_{t-1}/k_0)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)(\bar{g}^t z_t^n n_t/n_0)^{\frac{\sigma-1}{\sigma}}\big)^{\frac{\sigma}{\sigma-1}}, & \sigma \neq 1, \\ y_0(\bar{g}^t z_t^k k_{t-1}/k_0)^\alpha (z_t^n n_t/n_0)^{1-\alpha}, & \sigma = 1, \end{cases} \tag{3}$$

where k_{t-1} is the capital stock, n_t is the labor supply, and $\bar{g}$ is the average growth rate of labor-augmenting productivity. z_t^i, $i \in \{k, n\}$, is a stationary productivity shock that evolves according to

$$\hat{z}_t^i = \rho_{zi}\hat{z}_{t-1}^i + \sigma_{zi}\varepsilon_{zi,t}, \ \varepsilon_{zi} \sim \mathbb{N}(0,1). \tag{4}$$

Following the literature, we include normalizing constants, y_0, n_0 and k_0, to ensure α is the cost-share of capital with both the Cobb-Douglas ($\sigma = 1$) and the CES ($\sigma \neq 1$) production functions.[3]

There is monopolistic competition and an exogenous time-varying price markup given by

$$\hat{\mu}_t = (1 - \rho_\mu) \log \bar{\mu} + \rho_\mu \hat{\mu}_{t-1} + \sigma_\mu \varepsilon_{\mu,t}, \ \varepsilon_\mu \sim \mathbb{N}(0,1). \tag{5}$$

This shock breaks the assumption of a constant cost share of labor, which could have a large impact on log output uncertainty. In symmetric equilibrium, the firm's optimality conditions are given by

$$\mu_t r_t^k = \alpha (y_0 z_t^k / k_0)^{(\sigma-1)/\sigma} (y_t / k_{t-1})^{1/\sigma}, \tag{6}$$

$$\mu_t w_t = (1 - \alpha)(y_0 \bar{g}^t z_t^n / n_0)^{(\sigma-1)/\sigma} (y_t / n_t)^{1/\sigma}. \tag{7}$$

A representative household chooses $\{c_t, n_t, k_t\}_{t=0}^\infty$ to maximize expected lifetime utility, given by $E_0 \sum_{t=0}^\infty \beta^t a_t [c_t^{1-\gamma}/(1-\gamma) - \chi \xi_t n_t^{1+\eta}/(1+\eta)]$, where β is the subjective discount factor, χ determines the steady state labor supply, $1/\eta$ is the Frisch elasticity of labor supply, γ is the coefficient of relative risk aversion, c is consumption, and a and ξ are preference shocks that follow

$$\hat{a}_t = \rho_a \hat{a}_{t-1} + \sigma_a \varepsilon_{a,t}, \ \varepsilon_a \sim \mathbb{N}(0,1), \tag{8}$$

$$\hat{\xi}_t = \rho_\xi \hat{\xi}_{t-1} + \sigma_\xi \varepsilon_{\xi,t}, \ \varepsilon_\xi \sim \mathbb{N}(0,1). \tag{9}$$

The household's choices are constrained by $c_t + k_t = w_t n_t + r_t^k k_{t-1} + (1-\delta)k_{t-1} + d_t$, where δ is the depreciation rate of capital and d_t are firm dividends. The optimality conditions are given by

$$w_t = \chi \xi_t n_t^\eta c_t^\gamma, \tag{10}$$

$$1 = \beta E_t [(a_{t+1}/a_t)(c_t/c_{t+1})^\gamma (r_{t+1}^k + 1 - \delta)]. \tag{11}$$

In equilibrium, dividends are equal to $d_t = (1 - 1/\mu_t)y_t$, so the resource constraint is given by

$$c_t + k_t = (1 - s_t^g)y_t + (1 - \delta)k_{t-1}, \tag{12}$$

where s_t^g is the government spending-to-output ratio. Define $\zeta_t = 1/(1 - s_t^g)$, which follows

$$\hat{\zeta}_t = (1 - \rho_\zeta) \log \bar{\zeta} + \rho_\zeta \hat{\zeta}_{t-1} + \sigma_\zeta \varepsilon_{\zeta,t}, \ \varepsilon_\zeta \sim \mathbb{N}(0,1). \tag{13}$$

Competitive equilibrium includes sequences of quantities $\{y_t, k_t, c_t, n_t, s_t^g\}_{t=0}^\infty$, prices $\{w_t, r_t^k\}_{t=0}^\infty$, and exogenous variables, $\{z_t^n, \mu_t, a_t, \xi_t, \zeta_t\}_{t=1}^\infty$, that satisfy (3)-(13), given $\{k_{-1}, z_0^n, \mu_0, a_0, \xi_0, \zeta_0\}$. Except when estimating the model, the average growth rate, $\bar{g}$, is set to 1 to simplify the exposition.

---

[3]Notably, León-Ledesma et al. (2010) show that a normalized CES production function improves identification, in addition to its theoretical benefits. Gechert et al. (2019) find that normalization reduces the average estimate of the CES by 0.3. See the review article by Klump et al. (2012) for further information about the effects of normalization.

2.3 SOLUTION WITH PRODUCTIVITY SHOCKS   In general, the model does not lend itself to analytical solutions, but it is possible to obtain them when capital fully depreciates each period ($\delta = 1$) and $\gamma = 1/\sigma$. We first consider the model with only the labor productivity shock ($\hat{\mu}_t = \log \bar{\mu}$, $\hat{\zeta}_t = \hat{z}_t^k = \hat{a}_t = \hat{\xi}_t = 0$). The capital-augmenting productivity shock, $\hat{z}_t^k$, is common in the literature on CES production functions. However, it complicates the analytical derivations without providing additional insight, so we defer our examination of this shock to the quantitative sections.

To solve the model, first combine (6) with (11) and (7) with (10). Then guess $c_t = \theta y_t$ to obtain

$$c_t = (1 - (\alpha\beta/\bar{\mu})^\sigma (y_0/k_0)^{\sigma-1}) y_t, \tag{14}$$

$$k_t = (\alpha\beta/\bar{\mu})^\sigma (y_0/k_0)^{\sigma-1} y_t, \tag{15}$$

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

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

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

In contrast with the atheoretical example, only one term appears because α multiplies the capital stock, which is known at time t. It is straightforward to show that if we included both the capital- and labor-augmenting productivity shock, uncertainty would instead equal $\sqrt{\alpha\sigma_{zk}^2 + (1-\alpha)\sigma_{zn}^2}$.

This example satisfies the assumptions discussed above. The production function and the shock process are log-linear, the shock is homoskedastic, and the cost-shares are constant. Furthermore,

$$\partial \hat{y}_t / \partial \varepsilon_{zn,t} = (1 - \alpha) \sigma_{zn}.$$

As a result, a model with Cobb-Douglas production, log utility, and full depreciation cannot generate time-varying log-output uncertainty or state-dependent responses to productivity shocks.

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

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

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

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 (2), first combine (17) and (18) and reorganize to obtain

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

where $\hat{\kappa}_y$ collects the constant terms. Then the conditional variance is approximated by

$$\mathrm{Var}_t[\hat{y}_{t+1}] \approx [f_{\hat{z}}(k_t, E_t \hat{z}_{t+1})]^2 \sigma_z^2.$$

The second derivative of $\hat{y}$ with respect to $\hat{z}$, $f_{\hat{z}\hat{z}}$, determines how the first derivative of the production function changes with different values of $\hat{z}$, and the sign of $f_{\hat{z}\hat{z}}$ depends on the sign of $\sigma - 1$. The weight, $f_{\hat{z}}$, increases with $\hat{z}$ when $\sigma > 1$ and decreases when $\sigma < 1$, so uncertainty is time-varying under CES production and constant in the Cobb-Douglas case ($\sigma = 1$). Importantly, if output uncertainty is countercyclical in the data, then the model is capable of matching the data when $\sigma < 1$. Since the variation in $f_{\hat{z}}$ depends on the CES, data on uncertainty will help inform the CES. Another implication of CES production is that the effect of the shock on log output varies across states of the economy. This is due to the production function having non-zero second derivatives.

2.4 SOLUTIONS WITH OTHER SHOCKS

Our results show that log output uncertainty is constant under Cobb-Douglas production and time-varying under CES production. This section examines whether any of the other four shocks generate time-varying uncertainty in the Cobb-Douglas case.

Leisure Shock The consumption and capital solutions are unchanged but labor, (16), becomes

$$n_t = \left[\tfrac{1-\alpha}{\chi\xi_t} \left(\tfrac{1}{\bar{\mu}^\sigma - (\alpha\beta)^\sigma (y_0/k_0)^{\sigma-1}} \right)^{1/\sigma} \left(\tfrac{y_0 z_t^n}{n_0} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{1+\eta\sigma}}. \tag{19}$$

The shock enters linearly in logs, so log output uncertainty in the Cobb-Douglas case is given by

$$\mathcal{U}_{t,t+1}^{\hat{y}} = (1-\alpha)\sqrt{E_t[\phi_{\hat{z}^n,t+1}^2 + 2\phi_{\hat{z}^n,t+1}\phi_{\hat{n},t+1} + \phi_{\hat{n},t+1}^2]} = (1-\alpha)\sqrt{\sigma_{zn}^2 + \sigma_\xi^2/(1+\eta)^2}, \tag{20}$$

where $\phi_{\hat{x},t+1} \equiv \hat{x}_{t+1} - E_t[\hat{x}_{t+1}]$ is defined as the 1-quarter ahead forecast error for $\hat{x}$. Since the productivity and leisure shocks are uncorrelated, $E_t[\phi_{\hat{z}^n,t+1}\phi_{\hat{\xi},t+1}] = 0$. Output uncertainty is once again constant, but it depends on the variance of both shocks and the Frisch labor supply elasticity.

