JSPS Grants-in-Aid for Scientific Research (S)

Understanding Persistent Deflation in Japan

Working Paper Series

No. 032

December 2013

The Uncertainty Multiplier and Business Cycles

Hikaru Saijo

UTokyo Price Project
702 Faculty of Economics, The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
Tel: +81-3-5841-5595
E-mail: watlab@e.u-tokyo.ac.jp
http://www.price.e.u-tokyo.ac.jp/english/

The Uncertainty Multiplier and Business Cycles

Hikaru Saijo*

University of California, Santa Cruz

November 1, 2013

Abstract

I study a business cycle model where agents learn about the state of the economy by accumulating capital. During recessions, agents invest less, and this generates noisier estimates of macroeconomic conditions and an increase in uncertainty. The endogenous increase in aggregate uncertainty further reduces economic activity, which in turn leads to more uncertainty, and so on. Thus, through changes in uncertainty, learning gives rise to a multiplier effect that amplifies business cycles. I use the calibrated model to measure the size of this uncertainty multiplier.

*This work is based on the first chapter of my Ph.D. dissertation at the University of Pennsylvania. I am grateful to my advisor, Jesús Fernández-Villaverde, for his guidance and encouragement at every stage of this project. I also thank my dissertation committee, Dirk Krueger and Frank Schorfheide, and seminar participants at numerous places whose comments led to substantial improvement of the paper. *E-mail*: hsaijo@ucsc.edu.

1 Introduction

What drives business cycles? A rapidly growing literature argues that shocks to uncertainty are a significant source of business cycle dynamics—see, for example, Bloom (2009), Fernández-Villaverde et al. (2011), Gourio (2012), and Christiano et al. (forthcoming). However, the literature faces at least two important criticisms. In uncertainty shock theories, recessions are caused by exogenous increases in the volatility of structural shocks. First, fluctuations in uncertainty may be, at least partially, endogenous.[1] The distinction is crucial because if uncertainty is an equilibrium object that is coming from agents' actions, policy experiments that treat uncertainty as exogenous are subject to the Lucas critique. Second, some authors (Bachmann and Bayer 2013, Born and Pfeifer 2012, and Chugh 2012) have argued that, given small and transient fluctuations in observed ex-post volatility, changes in uncertainty have negligible effects. However, time-varying volatility need not be the only source of time-varying uncertainty.[2] If this is the case, these papers may be understating the contribution of changes in uncertainty to aggregate fluctuations.

In this paper, I present a business cycle model where the level of economic activity influences the level of aggregate uncertainty. The endogenous movement in uncertainty, in turn, affects the level of economic activity. My goal is to quantify the role of this two-way feedback between uncertainty and economic activity in explaining the fluctuations of macroeconomic variables.

I embed the idea of asymmetric learning (Veldkamp 2005 and Van Nieuwerburgh and Veldkamp 2006) into a standard DSGE framework with several real and nominal rigidities (Christiano et al. 2005). I introduce information frictions by subjecting the economy to aggregate shocks that agents cannot directly observe, namely, shocks to the marginal efficiency of investment and shocks to the depreciation rate of capital. Because the former are persistent while the latter are transitory, what matters for agents' optimal decision is the evolution of the efficiency of investment. Agents use the path of capital stock and investment to form their estimates in a Bayesian manner.[3] However, the capital stock is not perfectly revealing about the unobservable shocks because it is subject to a non-invertibility problem: Agents cannot tell whether an unexpectedly high realization of capital stock is due to a high efficiency of investment or to a low depreciation rate of capital.

In the model, the level of investment endogenously determines the informativeness of the capital stock about the shocks to the efficiency of investment. When agents invest less, their estimates are imprecise because the level of capital stock is largely determined by the realization of the depreciation shock. Conversely, when they invest more, their estimates are accurate because the current capital mostly reflects shocks to the efficiency of investment. Thus, aggregate uncertainty

[1] See Bachmann et al. (2013) for supporting VAR evidence.

[2] Orlik and Veldkamp (2013) show that changes in uncertainty measured using survey data are stronger than changes in GDP volatility.

[3] In the model, all information necessary for optimal learning is contained in the path of capital stock and investment. While agents have access to other endogenous variables, including prices, those variables do not reveal additional information about the unobservable shocks.

becomes endogenously countercyclical over the business cycle.