Preference Shock The consumption-to-output ratio is no longer constant, so there is no closed-form solution with CES production. In the Cobb-Douglas case, we can write the consumption Euler equation, (11), as a difference equation in the output-to-consumption ratio and iterate to obtain

$$y_t/c_t = \sum_{j=0}^\infty (\alpha\beta/\bar{\mu})^j E_t[\exp(\hat{a}_{t+j} - \hat{a}_t)] \equiv a_t^s,$$

where $E_t[\exp(\hat{a}_{t+j})] = \exp(\rho^j \hat{a}_t + \sum_{i=0}^{j-1} \rho^{2i} \sigma_a^2/2)$. Since a_t^s is exogenous, the solution is given by

$$c_t = y_t/a_t^s, \quad k_t = (1 - 1/a_t^s)y_t, \quad n_t = ((1-\alpha)a_t^s/(\chi\bar{\mu}))^{\frac{1}{1+\eta}},$$

and uncertainty about 1-period ahead log output is given by

$$\mathcal{U}_{t,t+1}^{\hat{y}} = (1-\alpha)\sqrt{E_t[\phi_{\hat{z}^n,t+1}^2 + \phi_{\hat{n},t+1}^2]} = (1-\alpha)\sqrt{\sigma_{zn}^2 + \phi_{\hat{a}^s,t+1}^2/(1+\eta)^2}. \qquad (21)$$

When $\rho_a = 0$, $E_t[\exp(\hat{a}_{t+j})] = 1$, so $\hat{a}_t^s$ linearly depends on $\hat{a}_t$. Therefore, $\phi_{\hat{a}^s,t+1}^2 = \sigma_a^2$ and log output uncertainty is constant. When $\rho_a > 0$, uncertainty is time-varying. The production function is log-linear, the weight on the volatilities, $1 - \alpha$, is constant, and the shocks are homoskedastic. However, the conditional volatility of $\hat{a}_t^s$ is time-varying, which causes time-varying uncertainty.[4]

Markup Shock This example is similar to the previous example. Iterate forward to obtain

$$y_t/c_t = 1 + \sum_{j=1}^{\infty}(\alpha\beta)^j E_t[\exp(-\sum_{i=1}^{j}\hat{\mu}_{t+i})] \equiv u_t.$$

Since u_t is exogenous, the solution is given by

$$c_t = y_t/u_t, \quad k_t = (1 - 1/u_t)y_t, \quad n_t = ((1-\alpha)u_t/(\chi\mu_t))^{\frac{1}{1+\eta}}.$$

We show in Appendix D that it is possible to derive an analytical expression for u_t that is a function of only the current value of the markup, μ_t, and parameters associated with its stochastic process.[5]

When $\rho_\mu = 0$, $u_t = 1/(1 - (\alpha\beta/\bar{\mu})\exp(\sigma_\mu^2/2))$. Therefore, the labor forecast error only depends on the markup, which is log-linear. Output uncertainty has the same expression as it does with an i.i.d preference shock, except it now depends on σ_μ^2. When $\rho_\mu > 0$, uncertainty is given by

$$\mathcal{U}_{t,t+1}^{\hat{y}} = (1-\alpha)\sqrt{\sigma_z^2 + (\sigma_\mu^2 + E_t\phi_{\hat{u},t+1}^2 - 2\,\mathrm{Cov}_t(\phi_{\hat{u},t+1}, \phi_{\hat{\mu},t+1}))/(1+\eta)^2}, \qquad (22)$$

which varies over time, just like in (21), because the conditional volatility of $\hat{u}_t$ is time-varying.

Government Spending Shock This example is similar to the markup shock, except $\mu_t = \bar{\mu}$ and

$$y_t/(\zeta_t c_t) = 1 + \sum_{j=1}^{\infty}(\alpha\beta)^j E_t\left[\exp(\sum_{i=1}^{j}\hat{\zeta}_{t+i})\right] = m_t.$$

Since m_t is exogenous, the solution is given by

$$c_t = y_t/(m_t\zeta_t), \quad k_t = (1 - 1/m_t)y_t/\zeta_t, \quad n_t = ((1-\alpha)m_t\zeta_t/(\chi\bar{\mu}))^{\frac{1}{1+\eta}}.$$

---

[4]If the preference shock process was written in levels, instead of logs, then uncertainty would not vary over time.
[5]It is possible to obtain an analytical solution with a CES production function ($\sigma \neq 1$) when the markup shock is i.i.d. ($\rho_\mu = 0$), since the consumption-to-output ratio becomes constant. That derivation is also shown in Appendix D.

Following the same derivations as the markup shock case, output uncertainty is given by

$$\mathcal{U}_{t,t+1}^{\hat{y}} = (1 - \alpha)\sqrt{\sigma_z^2 + (\sigma_\zeta^2 + E_t\phi_{\hat{m},t+1}^2 - 2\operatorname{Cov}_t(\phi_{\hat{m},t+1}, \phi_{\hat{\zeta},t+1})/(1 + \eta)^2}, \qquad (23)$$

which varies over time due to the conditional volatility of $\hat{m}_t$. If the shock was directly to the government spending share, $\hat{s}_t^g$, instead of $\hat{\zeta}_t$, it would introduce another nonlinearity and make it impossible to solve for log output uncertainty. However, the time-varying uncertainty would remain.

3 REAL BUSINESS CYCLE MODEL SIMULATIONS

The previous section assumed full depreciation of capital and set the coefficient of relative risk aversion to the inverse of the elasticity of substitution between capital and labor to facilitate closed-form solutions. This section relaxes those assumptions by numerically solving the nonlinear model.

Our first goal is to assess the amount of endogenous variation in uncertainty the model generates under different degrees of complementarity in the production function. To do so, we simulate the model and report summary statistics assuming different values for σ. We then show generalized impulse response functions and conditional distributions of future output to develop intuition for why the degree of complementarity generates state-dependency and time-varying uncertainty.

Subjective Discount Factor	β	0.9959	Average Markup	$\bar{\mu}$	1.11
Average Labor Share of Income	$\bar{w}\bar{n}/\bar{y}$	0.6031	Frisch Elasticity	$1/\eta$	0.5
Capital Depreciation Rate	δ	0.0247	Shock Persistence	ρ	0.9
Average Government Spending Share	$\bar{s}^g$	0.2081	Shock Standard Deviation	σ	0.02

Table 1: Parameter values.

Table 1 summarizes the parameter values. We specify log utility in consumption ($\gamma = 1$). The Frisch elasticity of labor supply matches the intensive margin estimate in Chetty et al. (2012). The average markup is consistent with Smets and Wouters (2007) and Farhi and Gourio (2018). Many of the other parameters are based on data from 1964Q1-2019Q4, which is the sample we will use to estimate our model in the next section.[6] The discount rate is equal to 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 depreciation rate matches the average rate for private fixed assets and consumer durable goods. The labor share is equal to the average for the total economy. The average government spending share is equal to the fraction of total output that excludes personal consumption expenditures and fixed private investment. For these exercises, the persistence and standard deviation of the shock are set to round values that are within the ranges used in the literature. In the next section, we will estimate the parameters governing the stochastic processes.

---

[6]See Appendix A for a detailed description of our data sources and how they were transformed for our calculations.

The model is the one introduced in Section 2.2, but we now focus on role of both labor and capital-augmenting productivity shocks. Other stochastic processes are set to their steady state values. We use the policy function iteration algorithm described in Richter et al. (2014) to solve the nonlinear model. To isolate the effects of each shock, only one shock enters at any give time. The algorithm is based on the theoretical work on monotone operators in Coleman (1991). We discretize the capital stock and approximate the shock process using the N-state Markov chain in Rouwenhorst (1995). This method is attractive because it does not require us to interpolate on the shock dimensions, which makes the solution more accurate and faster than quadrature methods.

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. The algorithm minimizes the Euler equation errors on every node in the state space and computes the maximum distance between the updated policy functions and the initial conjectures. It then replaces the initial conjectures with the updated policy functions and iterates 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 E for a detailed description of the algorithm.
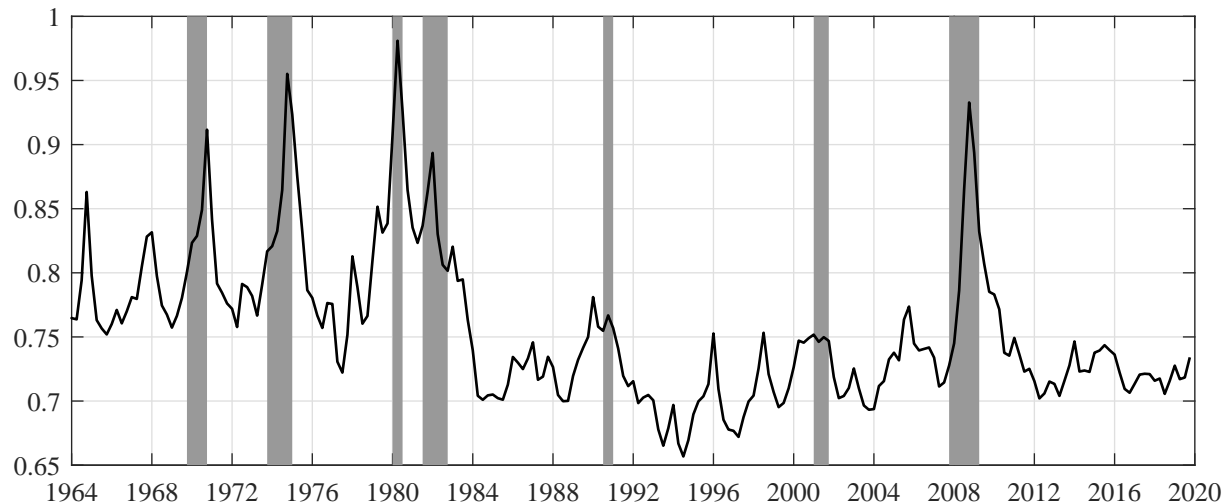
Figure 2: Real uncertainty series from Ludvigson et al. (2020). Shaded regions denote NBER recessions.

3.1 SIMULATED MOMENTS   We begin by simulating the model separately with each shock to assess the amount of endogenous variation in log output uncertainty and the source of the variation over time. Each simulation is initialized with a draw from the ergodic distribution and spans 224 quarters, the same number of observations as our quarterly data sample. For each shock and each simulation, we calculate the standard deviation of log output uncertainty, the standard deviation of output growth, and the correlation between uncertainty and output growth. We execute 1,000 simulations per specification and then report the average and $(5, 95)$ percentiles of the moments.

We compare the simulated moments to equivalent moments in the data. To measure uncertainty in the data, we use the real uncertainty series from Ludvigson et al. (2020), which is shown in Figure 2. This 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). For their analysis, most time series are transformed into growth rates and standard normalized. Repeated simulations of a factor augmented vector autoregression are used to obtain estimates of uncertainty for each real variable and then averaged to generate the aggregate real uncertainty series. The benefit of this series is that it uses the same definition of uncertainty as this paper, so it distinguishes between uncertainty and conditional volatility. To make the units from our model comparable to the real uncertainty series, we calculate $SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$, where $\hat{y}_t^g$ is the 1-quarter log-difference in output. We will refer to this statistic as the normalized volatility of output uncertainty. In the data, output growth equals the log difference in per capita real GDP.

	Data	CES				
		0.2	0.5	1.0	1.5	2.5
		Labor Productivity Shock (z^n)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	5.41 (4.56, 6.40)	1.52 (1.26, 1.82)	0.06 (0.04, 0.07)	0.49 (0.39, 0.62)	1.01 (0.76, 1.31)
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	−0.40	−0.44 (−0.47, −0.42)	−0.41 (−0.44, −0.39)	−0.37 (−0.42, −0.31)	0.36 (0.32, 0.40)	0.34 (0.29, 0.39)
$SD(\hat{y}_t^g)$	0.80	1.15 (1.06, 1.24)	1.50 (1.38, 1.62)	1.68 (1.55, 1.82)	1.76 (1.62, 1.90)	1.83 (1.69, 1.98)
		Capital Productivity Shock (z^k)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	5.54 (4.31, 7.11)	2.14 (1.64, 2.79)	0.03 (0.02, 0.04)	0.85 (0.61, 1.17)	1.59 (1.13, 2.22)
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	−0.40	−0.20 (−0.26, −0.15)	−0.18 (−0.25, −0.13)	−0.37 (−0.42, −0.31)	0.15 (0.09, 0.21)	0.15 (0.09, 0.21)
$SD(\hat{y}_t^g)$	0.80	1.10 (1.00, 1.20)	0.90 (0.83, 0.97)	0.83 (0.76, 0.90)	0.81 (0.75, 0.88)	0.80 (0.74, 0.87)

Table 2: Data and average model-implied moments. The data is quarterly from 1964-2019. The $(5, 95)$ percentiles are based on $1,000$ simulations of the model and shown in parentheses. The standard deviations are all multiplied by 100.

Table 2 shows the data and model-implied moments for different elasticities of substitution. For each shock, the first row shows the normalized volatility of uncertainty, the second row shows the correlation between log output uncertainty and output growth, and the third row shows the standard deviation of output growth. Below each moment, we report the $(5, 95)$ percentiles across the $1,000$ simulations of the model. Notably, the model generates almost no variation in uncertainty under Cobb-Douglas production ($\sigma = 1$), as our analytical results indicate. The insignificant variation that occurs stems from the minor non-linearity that exists when capital does not fully depreciate.

The variation in uncertainty sharply increases with the degree of complementarity. When $\sigma = 0.2$, slightly above the estimated value in Cantore et al. (2015), the normalized volatility of uncertainty is close to the volatility of the real uncertainty series. Importantly, the model achieves

this result while closely matching the volatility of output growth in the data. Both shocks also create a negative correlation between uncertainty and output growth, close to the value in the data.

Substitutability between the capital and labor inputs ($\sigma > 1$) also generates meaningful time-varying uncertainty, but there are two important caveats. First, matching the variation in the data would require unreasonably high substitutability, even over long horizons. With a labor productivity shock this would cause the model to significantly over-predict output growth volatility in the data. Second, both shocks create a positive correlation with output growth, in contrast with the data.
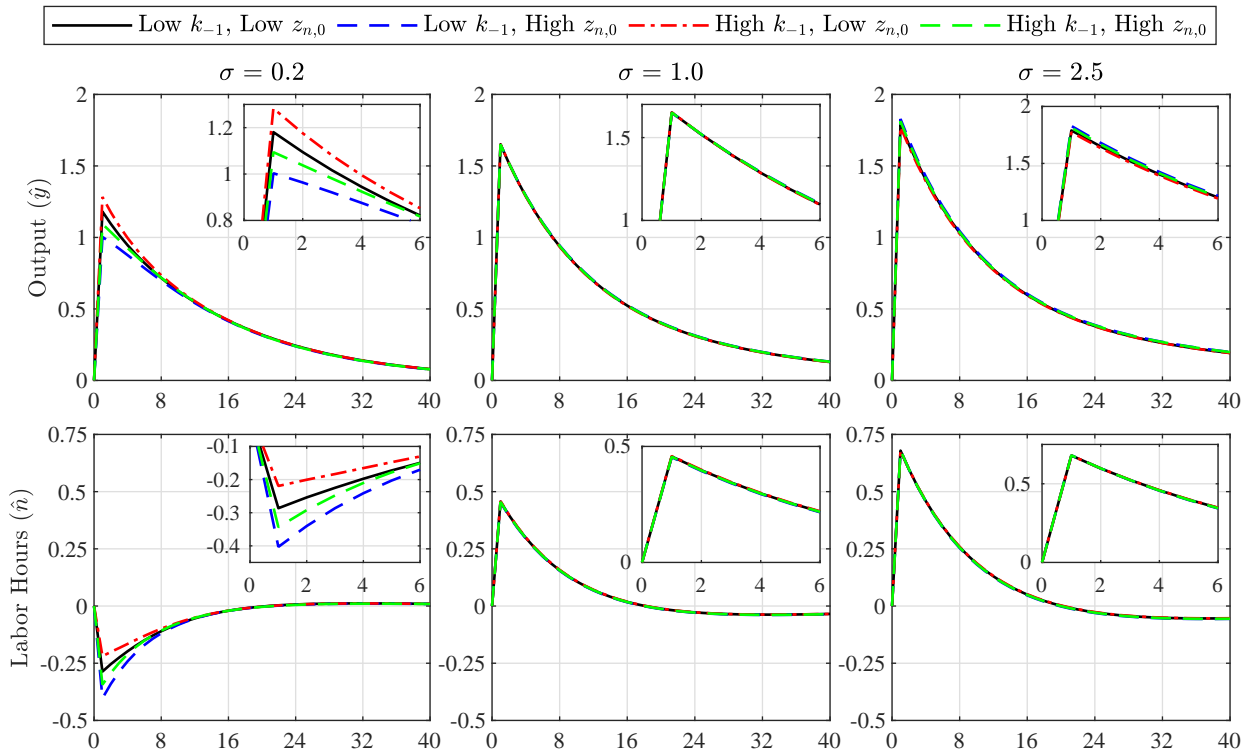
Figure 3: Generalized impulse responses to a 1 standard deviation shock to labor productivity in different states of the economy. The low (high) capital state is 2% below (above) steady state. The low (high) labor productivity state is 4.5% below (above) steady state. The values correspond to the $(16, 84)$ percentiles of the ergodic distribution when $\sigma = 0.2$.

3.2 STATE-DEPENDENT IMPULSE RESPONSES

Meaningful time-varying uncertainty indicates the economy is reacting differently to shocks depending up on the state of the economy. To illustrate the relationship between uncertainty and the degree of state-dependency, Figure 3 plots generalized impulse response functions (GIRFs) of a positive 1 standard deviation shock to labor productivity at different initial conditions and CES parameter values. To compute the GIRFs, we follow the procedure in Koop et al. (1996). We first calculate the mean of 10,000 simulations of a given specification, conditional on random shocks in every quarter. We then calculate a second mean from another set of 10,000 simulations, but this time the shock in the first quarter is replaced with the labor productivity shock. Each GIRF represents the difference between the two mean paths

14

in each specification. The different lines in the graph reflects different initial conditions, while the columns reflect different values of σ. The low (high) initial capital state is 2% below (above) steady state. The low (high) labor initial productivity state is 4.5% percent below (above) steady state. These states correspond to the $(16, 84)$ percentiles of the ergodic distribution when $\sigma = 0.2$.

First, it is important to note that lower CES values dampen the response of output to a labor productivity shock regardless of the initial state. The response of hours worked declines with lower values of σ and is negative when $\sigma = 0.2$. This is driven by the complementarity in the production function. Since capital cannot immediately adjust in response to the shock, higher complementarity (lower σ) reduces the increase in the marginal product of labor from a productivity shock and dampens the increase in labor demand. In the extreme case of perfect complementarity (Leontief production), a labor productivity shock causes a one-for-one decline in hours worked and no change in output, since the marginal product of labor is zero.[7] Complementarity has the opposite (amplifying) effect for the response of output to a capital productivity shock (z^k), because hours worked is free to increase in the same period as the shock. This mechanism explains why the volatility of output growth rises with σ in the top panel of Table 2 and declines in the bottom panel.

Second, state-dependency in the impulse responses occurs when there is a high degree of complementarity, as demonstrated in the left panel where $\sigma = 0.2$. Additionally, the response of output is dependent on the initial effective capital intensity ratio, $z_t^k k_{t-1}/(z_t^n n_t)$, rather than the initial level of output. There is little difference in the responses when the initial capital (k_{-1}) and labor productivity (z_0^n) states are both high or low. When the initial effective capital intensity ratio is low, as it is when k_{-1} is low and z_0^n is high, increases in effective labor increase output relatively little, as illustrated in the isoquants shown in Figure 1. Complementarity prevents firms from taking full advantage of the positive shock because capital cannot immediately increase with labor. In contrast, when k_{-1} is high and z_0^n is low, labor productivity shocks lead to a larger boost in output.

Complementarity also 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 other shocks is elevated, increasing uncertainty. When capital is low and a large positive productivity shock arrives at the start of an expansion, the responses are weak, and future outcomes relatively certain. This is the source of the negative correlation between uncertainty and output growth in the top panel of Table 2.

The pattern of state-dependency is altered in response to a capital productivity shock. Initial periods with low k_{-1} and z_0^k have the largest responses of output, while periods with high k_{-1}

---

[7]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 magnitude of the elasticity of substitution. Canova et al. (2010) shows the response in the data also depends on the type of productivity shock—neutral shocks reduce hours, while investment-specific shocks boost hours. Francis and Ramey (2009) find that labor productivity shocks reduce hours when adjusting the data for the changes in demographics over time.

and z_0^k have the smallest responses. This result is expected given the implied effective capital intensity ratio. The highest levels of output uncertainty then occur at low levels of output, rather than business cycle turning points as is the case with labor productivity shocks. This explains why the correlation of output uncertainty with output growth is weaker in the bottom panel than the top panel of Table 2. While not shown in the table, the correlation of uncertainty with the level of output is much more negative with capital productivity shocks than with labor productivity shocks.

Under Cobb-Douglas production ($\sigma = 1$), there is no state dependency in the impulse response functions, confirming the intuition from the analytical results and isoquant illustration. Similarly, when $\sigma > 1$ the state dependency is in the opposite direction, but by negligible amounts. Overall, the amount of state dependency is small even when $\sigma = 0.2$ and in the extreme states of Figure 3, and these regions are visited relatively infrequently. This suggests only a modest amount of state-dependency in the model is required to generate the variation in uncertainty observed in the data.