The countercyclical uncertainty gives rise to a novel multiplier effect that amplifies business cycles. Imagine that the economy is hit by a negative shock that lowers investment (for example, an exogenous tightening of monetary policy). Since agents learn less about the current period shock to the efficiency of investment, uncertainty increases. This, in turn, further reduces investment and other economic activity because of households' precautionary motive and countercyclical movements in markups. The opposite channel works when the economy is hit by a positive shock. I call this amplification mechanism the *uncertainty multiplier*.

To measure the size of the uncertainty multiplier, I perform numerical simulations. The model is calibrated to match the business cycle properties of the postwar U.S. quarterly data. An interesting challenge I face is that the choice of the variance parameters has important effects on the strength of learning dynamics. More specifically, when the variance of the depreciation shock is very small compared to that of the shock to the efficiency of investment, the capital stock is almost perfectly revealing about the shock to the efficiency of investment. Conversely, when the depreciation shock is very large, the capital stock is uninformative and little learning takes place. In both cases, fluctuations in aggregate uncertainty are negligible. To ensure that agents face a realistic amount of information frictions, I pin down the variance parameters so that the model replicates the properties of survey data on macroeconomic forecasts.

I find that, under the benchmark calibration with a full set of real and nominal rigidities, the standard deviation of output is amplified by 18%. Other real variables, such as investment and hours, are also amplified by a similar amount. The size of the amplification is nontrivial due to two main features of the model. First, in my model, changes in uncertainty generate positive comovements among real variables. Second, the uncertainty process is volatile and persistent because it is tied to the movement of investment.

Finally, I provide an external validation of my theory by showing that it replicates the VAR impulse response of the survey measure of uncertainty. In particular, it can account for the negative relationship between output and uncertainty and it also reproduces gradual responses of the two variables. This is because in the model uncertainty is inversely related to investment, which exhibits hump-shaped dynamics, and this uncertainty in turn induces gradual adjustments by households.

The rest of the paper is organized as follows. The next section describes my contributions with respect to the existing literature. Section 3 presents the model and Section 4 discusses its solution and calibration. Section 5 presents the results. Section 6 provides evidence of my theory from survey data. Section 7 concludes with some directions for future research.

2 Connections to the Literature

This paper is related to several strands of the literature. First, it is related to a growing literature on uncertainty shocks. A leading example is a paper by Bloom (2009), who shows that an exogenous increase in the volatility of firm-level productivity reduces output through a "wait-and-see" effect due to investment irreversibility. Fernández-Villaverde et al. (2011) show that volatility shocks to real interest rates generate sizable contractions in an otherwise standard small open economy model. Other examples include Arellano et al. (2012), Basu and Bundick (2011), Christiano et al. (forthcoming), Fernández-Villaverde et al. (2012), Gilchrist et al. (2010), Gourio (2012), Ilut and Schneider (2011), and Schaal (2012). I show that time-varying uncertainty could be an important amplification (rather than an impulse) mechanism of the business cycle. As stated in the Introduction, this distinction is important because now uncertainty is an equilibrium object.

Recently, some authors have argued that changes in uncertainty have negligible effects given small and transient fluctuations in observed realized volatility (Bachmann and Bayer 2013, Born and Pfeifer 2012, and Chugh 2012). In contrast, in my model, uncertainty features a large and persistent fluctuation that is not linked with movements in the observed volatility of macro variables.[4] As a result, unlike in these papers, changes in uncertainty have sizable effects.

Several papers attempt to account for the countercyclical firm-level volatility through conventional first-moment shocks. For example, in Bachmann and Moscarini (2011), recessions induce firms to price-experiment, which in turn raises the cross-sectional dispersion of price changes. See also D'Erasmo and Boedo (2012), Kehrig (2011), and Tian (2012). Their focus is on the explanation of the movement of ex-post volatility. I go one step further by highlighting the implications of ex-ante uncertainty. This is why I can show that uncertainty is not only a by-product of agents' response to first-moment shocks, but also an important factor that affects real allocations.

The main mechanism of this paper builds on a literature on asymmetric learning, for example, Veldkamp (2005), Van Nieuwerburgh and Veldkamp (2006), Ordoñez (2012), and Görtz and Tsoukalas (2013). They argue that the time-varying speed of learning about the macroeconomic conditions could explain the asymmetries in growth rates over the business cycle. When the economy passes the peak of a boom, agents are able to precisely detect the slowdown, leading to an abrupt crash. At the end of the recession, agents' estimates about the extent of recovery are noisy, slowing reactions and delaying booms. My contribution is to explore the direct effects of endogenous fluctuations in uncertainty that shift the levels of macro variables. Recessions are deeper because high uncertainty leads precautionary households to cut consumption. Booms are stronger for the opposite reason. This channel has been overlooked in the previous literature.