3.3 CONDITIONAL 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 10,000 1-period simulations of the model for each of the four states used to generate the GIRFs. We then use the simulated values from each specification to construct a kernel density estimator of the distribution of next period's log output.
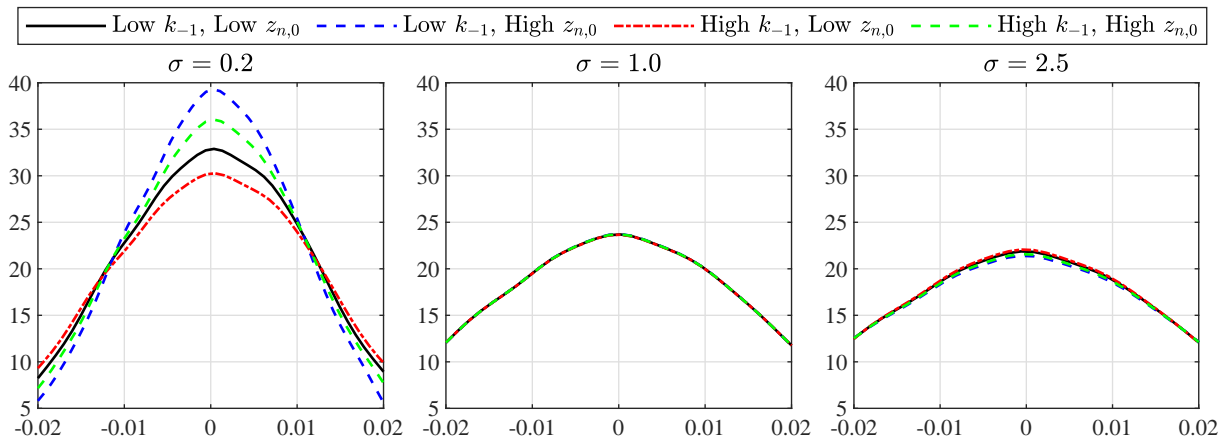
Figure 4: Conditional distributions for log output one period in the future ($\hat{y}_{t+1}$) in different states of the economy.

Figure 4 plots the conditional distribution for the model with labor productivity shocks. Consistent with our previous results, the conditional distributions are essentially invariant to the state of the economy when σ is near 1. However, when $\sigma = 0.2$, the distributions differ depending upon the initial state. In cases with a low initial capital ratio, the distributions are tighter around the conditional mean of the distribution, whereas the distributions are wider in cases where the ratio is initially high. This reflects that in states of the economy with a depressed level of capital relative to labor productivity, changes in labor productivity next period will have a smaller impact on output.

4 REAL BUSINESS CYCLE MODEL ESTIMATES

This section takes the model to the data using a simulated method of moments. We introduce habit persistence in consumption and investment adjustment costs following Boldrin et al. (2001) and Christiano et al. (2005) to match the positive autocorrelations of real activity growth in the data.

The production function, (3), as well as the firm's optimality conditions, (6) and (7), are unchanged. The representative household now chooses $\{c_t, n_t, k_t, x_t\}_{t=0}^{\infty}$ to maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \beta^t [\log(c_t - hc_{t-1}^a) - \chi n_t^{1+\eta}/(1+\eta)]$, where x is investment, c^a is aggregate consumption and h is the degree of external habit persistence. Those choices are constrained by

$$c_t + x_t = w_t n_t + r_t^k k_{t-1} + d_t,$$
$$k_t = (1-\delta)k_{t-1} + x_t(1 - \varphi^x(x_t^{gap} - 1)^2/2),$$

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

$$\lambda_t = c_t - hc_{t-1},$$
$$w_t = \chi n_t^{\eta} \lambda_t,$$
$$1 = q_t[1 - \varphi^x(x_t^{gap} - 1)(3x_t^{gap} - 1)/2] + \beta\bar{g}\varphi^x E_t[(\lambda_t/\lambda_{t+1})q_{t+1}(x_{t+1}^{gap})^2(x_{t+1}^{gap} - 1)],$$
$$q_t = \beta E_t[(\lambda_t/\lambda_{t+1})(r_{t+1}^k + (1-\delta)q_{t+1})].$$

The aggregate resource constraint is given by $c_t + x_t = (1 - \bar{s}^g)y_t$. The model does not have a steady-state due to growth in labor productivity. Variables with a trend are defined in terms of labor productivity (i.e., $\tilde{x}_t \equiv x_t/\bar{g}^t$). The detrended equilibrium system is provided in Appendix B.

Competitive equilibrium consists of sequences of quantities, $\{\tilde{y}, \tilde{k}, \tilde{c}, n, \tilde{x}, \tilde{\lambda}, x_t^{gap}\}_{t=0}^{\infty}$, prices, $\{r^k, \tilde{w}, q\}_{t=0}^{\infty}$, and exogenous variables, $\{z_t^k, z_t^n\}_{t=0}^{\infty}$, that satisfy the detrended equilibrium system, given the initial conditions, $\{\tilde{c}_{-1}, \tilde{k}_{-1}, \tilde{x}_{-1}, z_0^k, z_0^n\}$ and the two shock sequences, $\{\varepsilon_{zk,t}, \varepsilon_{zn,t}\}_{t=1}^{\infty}$.

We estimate the model using data from 1964Q1-2019Q4. Using the transformed data described in Appendix A, the target moments, $\hat{\Psi}_T^D$, are estimated with a two-step Generalized Method of Moments (GMM) estimator, where $T = 224$ is the sample size. Given the GMM estimates, the model is estimated with Simulated Method of Moments (SMM). For parameterization θ and shocks $\mathcal{E}$, we solve the nonlinear model and simulate it $R = 1{,}000$ times for T periods. The model-implied analogues of the target moments 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. We use Monte

Carlo methods to calculate the standard errors on the parameter estimates. For different sequences of shocks, we re-estimate the structural model $N_s = 200$ times and report the mean and $(5, 95)$ percentiles of the parameter estimates. Appendix F provides a detailed description of methodology.

The SMM procedure targets nine moments. They include the standard deviations and first-order autocorrelations of output, consumption, and investment growth, as well as three moments related to uncertainty: the standard deviation of (normalized) output uncertainty, the first-order autocorrelation of output uncertainty, and the cross-correlation of output uncertainty with output growth.

We use a limited information approach to match empirical targets and SMM to account for short-sample bias. To improve on the current methodology, we solve the fully nonlinear model and repeat the estimation procedure for different shock sequences. The solutions and estimations are run in parallel on a supercomputer. The literature typically estimates models once based on a particular seed and uses the Delta method to compute standard errors. While our approach has a higher computational burden, our estimates are independent of the seed and have precise numerical standard errors. The estimates allow us to numerically approximate the sampling distribution of the parameters and test whether they are significantly different across models.[8] We also obtain a distribution of J values, which determine whether a model provides a significant improvement in fit over another model, and the corresponding p-values from a test of over-identifying restrictions.[9]

Table 3 shows the parameter estimates and model-implied moments under four specifications: the CES model with only a labor productivity shock (z^n), only a capital productivity shock (z^k), and both productivity shocks (z^k & z^n) and the Cobb-Douglas model with only the labor productivity shock. In our model, the Cobb-Douglas specification is unable to distinguish between a capital, labor, and total factor productivity shock, so all three versions yield the same overall fit and parameter estimates. The only difference that would occur is that the shock standard deviation would scale to account for differences in the labor cost share in the production function. We report the results with the labor productivity shock. For each parameter, we show the average and $(5, 95)$ percentiles across 200 estimations of the model. For each moment, we report the mean and t-statistic for the null hypothesis that a model-implied moment equals its counterpart in the data.

We begin with the Cobb-Douglas specification.[10] There are only four estimated parameters because the CES is fixed at 1. The estimates are consistent with the literature. The data prefers a

---

[8] 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.

[9] The test statistic is given by $\hat{J}^s = J(\hat{\theta}, \mathcal{E}^s)$, where $\mathcal{E}^s$ is a matrix of shocks given seed s. $J(\hat{\theta}, \mathcal{E})$ converges to a χ^2 distribution with $N_m - N_p$ degrees of freedom, where N_m is the number of empirical targets and N_p is the number of estimated parameters. The $(5, 95)$ percentiles of the p-values determine whether a model reliably passes the test.

[10] Table 6 in Appendix G shows the results of adding fourth-order autocorrelations in real activity growth to our list of targeted moments. These values are informative because the autocorrelation functions in the data are decreasing. The qualitative results are unchanged. All three CES specifications perform significantly better than the Cobb-Douglas specification and the CES model with both productivity shocks passes a test of over-identifying restrictions. The only difference from our main results is that the CES model with only a capital productivity shock no longer passes the test.

| Parameter | C-D | CES | | |
	$\hat{z}^n$	$\hat{z}^n$	$\hat{z}^k$	$\hat{z}^n$ & $\hat{z}^k$
ρ_{zn}	0.9571 (0.9529, 0.9609)	0.9839 (0.9822, 0.9859)	–	0.9626 (0.9555, 0.9724)
σ_{zn}	0.0211 (0.0209, 0.0214)	0.0166 (0.0159, 0.0173)	–	0.0199 (0.0182, 0.0211)
ρ_{zk}	–	–	0.9152 (0.9125, 0.9182)	0.9664 (0.9414, 0.9788)
σ_{zk}	–	–	0.0316 (0.0311, 0.0321)	0.0071 (0.0066, 0.0076)
φ_x	2.9470 (2.8521, 3.0550)	0.9074 (0.7465, 1.0527)	2.1571 (2.0987, 2.2216)	1.7650 (1.3480, 2.0866)
h	0.8998 (0.8965, 0.9030)	0.5075 (0.4053, 0.5893)	0.8859 (0.8843, 0.8876)	0.7798 (0.7224, 0.8143)
σ	–	0.1739 (0.1713, 0.1767)	0.3945 (0.3848, 0.4040)	0.1325 (0.1178, 0.1430)
J	53.41 (52.94, 53.91)	12.51 (12.01, 12.95)	6.75 (6.62, 6.85)	1.88 (1.72, 2.10)
pval	0.00 (0.00, 0.00)	0.01 (0.01, 0.02)	0.15 (0.14, 0.16)	0.39 (0.35, 0.42)
df	5	4	4	2

(a) Average and $(5, 95)$ percentiles of the parameter estimates.

| Moment | Data | C-D | CES | | |
		$\hat{z}^n$	$\hat{z}^n$	$\hat{z}^k$	$\hat{z}^n$ & $\hat{z}^k$
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.76	1.98 (-6.15)	6.24 (0.78)	5.71 (-0.09)	5.71 (-0.08)
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	-0.40	-0.33 (0.84)	-0.58 (-1.95)	-0.35 (0.56)	-0.46 (-0.65)
$AC(\mathcal{U}^{\hat{y}}_{t,t+1})$	0.91	0.90 (-0.12)	0.91 (-0.07)	0.88 (-0.74)	0.88 (-0.82)
$SD(\hat{y}^g)$	0.80	0.95 (2.33)	0.80 (0.04)	0.88 (1.32)	0.81 (0.19)
$SD(\hat{c}^g)$	0.51	0.40 (-2.20)	0.41 (-1.97)	0.49 (-0.45)	0.52 (0.20)
$SD(\hat{x}^g)$	2.02	2.11 (0.44)	2.13 (0.54)	1.76 (-1.35)	1.95 (-0.40)
$AC(\hat{y}^g)$	0.29	0.46 (1.89)	0.17 (-1.38)	0.35 (0.60)	0.24 (-0.58)
$AC(\hat{c}^g)$	0.41	0.46 (0.77)	0.31 (-1.35)	0.46 (0.75)	0.43 (0.33)
$AC(\hat{x}^g)$	0.42	0.47 (0.56)	0.41 (-0.09)	0.32 (-1.10)	0.41 (-0.07)

(b) Data and average model-implied moments. t-statistics are in parentheses.

Table 3: Parameter estimates and targeted moments.

highly persistent labor productivity process ($\rho_{zn} = 0.96$), strong habit persistence ($h = 0.9$), and significant investment adjustment costs ($\varphi^x = 2.9$). The model has a poor fit. The J value is 53.4 and the p-value is essentially 0, meaning the data soundly rejects this specification. The biggest reason is because the model cannot generate the volatility in uncertainty that occurs in the data. Normalized uncertainty is considerably larger than the near zero values reported in Section 3 because the real frictions introduce two additional backward looking components that generate greater non-

linearities and state-dependency than in the frictionless model. However, the t-statistic is still 6.15, meaning the volatility of uncertainty in the model is significantly different from the data. The model also has trouble matching the standard deviations of output and consumption growth, as both t-statistics exceed 2. Overall, these three moments account for roughly 90% of the J value.

All three specifications with a CES production function perform significantly better than the Cobb-Douglas model. While there is some variation in the estimates, the data prefers strong complementarity in the production function ($\sigma < 0.4$) under all three shock specifications. In the model with a productivity shock, the estimated CES is 0.17 and the J value drops to 12.5. Despite the significant improvement in fit, the p-value is still only 0.01, mainly because the model over-predicts the correlation between output growth and uncertainty and under-predicts the standard deviation of consumption growth. The model with a capital productivity shock performs better, even though the estimated CES is higher ($\sigma = 0.39$). The J value is 6.8 and the p-value increases to 0.15, passing the test of over-identifying restrictions at the 10% level. The model with both labor and capital productivity shocks performs the best. Despite only two degrees of freedom, the J value declines to 1.9 and the p-value further increases to 0.39. None of the t-statistics exceed 1, indicating that none of the model-implied moments are statistically different from their counterparts in the data. Furthermore, the CES is only 0.13, pointing to strong complementarity between capital and labor.

5 MOVING BEYOND PRODUCTIVITY SHOCKS

Table 4 repeats the exercise in Table 2 for the remaining shocks discussed in Section 2 under our baseline calibration. We begin with a discussion of the results under Cobb-Douglas production ($\sigma = 1$), which confirm that our analytical results carry through for our model without full depreciation. In particular, shocks to the price markup (μ), government spending (ζ), and preferences (a) generate noticeable time variation in output uncertainty. This stems from the conditional volatility that is created by the way the exogenous shock enters the equilibrium system, rather than state dependence through the endogenous capital state (k_{t-1}). When μ_0 or a_0 is elevated, further increases in the shock have a smaller impact on output than when either of those states are low. The opposite holds for the government spending shock—an elevated ζ_0 amplifies the response of output to further increases in government spending. Shocks to leisure preferences (ξ) do not generate meaningful variation in uncertainty with Cobb-Douglas production, similar to the productivity shocks.

Government spending shocks generate the largest variation in uncertainty, nearly 20% of what is in the data and several times larger than what the markup shock generates. This may seem at odds with our analytical results, since the expressions for uncertainty were nearly identical for the two shocks. However, both produce nearly the same value for the un-normalized level of uncertainty ($SD(\mathcal{U}_{t,t+1}^{\hat{y}})$). The differences occur because capital and hours worked move in opposite directions in response to a government spending shock, which dampens the volatility of output

	Data	CES				
		0.2	0.5	1.0	1.5	2.5
		Markup Shock (μ)				
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.76	0.93 $(0.77, 1.12)$	0.22 $(0.15, 0.31)$	0.17 $(0.14, 0.21)$	0.27 $(0.22, 0.33)$	0.38 $(0.30, 0.48)$
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	-0.40	-0.48 $(-0.51, -0.44)$	-0.31 $(-0.40, -0.22)$	0.35 $(0.31, 0.39)$	0.38 $(0.35, 0.41)$	0.36 $(0.32, 0.40)$
$SD(\hat{y}^g_t)$	0.80	0.22 $(0.20, 0.24)$	0.29 $(0.27, 0.32)$	0.33 $(0.31, 0.36)$	0.35 $(0.33, 0.38)$	0.37 $(0.34, 0.40)$
		Government Spending Shock (ζ)				
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.76	2.85 $(2.03, 3.93)$	0.67 $(0.55, 0.80)$	1.02 $(0.72, 1.41)$	1.34 $(0.89, 1.95)$	1.61 $(1.05, 2.40)$
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	-0.40	0.03 $(-0.03, 0.10)$	-0.13 $(-0.21, -0.05)$	-0.17 $(-0.23, -0.10)$	-0.13 $(-0.20, -0.07)$	-0.08 $(-0.16, -0.00)$
$SD(\hat{y}^g_t)$	0.80	0.25 $(0.23, 0.27)$	0.25 $(0.23, 0.27)$	0.22 $(0.20, 0.23)$	0.19 $(0.18, 0.21)$	0.17 $(0.16, 0.18)$
		Preference Shock (a)				
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.76	3.11 $(2.07, 4.43)$	1.39 $(0.87, 2.04)$	0.61 $(0.45, 0.82)$	0.50 $(0.42, 0.61)$	0.74 $(0.55, 1.00)$
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	-0.40	-0.25 $(-0.34, -0.16)$	-0.07 $(-0.12, -0.03)$	0.22 $(0.16, 0.29)$	0.43 $(0.41, 0.45)$	0.39 $(0.32, 0.46)$
$SD(\hat{y}^g_t)$	0.80	0.22 $(0.20, 0.24)$	0.34 $(0.31, 0.36)$	0.42 $(0.39, 0.45)$	0.46 $(0.43, 0.50)$	0.51 $(0.47, 0.55)$
		Leisure Preference Shock (ξ)				
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.76	1.84 $(1.55, 2.18)$	0.51 $(0.42, 0.61)$	0.02 $(0.02, 0.03)$	0.16 $(0.13, 0.21)$	0.33 $(0.25, 0.44)$
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	-0.40	-0.44 $(-0.47, -0.43)$	-0.41 $(-0.44, -0.39)$	-0.37 $(-0.42, -0.31)$	0.36 $(0.32, 0.40)$	0.34 $(0.29, 0.39)$
$SD(\hat{y}^g_t)$	0.80	0.38 $(0.35, 0.41)$	0.50 $(0.46, 0.54)$	0.56 $(0.52, 0.61)$	0.59 $(0.54, 0.63)$	0.61 $(0.56, 0.66)$

Table 4: Data and average model-implied moments. The data is quarterly from 1964-2019. The $(5, 95)$ percentiles are based on 1,000 simulations of the model and shown in parentheses. The standard deviations are all multiplied by 100. The persistence of each process is 0.9. The shock standard deviations are 0.02, except the markup shock which is 0.01.

$(SD(\hat{y}^g_t))$. While government spending shocks generate sizable volatility in uncertainty relative to the volatility of output growth, they would require unreasonably large variations in the output share of government spending to generate meaningful variation in the un-normalized level of uncertainty.

When deviating from Cobb-Douglas production, the moments in Table 4, with the exception of those generated from government spending shocks, move in a similar direction as shown with the labor productivity shock in Table 2. As σ falls, the volatility of output growth declines, the volatility of uncertainty increases, and the correlation of uncertainty and output growth becomes more negative. Complementarity introduces state dependency in the impulse responses in a similar way as the labor productivity shock, through the effective capital ratio. Initial periods with a higher effective capital ratio have a larger response in hours worked and output. When $\sigma = 0.2$, these shocks generate variation in uncertainty between one-third and one-half of what is in the data. The correlation between uncertainty and the effective capital ratio exceeds 0.9 for each shock, confirm-

ing the mechanism that generates uncertainty. These results demonstrate that complementarity has meaningful implications for time-varying uncertainty for most shocks, not just productivity shocks.

Government spending shocks are the exception to this pattern. The variation in uncertainty is relatively stable across the values of σ, suggesting the conditional volatility that stems from the way the shock hits the equilibrium system dominates the effect of complementarity. The correlation of output growth and uncertainty is positive and increases as σ falls, in contrast with the other shocks.

6 NEW KEYNESIAN MODEL SIMULATIONS

This section extends the real business cycle model in Section 2.2 to include sticky prices. Following Rotemberg (1982), firms choose their labor, capital, and price level to maximize the expected discounted present value of future dividends subject to a quadratic price adjustment cost. Using the functional form in Ireland (1997), the optimality conditions in symmetric equilibrium are given by

$$r_t^k = \alpha m c_t \left(y_0 z_t^k / k_0 \right)^{\frac{\sigma-1}{\sigma}} \left(y_t / k_{t-1} \right)^{\frac{1}{\sigma}},$$

$$w_t = (1 - \alpha) m c_t \left(y_0 \bar{g}^t z_t^n / n_0 \right)^{\frac{\sigma-1}{\sigma}} \left(y_t / n_t \right)^{\frac{1}{\sigma}},$$

$$\varphi^p (\pi_t / \bar{\pi} - 1) \pi_t / \bar{\pi} = (\mu_t m c_t - 1) / (\mu_t - 1) + \beta \varphi^p E_t[(c_t / c_{t+1})(\pi_{t+1} / \bar{\pi} - 1)(\pi_{t+1} / \bar{\pi})(y_{t+1} / y_t)],$$

where π is the gross inflation rate and φ^p scales the size of the price adjustment cost. When $\varphi^p > 0$, the real marginal cost of producing an additional unit of output, mc, is endogenously determined by the Phillips Curve. When $\varphi^p = 0$, $mc_t = 1/\mu_t$ and the model collapses the version in Section 2.2.