A recent work by Fajgelbaum et al. (2013) independently develops similar ideas. There are

[4] Ilut and Schneider (2011) also propose a business cycle model where changes in uncertainty are not followed by changes in volatility by assuming ambiguity-averse preferences.

two key distinctions. First, in their paper the level of aggregate uncertainty is related to the number of firms investing (extensive margin), while in my paper the level of investment influences the level of uncertainty (intensive margin). As in Van Nieuwerburgh and Veldkamp (2006), this specification allows me to perform realistic quantitative analysis without losing tractability. Second, in their paper time-varying uncertainty feeds back into the level of economic activity through irreversible investment, while in my paper uncertainty influences business cycles through countercyclical markups due to nominal rigidities. The advantage of my approach is that the markup channel generates comovements among real variables that are consistent with U.S. business cycles (Basu and Bundick 2011).

Finally, this paper joins a long tradition in macroeconomics, starting from Lucas (1972), by considering the role of imperfect information and expectations in shaping business cycle dynamics. Recent contributions include Barsky and Sims (2012), Beaudry and Portier (2004), Eusepi and Preston (2011), Lorenzoni (2009), Jaimovich and Rebelo (2009), and Schmitt-Grohe and Uribe (2012). These papers emphasize changes in the *mean* of agents' subjective estimates about fundamentals. The current paper, instead, demonstrates the importance of changes in the *variance* of estimates about fundamentals. As a methodological contribution, I show how to apply higher-order approximation methods to a model with linear information frictions.

3 The Model

I embed a learning problem into the capital accumulation process of a standard DSGE framework (Christiano et al. 2005, Justiniano et al. 2010, and Smets and Wouters 2007). This framework is a natural laboratory for my quantitative investigation, since it has now become the foundation of applied research in both academic and government institutions.

In the first subsection, I describe the information frictions. In the second subsection, I present the standard part of the model.

3.1 Learning and Endogenously Countercyclical Uncertainty

I divide this subsection into several parts. First, I describe the setup. Second, I provide concrete examples that help interpret the baseline environment. Third, I express the learning process as a Kalman filtering problem. Fourth, I present a simple example that illustrates the key properties of the filtering problem. Finally, I rewrite the capital accumulation process from the perspective of the agents. This clarifies the impact of changes in uncertainty on the agents' decision making.

3.1.1 Setup

The law of motion for capital, K_t, is subject to two types of aggregate disturbances:

$$K_t = (1 - \delta_t)K_{t-1} + \mu_t I_{t-1}.$$

The depreciation shock, δ_t, follows

$$\delta_t = \delta - \epsilon_{\delta,t},$$

where $\epsilon_{\delta,t}$ is i.i.d. distributed from a normal distribution with mean zero and variance σ_δ^2. The investment shock, μ_t, determines the marginal efficiency of investment. I assume that μ_t follows the stochastic process

$$\begin{aligned} \mu_t &= g_{t-1} + (1 - \rho_\mu)\mu + \rho_\mu \mu_{t-1} + \epsilon_{\mu,t}, \\ g_t &= \rho_g g_{t-1} + \epsilon_{g,t}, \end{aligned}$$

where $\epsilon_{\mu,t}$ and $\epsilon_{g,t}$ are i.i.d. distributed from a normal distribution with mean zero and variance σ_μ^2 and σ_g^2, respectively. The growth shock, g_t, controls the growth rate of μ_t.[5] Agents cannot directly observe the current or previous values of δ_t, μ_t, and g_t. This informational assumption gives rise to a non-invertibility problem: Agents cannot tell whether an unexpectedly high realization of capital stock is due to a high efficiency of investment or to a low depreciation rate of capital. As a result, they face a signal-extraction problem in forecasting the evolution of the shocks. Agents use all available information, including the path of capital stock, to form their estimates.

A literal interpretation of the depreciation shock is that it represents an exogenous change in the physical depreciation rate of capital. However, as in Gourio (2012), Gertler and Karadi (2011), and Liu et al. (2011), a broader interpretation is possible. For example, it can represent an economic obsolescence of capital. Alternatively, reallocation of capital may be subject to temporary frictions and could show up as a change in the "quality" of aggregate capital.