The household's budget constraint is given by $c_t + b_t / i_t = w_t n_t + b_{t-1} / \pi_t + d_t$, where i is the gross nominal interest rate set by the central bank and b is the real value of a privately-issued 1-period nominal bond in zero net supply. The optimality conditions are given by (10), (11), and

$$1 = \beta E_t[(c_t / c_{t+1})(i_t / \pi_{t+1})].$$

The monetary policy rule is given by

$$i_t = i_{t-1}^{\rho_i} (\bar{i}(\pi_t / \bar{\pi})^{\phi_\pi} (y_t / (\bar{g} y_{t-1}))^{\phi_y})^{1-\rho_i} \exp(\varepsilon_{i,t}), \ \varepsilon_i \sim \mathbb{N}(0, 1),$$

where $\bar{\pi}$ is the inflation target, ϕ_π and ϕ_y are the responses to the inflation and output growth gaps, and ρ_i is the persistence of the policy rule. Following Eggertsson and Singh (2019), we exclude the price adjustment costs from the aggregate resource constraint, so it is unchanged from Section 2.2.

We set $\phi_\pi = 2.0$, $\phi_y = 0.2$, $\rho_i = 0.8$, and $\sigma_i = 0.0025$, which are all common values. The price adjustment cost parameter is set to 51.49. This value corresponds to a Calvo price setting where prices change once every three quarters, the mean estimate in Smets and Wouters (2007). The inflation target is set to 1.0084, which is the average GDP deflator inflation rate in our sample.

	Data	CES				
		0.2	0.5	1.0	1.5	2.5
		Labor Productivity Shock (z^n)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	4.37 $(3.65, 5.20)$	1.08 $(0.90, 1.30)$	0.10 $(0.08, 0.13)$	0.53 $(0.42, 0.67)$	0.93 $(0.72, 1.20)$
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	-0.40	-0.38 $(-0.41, -0.37)$	-0.39 $(-0.42, -0.37)$	0.38 $(0.33, 0.43)$	0.38 $(0.34, 0.42)$	0.37 $(0.31, 0.42)$
$SD(\hat{y}_t^g)$	0.80	1.43 $(1.31, 1.55)$	1.58 $(1.46, 1.71)$	1.63 $(1.50, 1.76)$	1.64 $(1.52, 1.77)$	1.66 $(1.53, 1.79)$
		Capital Productivity Shock (z^k)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	9.27 $(7.28, 11.78)$	2.82 $(2.17, 3.68)$	0.05 $(0.04, 0.06)$	0.89 $(0.68, 1.16)$	1.54 $(1.16, 2.05)$
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	-0.40	-0.26 $(-0.34, -0.20)$	-0.22 $(-0.29, -0.16)$	0.38 $(0.33, 0.43)$	0.20 $(0.15, 0.26)$	0.18 $(0.13, 0.24)$
$SD(\hat{y}_t^g)$	0.80	0.82 $(0.74, 0.91)$	0.76 $(0.70, 0.82)$	0.80 $(0.74, 0.87)$	0.84 $(0.77, 0.91)$	0.88 $(0.81, 0.96)$
		Markup Shock (μ)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	4.66 $(3.73, 5.74)$	3.46 $(2.74, 4.30)$	2.82 $(2.22, 3.55)$	2.54 $(2.00, 3.20)$	2.28 $(1.78, 2.89)$
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	-0.40	-0.25 $(-0.29, -0.22)$	-0.24 $(-0.28, -0.19)$	-0.24 $(-0.28, -0.19)$	-0.24 $(-0.29, -0.19)$	-0.24 $(-0.29, -0.19)$
$SD(\hat{y}_t^g)$	0.80	0.36 $(0.32, 0.39)$	0.42 $(0.38, 0.46)$	0.45 $(0.41, 0.49)$	0.46 $(0.42, 0.50)$	0.47 $(0.43, 0.51)$
		Monetary Policy Shock (ε_i)				
$SD(\mathcal{U}_{t,t+1}^{\hat{y}})/SD(\hat{y}_t^g)$	5.76	0.17 $(0.12, 0.22)$	0.06 $(0.05, 0.08)$	0.04 $(0.04, 0.05)$	0.05 $(0.05, 0.07)$	0.07 $(0.05, 0.09)$
$Corr(\mathcal{U}_{t,t+1}^{\hat{y}}, \hat{y}_t^g)$	-0.40	-0.29 $(-0.36, -0.21)$	-0.56 $(-0.66, -0.46)$	-0.65 $(-0.70, -0.60)$	-0.46 $(-0.54, -0.36)$	-0.32 $(-0.42, -0.22)$
$SD(\hat{y}_t^g)$	0.80	0.83 $(0.75, 0.90)$	1.12 $(1.02, 1.22)$	1.28 $(1.16, 1.39)$	1.34 $(1.22, 1.46)$	1.39 $(1.27, 1.52)$

Table 5: Data and average model-implied moments. The data is quarterly from 1964-2019. The $(5, 95)$ percentiles are based on 1,000 simulations of the model and shown in parentheses. The standard deviations are all multiplied by 100. The shock standard deviations are set to 0.02 (productivity), 0.01 (price markup), and 0.0025 (monetary policy).

Table 5 once again repeats the exercise in table 2 for the model with sticky prices. The results with the capital and labor productivity shocks are qualitatively unchanged. There is little time-varying uncertainty under Cobb-Douglas production and meaningful, counter-cyclical variation when capital and labor are gross complements. Quantitatively, however, there are some important differences. Strong complementarity causes capital productivity shocks to generate considerably more time-varying uncertainty and less output growth volatility than our baseline model. These dynamics are driven almost entirely by the interest rate inertia in the monetary policy rule. When k_0 is low, there is a large response of output. The persistence in the policy rule keeps interest rates lower for longer, amplifying the response of output. Conversely, when k_0 is high, the response of output is weak due to complementarity, which mitigates the effects of the interest rate persistence.

Price markup shocks generate meaningful time-varying uncertainty even under Cobb-Douglas production. This stems from the endogeneity of marginal costs. Positive markup shocks always

reduce factor prices, but when prices are sticky, the shock also increases current and future infla-tion. This feeds into marginal costs and further reduces factor prices. Since price adjustment costs are quadratic, marginal costs and hence factor prices decline more in states that are farther away from steady state. This generates state-dependent output responses and time-varying uncertainty.

Monetary policy shocks cause almost no variation in uncertainty, regardless of the elasticity of substitution in production. The small variation in uncertainty occurs even though the volatility of output growth is larger than it is in the data. Overall, these results reinforce that productivity shocks under complementarity are an important driver of the variation in uncertainty observed in the data.

7 CONCLUSION

Much of the uncertainty literature uses structural models to examine the effects of shocks to the variance of exogenous variables, while holding their conditional means fixed. In this paper, our fo-cus is on endogenous movements in uncertainty that are due to changes in the state of the economy. This can occur in any model where first moment shocks have state-dependent effects on output. However, this channel is weak in many business cycle models because many are close to log linear.

One source of the near-linearity in business cycle models stems from the common assumption of Cobb-Douglas production, which implies a unitary elasticity of substitution between inputs in production. This paper shows that if capital and labor are gross complements, first moment pro-ductivity shocks can generate time-varying and counter-cyclical uncertainty consistent with the data. While there are other mechanisms in the literature, ours is particularly compelling for two main reasons. First, there is overwhelming evidence in the literature that the elasticity of substitu-tion between inputs in the production process is well below unity. Second, complementarity can explain why uncertainty fluctuates at all points of the business cycle, not just during recessions or isolated events. More broadly, our results demonstrate that failing to account for the effects of first-moments shocks on uncertainty will lead to an over-statement of the causal effects of uncertainty.

There are several avenues to build on our results. We focus on uncertainty about future real activity. Developing richer models that include a financial sector or additional frictions would make it possible to examine whether the model can match alternative sources of uncertainty, such as those surrounding prices and interest rates. Building exogenous volatility shocks into any model where there are meaningful state-dependent effects of first moment shock would make it possible to examine the causal effect of various sources of uncertainty, while controlling for feedback from the state of the economy. Both of these extensions would further advance the uncertainty literature.

REFERENCES

ANTRÀS, P. (2004): "Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution," *The B.E. Journal of Macroeconomics*, 4, 1–36, https://doi.org/10.2202/1534-6005.1161.

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.

BOLDRIN, M., L. J. CHRISTIANO, AND J. D. M. FISHER (2001): "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 91, 149–166, https://doi.org/10.1257/aer.91.1.149.

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.

CANOVA, F., D. LOPEZ-SALIDO, AND C. MICHELACCI (2010): "The effects of technology shocks on hours and output: a robustness analysis," *Journal of Applied Econometrics*, 25, 755–773, https://doi.org/10.1002/jae.1090.

CANTORE, C., M. LEN-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.

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, https://doi.org/10.1086/426038.

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

EGGERTSSON, G. B. AND S. R. SINGH (2019): "Log-linear approximation versus an exact solution at the ZLB in the New Keynesian model," *Journal of Economic Dynamics and Control*, 105, 21–43, 10.1016/j.jedc.2018.09.01.

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.

FARHI, E. AND F. GOURIO (2018): "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia," *Brookings Papers on Economic Activity*, 49, 147–250, https://doi.org/10.1353/eca.2018.0024.

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.

——— (2009): "Measures of per Capita Hours and Their Implications for the Technology-Hours Debate," *Journal of Money, Credit and Banking*, 41, 1071–1097, https://doi.org/10.1111/j.1538-4616.2009.00247.x.

GECHERT, S., T. HAVRANEK, Z. IRSOVA, AND D. KOLCUNOVA (2019): "Death to the Cobb-Douglas Production Function," IMK Working Paper 201.

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

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

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.

IRELAND, P. N. (1997): "A Small, Structural, Quarterly Model for Monetary Policy Evaluation," *Carnegie-Rochester Conference Series on Public Policy*, 47, 83–108, https://doi.org/10.1016/S0167-2231(98)00005-0.

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.

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.

LEÓN-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.

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.

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.

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.

ROTEMBERG, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211, https://doi.org/10.1086/261117.

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.

SMETS, F. AND R. WOUTERS (2007): "Shocks and frictions in us business cycles: a bayesian dsge approach," *American Economic Review*, 97, 586–606, https://doi.org/10.1257/aer.97.3.586.

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 DATA SOURCES

We use the following quarterly time-series from 1964Q1-2019Q4 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. **Gross Domestic Product**, Seasonally Adjusted, Billions of Dollars, (GDP@USECON)

4. **Personal Consumption Expenditures: Nondurable Goods**,
   Seasonally Adjusted, Billions of Dollars (CN@USECON)

5. **Personal Consumption Expenditures: Services**,
   Seasonally Adjusted, Billions of Dollars (CS@USECON)

6. **Personal Consumption Expenditures: Durable Goods**,
   Seasonally Adjusted, Billions of Dollars (CD@USECON)

7. **Private Fixed Investment**, Seasonally Adjusted, Billions of Dollars (F@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 Ludvigson et al. (2020) Real Uncertainty Index, available on the authors' personal websites. We use a quarterly average of monthly values for $h = 3$ (one quarter forecast horizon).