The investment shock was originally proposed by Greenwood et al. (1988). In a medium-scale DSGE model similar to the one employed in this paper, Justiniano et al. (2010) have found that the shock is the most important driver of the U.S. business cycle. In general, there are two ways to think about the investment shocks. The first interpretation is that they represent disturbances that affect the transformation of consumption goods into investment goods. The second interpretation is that they are shocks that affect the transformation of investment goods into installed capital. In this paper, I adopt the second interpretation.[6] An important implication of this interpretation is that, unlike Fisher (2006), the investment shock does not affect the price of investment goods relative to consumption goods. Thus, agents cannot back out the shocks by observing the price.

[5] The growth shock is not necessary for the main qualitative results. However, as I show below the shock helps match some of the survey data moments.

[6] See Justiniano et al. (2011) for supportive evidence based on a DSGE model estimation.

As summarized in Figure 1, the timing of events is as follows: At the end of period $t-1$, agents choose their investment level I_{t-1} given the current capital level K_{t-1} and their estimates about the unobservable state. Then, at the beginning of period t, unobservable shocks are realized. Finally, after observing the level of new capital K_t, agents update the estimates.[7]

3.1.2 More on the Setup: Concrete Examples

The information friction in the model is set up in a fairly stylistic way. However, the same qualitative feature arises in a more concrete environment. Below, I provide two examples.

First, one can consider a setting where the aggregate capital stock is composed of distinct vintages that are hit by stochastic depreciations.[8] The process of μ_t is persistent, reflecting the view that the productivity of "close-by" vintages are similar.[9] Let K_t denote the sum of all efficiency units of capital available for production in period t:

$$K_t = \sum_{s=0}^{S} K_{t,s}$$

where $K_{t,s}$ is capital of vintage s that is available at time t. The vintages of capital evolve according to

$$K_{t,s} = \begin{cases} \mu_t I_{t-1} & \text{if } s = 0 \\ (1-\delta)K_{t-1,s-1} + \epsilon^{\delta}_{t,s} & \text{if } s \geq 1 \end{cases}$$

where $\epsilon^{\delta}_{t,s}$ is a stochastic capital depreciation to vintage s at time t. $\epsilon^{\delta}_{t,s}$, $s = 1, \ldots, S$ are distributed from a multivariate normal distribution with mean zero and covariance matrix Σ_δ. While I assume that draws are independent across time, I still allow for contemporaneous correlations among $\epsilon^{\delta}_{t,s}$.[10] The evolution of aggregate capital is then given by

$$K_t = (1-\delta)K_{t-1} + \tilde{\epsilon}^{\delta}_t + \mu_t I_{t-1}$$

where $\tilde{\epsilon}^{\delta}_t \equiv \sum_{s=1}^{S} \epsilon^{\delta}_{t,s}$ are distributed independently across time from a normal distribution with mean zero and some variance $\tilde{\sigma}_\delta$.[11] Thus, in this example capital depreciation is transitory while the marginal efficiency of investment is persistent.

Second, as will be discussed in detail below, an important feature of the capital accumulation

[7] In the Appendix, I discuss how to implement the investment problem in a decentralized competitive equilibrium. Importantly, relative prices like the price of capital do not reveal the unobservable states and hence the information frictions survive.

[8] Görtz and Tsoukalas (2013) also consider a model with the vintage view of capital stock.

[9] For example, the functions of iPhone 5 would be more similar to iPhone 4 than the 5 to the original iPhone.

[10] Continuing with the iPhone example, it is natural to think that iPhone 4 and iPhone 5 becomes obsolete by a similar amount when iPhone 6 is introduced.

[11] The depreciation shock is additive in this example while it is attached to capital stock in the baseline model. The difference matters little because capital stock is acyclical.

technology is that it generates a procyclical signal-to-noise ratio. In the Appendix, I show how this feature also arises from aggregation of investment units with common and idiosyncratic shocks.