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**:

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

5. **Capital Depreciation Rate**:

$$\delta = (KPT + KDT) / (EPT + EDT).$$

B DETRENDED EQUILIBRIUM SYSTEM

The detrended real business cycle model includes the exogenous productivity processes, (4), and

$$\mu_t r_t^k = \alpha \left(y_0 z_t^k / k_0 \right)^{\frac{\sigma-1}{\sigma}} \left(\bar{g} \tilde{y}_t / \tilde{k}_{t-1} \right)^{\frac{1}{\sigma}},$$

$$\tilde{y}_t = y_0 \left[\alpha \left(z_t^k \tilde{k}_{t-1} / (\bar{g} k_0) \right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha) \left(z_t^n n_t / n_0 \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

$$\mu_t \tilde{w}_t = (1-\alpha) \left(y_0 z_t^n / n_0 \right)^{\frac{\sigma-1}{\sigma}} \left(\tilde{y}_t / n_t \right)^{\frac{1}{\sigma}},$$

$$\tilde{w}_t = \chi n_t^{\eta} \tilde{\lambda}_t,$$

$$\tilde{\lambda}_t = \tilde{c}_t - (h/\bar{g}) \tilde{c}_{t-1},$$

$$\tilde{c}_t + \tilde{x}_t = (1 - \bar{s}^g) \tilde{y}_t,$$

$$\tilde{k}_t = (1-\delta) \tilde{k}_{t-1} / \bar{g} + \tilde{x}_t (1 - \varphi^x (x_t^{gap} - 1)^2 / 2),$$

$$1 = q_t [1 - \varphi^x (x_t^{gap} - 1)(3x_t^g - 1)/2] + \beta \varphi^x E_t [(\tilde{\lambda}_t / \tilde{\lambda}_{t+1}) q_{t+1} (x_{t+1}^{gap})^2 (x_{t+1}^{gap} - 1)],$$

$$q_t = (\beta / \bar{g}) E_t [(\tilde{\lambda}_t / \tilde{\lambda}_{t+1})(r_{t+1}^k + (1-\delta) q_{t+1})],$$

$$x_t^{gap} = \tilde{x}_t / \tilde{x}_{t-1}.$$

29

The New Keynesian model also includes the following equations:

$$1 = (\beta/\bar{g})E_t[(a_{t+1}/a_t)(\tilde{c}_t/\tilde{c}_{t+1})(s_t i_t/(\bar{\pi}\pi_{t+1}^{gap}))],$$

$$\varphi(\pi_t^{gap} - 1)\pi_t^{gap} = (\mu_t mc_t - 1)/(\mu_t - 1) + \beta\varphi E_t[(a_{t+1}/a_t)(\tilde{c}_t/\tilde{c}_{t+1})(\pi_{t+1}^{gap} - 1)\pi_{t+1}^{gap}(\tilde{y}_{t+1}/\tilde{y}_t)],$$

$$i_t = \ell_{t-1}(\bar{\imath}(\pi_t^{gap})^{\phi_\pi}\tilde{y}_t^{\phi_y})^{1-\rho_i}\exp(mp_t),$$

$$\ell_t = i_t^{\rho_i}\tilde{y}_t^{\phi_y(\rho_i-1)}.$$

C DERIVATION OF APPROXIMATIONS

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

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

so the unconditional expectation is given by

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

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

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

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

Now suppose x follows a stochastic process. An approximation of the time-t conditional variance of y_{t+1}, denoted by $\text{Var}_t[y_{t+1}]$, follows the same deviation, except it is approximated around the conditional mean of x_{t+1}, denoted by $E_t[x_{t+1}]$. Therefore, $\text{Var}_t[y_{t+1}] \approx [g'(E_t[x_{t+1}])]^2\sigma_{x,c}^2$.

Real Business Cycle Model with CES Production Substituting (18) into (17) implies

$$\hat{y}_{t+1} = \hat{y}_0 + \frac{\sigma}{\sigma-1}\ln\left(\alpha\exp(\frac{\sigma-1}{\sigma}(\hat{k}_t - \hat{k}_0)) + (1-\alpha)\exp(\frac{\sigma-1}{1+\eta\sigma}\hat{\kappa}_n + \frac{(\sigma-1)(1+\eta)}{1+\eta\sigma}\hat{z}_{t+1} - \frac{\sigma-1}{\sigma}\hat{n}_0)\right).$$

Denote $\hat{y}_z$ and $\hat{y}_{zz}$ as the first and second partial derivatives with respect to productivity and define

$$\exp(\hat{f}_1) = \exp(\frac{\sigma-1}{\sigma}(\hat{k}_t - \hat{k}_0)), \quad \exp(\hat{f}_2) = \exp(\frac{\sigma-1}{1+\eta\sigma}\hat{\kappa}_n + \frac{(\sigma-1)(1+\eta)}{1+\eta\sigma}\hat{z}_{t+1} - \frac{\sigma-1}{\sigma}\hat{n}_0),$$

Then the first and second derivatives are given by

$$\hat{y}_{\hat{z}} = \frac{\sigma(1+\eta)}{1+\eta\sigma}\frac{(1-\alpha)\exp(\hat{f}_2)}{\alpha\exp(\hat{f}_1)+(1-\alpha)\exp(\hat{f}_2)}, \tag{24}$$

$$\hat{y}_{\hat{z}\hat{z}} = \frac{\sigma(\sigma-1)(1+\eta)^2}{(1+\eta\sigma)^2}\frac{\alpha(1-\alpha)\exp(\hat{f}_1)\exp(\hat{f}_2)}{(\alpha\exp(\hat{f}_1)+(1-\alpha)\exp(\hat{f}_2))^2}, \tag{25}$$

where $\hat{y}_{\hat{z}} > 0$ and $\hat{y}_{\hat{\mu}} < 0$. The sign of the second derivative depends on the $\sigma - 1$ term.

D ANALYTICAL SOLUTIONS

D.1 COBB-DOUGLAS SOLUTION WITH MARKUP SHOCKS

Forward iteration implies

$$\tfrac{c_t+k_t}{c_t} = 1 + \sum_{j=1}^{\infty}(\alpha\beta)^j E_t\left[\exp\left(-\sum_{i=1}^{j}\hat{\mu}_{t+i}\right)\right] \equiv u_t.$$

The price markup in period $t+j$, written as a function of the current markup, is given by

$$\hat{\mu}_{t+1} = \mu_0 + \rho_\mu\hat{\mu}_t + \sigma_\mu\varepsilon_{t+1},$$

$$\hat{\mu}_{t+2} = \mu_0(1+\rho_\mu) + \rho_\mu^2\hat{\mu}_t + \rho_\mu\sigma_\mu\varepsilon_{t+1} + \sigma_\mu\varepsilon_{t+2},$$

$$\hat{\mu}_{t+j} = \mu_0\sum_{i=1}^{j}\rho_\mu^{i-1} + \rho_\mu^j\hat{\mu}_t + \sum_{i=1}^{j}\rho_\mu^{j-i}\sigma_\mu\varepsilon_{t+i}.$$

where $\mu_0 \equiv (1-\rho_\mu)\ln\bar{\mu}$. Therefore, the sum of the markup over j periods is equal to

$$\sum_{i=1}^{1}\hat{\mu}_{t+i} = \mu_0 + \rho_\mu\hat{\mu}_t + \sigma_\mu\varepsilon_{t+1},$$

$$\sum_{i=1}^{2}\hat{\mu}_{t+i} = \mu_0(2+\rho_\mu) + \rho_\mu(1+\rho_\mu)\hat{\mu}_t + \sigma_\mu(1+\rho_\mu)\varepsilon_{t+1} + \sigma_\mu\varepsilon_{t+2},$$

$$\sum_{i=1}^{j}\hat{\mu}_{t+i} = \mu_0\sum_{i=1}^{j}(j+1-i)\rho_\mu^{i-1} + \rho_\mu\sum_{i=1}^{j}\rho_\mu^{i-1}\hat{\mu}_t + \sum_{h=1}^{j}(\sum_{i=0}^{j-h}\rho_\mu^i)\sigma_\mu\varepsilon_{t+h}.$$

Since $\hat{\mu}_{t+1}$ is normally distributed with mean $\mu_0+\rho_\mu\hat{\mu}_t$ and variance σ_μ^2, $\exp(\hat{\mu}_{t+j})$ is log-normally distributed with mean $\exp(\mu_0+\rho_\mu\hat{\mu}_t+\sigma_\mu^2/2)$. Applying this property to $\exp(-\sum_{i=1}^{j}\hat{\mu}_{t+i})$ implies

$$\tfrac{c_t+k_t}{c_t} = 1 + \sum_{j=1}^{\infty}(\alpha\beta)^j\exp\left(-\mu_0\sum_{i=1}^{j}(j+1-i)\rho_\mu^{i-1} - \rho_\mu\sum_{i=1}^{j}\rho_\mu^{i-1}\hat{\mu}_t + \sum_{h=1}^{j}(\sum_{i=0}^{j-h}\rho_\mu^i)^2\sigma_\mu^2/2\right).$$

In the special case there markup shocks are i.i.d., this result simplifies to

$$\tfrac{c_t+k_t}{c_t} = \sum_{j=0}^{\infty}[(\alpha\beta/\bar{\mu})\exp(\sigma_\mu^2/2)]^j = 1/(1-(\alpha\beta/\bar{\mu})\exp(\sigma_\mu^2/2)),$$

so the output-to-consumption ratio is constant and does not cause equilibrium hours to vary.

D.2 CES SOLUTION WITH I.I.D. MARKUP SHOCKS

The Euler equation is given by

$$1 = \alpha\beta(y_0/k_0)^{(\sigma-1)/\sigma}E_t[(c_t/c_{t+1})^\gamma(y_{t+1}/k_t)^{1/\sigma}(1/\mu_{t+1})].$$

To solve the model, assume $\gamma = 1/\sigma$ and guess $c_t = \theta y_t$. Then

$$1 = \alpha\beta(y_0/k_0)^{(\sigma-1)/\sigma}(y_t/k_t)^{1/\sigma}E_t[\exp(-\hat{\mu}_{t+1})].$$

With i.i.d. markup shocks, $E_t[\exp(-\hat{\mu}_{t+1})] = (1/\bar{\mu})\exp(\sigma_\mu^2/2)$. Therefore,

$$c_t = (1-((\alpha\beta/\bar{\mu})\exp(\sigma_\mu^2))^\sigma(y_0/k_0)^{\sigma-1}y_t, \quad k_t = ((\alpha\beta/\bar{\mu})\exp(\sigma_\mu^2))^\sigma(y_0/k_0)^{\sigma-1}y_t,$$

which verifies our initial conjecture that the consumption-to-output ratio is constant.

31

E 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, $\mathbf{s}_t$ is a vector of variables, ε_t are the shocks, $\mathbf{z}_t$ are the states ($\mathbf{z}_t \equiv [k_{t-1}, z_t^n]'$ for the baseline model and $\mathbf{z}_t \equiv [\tilde{k}_{t-1}, \tilde{c}_{t-1}, \tilde{x}_{t-1}, z_t^n, z_t^k]'$ for the estimated model).

There are many ways to discretize the exogenous states, z_t^n and z_t^k. We use the Markov chain in Rouwenhorst (1995), which Kopecky and Suen (2010) show outperforms other methods for approximating autoregressive processes. The bounds on the endogenous states are set as a percentage of their deterministic steady state values (k_{t-1} is set to $\pm 15\%$ for the baseline model and $\tilde{k}_{t-1}, \tilde{c}_{t-1}$, and $\tilde{x}_{t-1}$, are respectively set to $\pm 10\%, \pm 25\%, \pm 10\%$ for the estimated model). These values were chosen so the grids contain 99% of the simulated values for each state. We discretize the states into 9 evenly-spaced points. The product of the points in each dimension, D, represents the total nodes in the state space ($D = 81$ for the baseline model and $D = 59{,}049$ for the estimated model). The realization of $\mathbf{z}_t$ on node d is denoted $\mathbf{z}_t(d)$. The Rouwenhorst method provides integration nodes, $[z_{t+1}^n(m), z_{t+1}^k(m)]$, with weights, $\phi(m)$, for $m \in \{1, \ldots, M\}$. Since the exogenous variables evolve according to a Markov chain, the number of future realizations is the same as the state (9).

The vector of policy functions is denoted $\mathbf{pf}_t$ and the realization on node d is denoted $\mathbf{pf}_t(d)$ ($\mathbf{pf}_t \equiv [n_t(\mathbf{z}_t)]$ for the baseline model and $\mathbf{pf}_t \equiv [n_t(\mathbf{z}_t), q_t(\mathbf{z}_t)]$ for the estimated model). 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 policy function iteration 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

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

   When `csolve` converges, set $\mathbf{pf}_j(d) = \mathbf{pf}_t(d)$.

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

F 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 a sampling distribution for each parameter. 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 real business cycle model. Given a random seed, s, 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 in the data. Denote the shock matrix $\mathcal{E}_{B+T}^s$ (e.g., in the 2-shock model $\mathcal{E}_{B+T}^s = [\varepsilon_{\hat{z}^k,t}^s, \varepsilon_{\hat{z}^n,t}^s]_{t=1}^{B+T}$). For $s \in \{1,\dots,N_s\}$, run the following steps:

   (a) Specify a guess, $\hat{\theta}_0$, for the N_p estimated parameters and the covariance matrix, $\Sigma_P^{s,0}$.

      i. For all $i \in \{1,\dots,N_m\}$, apply the following steps:

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

         B. Solve the nonlinear model with the policy function iteration algorithm in Appendix E given $\hat{\theta}_i$. Repeat 3(a)i.A. if the algorithm does not converge.

         C. Given $\mathcal{E}_{B+T}^s(r)$, 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 the remaining T periods, $\Psi_T^M(\hat{\theta}_i, \mathcal{E}_T^s(r))$, the same way they are calculated in the data.

         D. Calculate the median moments across the R simulations, $\bar{\Psi}_{R,T}^M(\hat{\theta}_i, \mathcal{E}_T^s) = \mathrm{median}\{\Psi_T^M(\hat{\theta}_i, \mathcal{E}_T^s(r))\}_{r=1}^R$, and evaluate the loss function:

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

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

         A. 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 $\tilde{\Theta}^s = \hat{\Theta}^s - \mathbf{1}_{N_{best} \times 1}\sum_{i=N_{best}}^{N_d}\hat{\theta}_i^s/(N_{best})$.

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

   (b) Use simulated annealing to minimize the loss function.

      i. For $i \in \{0,\dots,N_d\}$, repeat the following steps:

A. 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.

B. Solve the nonlinear model with policy function iteration given $\hat{\theta}_i^{cand}$.

C. Given $\mathcal{E}_{B+T}^s(r)$, 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 the remaining T periods, $\Psi_T^M(\hat{\theta}_i^{cand}, \mathcal{E}_T^s(r))$, the same way they are calculated in the data.

D. Calculate the median moments across the R simulations, $\bar{\Psi}_{R,T}^M(\hat{\theta}_i^{cand}, \mathcal{E}_T^s) = \text{median}\{\Psi_T^M(\hat{\theta}_i^{cand}, \mathcal{E}_T^s(r))\}_{r=1}^R$, and evaluate the loss function:

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

E. 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.

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

A. 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 $\tilde{\Theta}^s = \hat{\Theta}^s - \mathbf{1}_{N_d/2 \times 1} \sum_{i=N_d/2}^{N_d} \hat{\theta}_i^s/(N_d/2)$.

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

(c) 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 N_J lowest J values, denoted $\{J_{guess}^{s,j}\}_{j=1}^{N_J}$, and the corresponding draws, $\{\theta_{guess}^{s,j}\}_{j=1}^{N_J}$.

(d) For $j \in \{1, \ldots, N_J\}$, minimize the same loss function with MATLAB's `fminsearch` starting at $\theta_{guess}^{s,j}$. The resulting minimum is $\hat{\theta}_{min}^{s,j}$ with a loss function value of $J_{min}^{s,j}$. Repeat, each time updating the guess, until $J_{guess}^{s,j} - J_{min}^{s,j} < 0.001$. The parameter estimates reported in the tables in the main paper, denoted $\hat{\theta}^s$, correspond to $\min\{J_{min}^{s,j}\}_{j=1}^{N_J}$.

3. The set of SMM parameter estimates $\{\hat{\theta}^s\}_{s=1}^{N_s}$ approximate the joint sampling distribution of the parameters. We report its mean, $\bar{\theta} = \sum_{s=1}^{N_s} \hat{\theta}^s/N_s$, and $(5, 95)$ percentiles.

We set $N_s = 200$, $R = 1{,}000$, $N_{SMM} = 3$, and $N_J = 1$. N_m, N_d, N_p, and c_1 are all model-specific.

G ADDITIONAL RESULTS

Parameter	C-D	CES		
	$\hat{z}^n$	$\hat{z}^n$	$\hat{z}^k$	$\hat{z}^n$ & $\hat{z}^k$
ρ_{zn}	0.9629 (0.9605, 0.9653)	0.9582 (0.9496, 0.9640)	−	0.9671 (0.9573, 0.9770)
σ_{zn}	0.0212 (0.0209, 0.0214)	0.0230 (0.0216, 0.0243)	−	0.0196 (0.0176, 0.0216)
ρ_{zk}	−	−	0.9325 (0.9297, 0.9351)	0.9497 (0.9106, 0.9798)
σ_{zk}	−	−	0.0350 (0.0341, 0.0358)	0.0069 (0.0062, 0.0076)
φ_x	3.0369 (2.9041, 3.1527)	3.4135 (2.8463, 4.0188)	2.5592 (2.3614, 2.7270)	2.0112 (1.3984, 2.6369)
h	0.9001 (0.8965, 0.9034)	0.8176 (0.7913, 0.8392)	0.8912 (0.8875, 0.8945)	0.7502 (0.6700, 0.8125)
σ	−	0.1692 (0.1658, 0.1723)	0.4738 (0.4582, 0.4875)	0.1294 (0.1071, 0.1480)
J	55.89 (55.29, 56.52)	18.19 (16.81, 20.45)	21.51 (20.94, 22.08)	8.06 (7.67, 8.58)
pval	0.00 (0.00, 0.00)	0.01 (0.00, 0.02)	0.00 (0.00, 0.00)	0.15 (0.13, 0.18)
df	8	7	7	5

(a) Average and $(5, 95)$ percentiles of the parameter estimates.

Moment	Data	C-D	CES		
		$\hat{z}^n$	$\hat{z}^n$	$\hat{z}^k$	$\hat{z}^n$ & $\hat{z}^k$
$SD(\mathcal{U}^{\hat{y}}_{t,t+1})/SD(\hat{y}^g_t)$	5.75	1.96 (−5.99)	5.97 (0.36)	5.47 (−0.44)	5.70 (−0.07)
$Corr(\mathcal{U}^{\hat{y}}_{t,t+1}, \hat{y}^g_t)$	−0.41	−0.33 (0.84)	−0.64 (−2.63)	−0.34 (0.67)	−0.48 (−0.83)
$AC(\mathcal{U}^{\hat{y}}_{t,t+1})$	0.92	0.90 (−0.45)	0.88 (−1.15)	0.89 (−0.67)	0.88 (−0.99)
$SD(\hat{y}^g)$	0.80	0.95 (2.19)	0.84 (0.59)	0.89 (1.35)	0.80 (0.00)
$SD(\hat{c}^g)$	0.51	0.41 (−2.01)	0.47 (−0.78)	0.50 (−0.07)	0.54 (0.61)
$SD(\hat{x}^g)$	2.04	2.09 (0.26)	1.81 (−1.12)	1.74 (−1.48)	1.89 (−0.74)
$AC(\hat{y}^g)$	0.29	0.47 (1.98)	0.33 (0.40)	0.39 (1.12)	0.24 (−0.58)
$AC(\hat{c}^g)$	0.39	0.47 (1.15)	0.26 (−1.87)	0.49 (1.41)	0.38 (−0.21)
$AC(\hat{x}^g)$	0.42	0.48 (0.64)	0.49 (0.73)	0.36 (−0.61)	0.45 (0.31)
$AC4(\hat{y}^g)$	0.11	0.03 (−1.27)	0.01 (−1.54)	−0.02 (−1.97)	0.02 (−1.35)
$AC4(\hat{c}^g)$	0.06	0.07 (0.09)	0.05 (−0.24)	0.11 (0.62)	0.15 (1.28)
$AC4(\hat{x}^g)$	0.15	0.03 (−1.70)	0.09 (−0.89)	−0.05 (−2.92)	0.07 (−1.12)

(b) Data and average model-implied moments. t-statistics are in parentheses.

Table 6: Parameter estimates and targeted moments.