3.1.3 The Kalman Filtering Problem

Agents update their estimates about μ_t and g_t in an optimal (Bayesian) manner. The learning process can be expressed as a Kalman filtering problem:

$$\begin{bmatrix} \mu_t \\ g_t \end{bmatrix} = \begin{bmatrix} (1-\rho_\mu)\mu \\ 0 \end{bmatrix} + \begin{bmatrix} \rho_\mu & 1 \\ 0 & \rho_g \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ g_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_{\mu,t} \\ \epsilon_{g,t} \end{bmatrix}, \tag{1}$$

$$K_t - (1-\delta)K_{t-1} = [I_{t-1} \quad 0] \begin{bmatrix} \mu_t \\ g_t \end{bmatrix} + K_{t-1}\epsilon_{\delta,t}. \tag{2}$$

Equation (1) is the state equation that characterizes the evolution of the unobservable state. Equation (2) is the measurement equation that describes the observables as a linear function of the underlying state. I point out two things regarding the measurement equation. First, $\epsilon_{\delta,t}$ serves as a measurement error in the filtering system. Second, unlike standard time-invariant systems, the coefficient matrices are time-varying.[12]

The key property of the system is that the signal-to-noise ratio is procyclical, which follows from the fact that $\frac{I_{t-1}}{K_{t-1}}$ is procyclical. The flip side implication of this property is that *uncertainty is countercyclical*. Denote Σ_t as the error-covariance matrix of the unobservable states,

$$\Sigma_t = \begin{bmatrix} \mathrm{Var}_t(\mu_t - \tilde{\mu}_t) & \mathrm{Cov}_t(\mu_t - \tilde{\mu}_t, g_t - \tilde{g}_t) \\ \cdots & \mathrm{Var}_t(g_t - \tilde{g}_t) \end{bmatrix},$$

then the elements of Σ_t are decreasing in $\frac{I_{t-1}}{K_{t-1}}$. Intuitively, when agents invest less, their estimates about the efficiency of investment are imprecise because the level of capital stock is largely determined by the realization of the depreciation shock. Conversely, their estimates are accurate when they invest more because the current capital mostly reflects shocks to the efficiency of investment.

3.1.4 Understanding Why Uncertainty Is Countercyclical

I explain how a procyclical signal-to-noise ratio leads to countercyclical uncertainty by going through a simple example. In particular, assume that there is no growth shock.[13] Then the

[12] As in Veldkamp (2005) and Van Nieuwerburgh and Veldkamp (2006), I rule out active experimentation for computational reasons. Cogley et al. (2007) have shown, in the context of U.S. monetary policy making, that the two approaches (learning with and without experimentation) produce very similar decision rules.

[13] In the Appendix, I provide a full derivation with the growth shock.

filtering problem reduces to

$$\mu_t = (1-\rho_\mu)\mu + \rho_\mu \mu_{t-1} + \epsilon_{\mu,t}, \tag{3}$$

$$y_t = I_{t-1}\mu_t + K_{t-1}\epsilon_{\delta,t}, \tag{4}$$

where (3) is the state equation and (4) is the measurement equation. I define $y_t \equiv K_t - (1-\delta)K_{t-1}$. In period $t-1$, agents enter with the mean estimate $\tilde{\mu}_{t-1}$ and its associated error variance $\Sigma_{t-1} \equiv \mathrm{Var}_{t-1}(\mu_{t-1} - \tilde{\mu}_{t-1})$. Then, the period $t-1$ prediction of μ_t and its associated error variance is given by

$$\tilde{\mu}_{t|t-1} = (1-\rho_\mu)\mu + \rho_\mu \tilde{\mu}_{t-1}$$
$$\Sigma_{t|t-1} = \rho_\mu^2 \Sigma_{t-1} + \sigma_\mu^2$$

After observing the outcome y_t, they update their estimates according to

$$\tilde{\mu}_t = \tilde{\mu}_{t|t-1} + Gain_t(y_t - I_{t-1}\tilde{\mu}_{t|t-1}),$$

where $Gain_t$ is the Kalman gain and is given by

$$Gain_t = \underbrace{\frac{I_{t-1}^2 \Sigma_{t|t-1}}{I_{t-1}^2 \Sigma_{t|t-1} + K_{t-1}^2 \sigma_\delta^2}}_{\text{Informativeness of observation}} \cdot \underbrace{\frac{1}{I_{t-1}}}_{\text{Adjustment}} .$$

The first term represents the informativeness of observation y_t and is given by the variance of the signal divided by the total variance (the variance of the signal and noise). The term is increasing in $\frac{I_{t-1}}{K_{t-1}}$. The second term is the scale adjustment term reflecting the fact that μ_t is multiplied by I_{t-1} in the observation.

The error variance associated with $\tilde{\mu}_t$ is given by

$$\Sigma_t = (1 - Gain_t I_{t-1})\Sigma_{t|t-1}$$
$$= \underbrace{\frac{K_{t-1}^2 \sigma_\delta^2}{I_{t-1}^2 \Sigma_{t|t-1} + K_{t-1}^2 \sigma_\delta^2}}_{\text{Un-informativeness of observation}} \cdot \Sigma_{t|t-1}.$$

The first line says that the error shrinks as we learn more from the observation; the error is decreasing in the size of the Kalman gain. The second line says that the error variance is increasing in the un-informativeness of the observation (the variance of noise divided by the total variance). Since the un-informativeness term is decreasing in $\frac{I_{t-1}}{K_{t-1}}$, Σ_t is decreasing in $\frac{I_{t-1}}{K_{t-1}}$. Since investment is much more volatile than capital, $\frac{I_{t-1}}{K_{t-1}}$ moves almost proportionally to I_{t-1}. Thus, less investment leads to more uncertainty.

3.1.5 Implications of Time-Varying Uncertainty From the Perspective of the Agents

How do changes in uncertainty about the current efficiency of investment affect agents' decision making? The key insight here is that, because shocks to the efficiency of investment are persistent, uncertainty about the current state translates into uncertainty about the future realization of capital.

To see this, it is useful to rewrite the capital accumulation equation from the perspective of the agent at period $t-1$:

$$K_t = (1-\delta_t)K_{t-1} + (\tilde{\mu}_{t|t-1} + u_t)I_{t-1},$$

where $\tilde{\mu}_{t|t-1}$ is the mean forecast of μ_t at time $t-1$ and u_t is normally distributed with mean zero and variance $\sigma^2_{u,t}$. The innovation u_t takes into account not only the exogenous innovation to μ_t, but also its estimation error:

$$\begin{aligned} u_t &= \mu_t - \tilde{\mu}_{t|t-1} \\ &= (g_{t-1} - \tilde{g}_{t-1}) + \rho_\mu(\mu_{t-1} - \tilde{\mu}_{t-1}) + \epsilon_{\mu,t}, \end{aligned}$$

and hence its variance is given by

$$\sigma^2_{u,t} = \rho^2_\mu \Sigma^{11}_{t-1} + 2\rho_\mu \Sigma^{12}_{t-1} + \Sigma^{22}_{t-1} + \sigma^2_\mu.$$

Thus, the fluctuation in uncertainty shows up as a fluctuation in the variance of the innovation to the mean forecast of the marginal efficiency of investment. Moreover, this fluctuation in variance is persistent to the extent that investment is persistent.

3.2 Standard Part of the Model

I now describe other components of the model. The economy is composed of a final-goods sector, intermediate-goods sector, household sector, employment sector, and a central bank. I start by describing the production side of the economy.

3.2.1 The Final-Goods Sector

In each period t, the final goods, Y_t, are produced by a perfectly competitive representative firm that combines a continuum of intermediate goods, indexed by $j \in [0,1]$, with technology

$$Y_t = \left[\int_0^1 Y_{j,t}^{\frac{\theta_p - 1}{\theta_p}} dj\right]^{\frac{\theta_p}{\theta_p - 1}}.$$

$Y_{j,t}$ denotes the time t input of intermediate good j and θ_p controls the price elasticity of demand for each intermediate good. The demand function for good j is

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t}\right)^{-\theta_p} Y_t,$$

where P_t and $P_{j,t}$ denote the price of the final good and intermediate good j, respectively. Finally, P_t is related to $P_{j,t}$ via the relationship

$$P_t = \left[\int_0^1 P_{j,t}^{1-\theta_p} dj\right]^{\frac{1}{1-\theta_p}}.$$

3.2.2 The Intermediate-Goods Sector

The intermediate-goods sector is monopolistically competitive. In period t, each firm j rents $K_{j,t}$ units of capital stock from the household sector and buys $H_{j,t}$ units of aggregate labor input from the employment sector to produce intermediate good j using technology

$$Y_{j,t} = z_t K_{j,t}^{\alpha} H_{j,t}^{1-\alpha}.$$

z_t is the level of total factor productivity that follows

$$z_t = (1-\rho_z)z + \rho_z z_{t-1} + \epsilon_{z,t},$$

where $\epsilon_{z,t}$ is i.i.d. distributed from a normal distribution with mean zero and variance σ_z^2.[14]

Firms face a Calvo-type price-setting friction: In each period t, a firm can reoptimize its intermediate-goods price with probability $(1-\xi_p)$. Firms that cannot reoptimize index their price according to the steady-state inflation rate, π.

3.2.3 The Household Sector

There is a continuum of households, indexed by $i \in [0,1]$. In each period, household i chooses consumption C_t, investment I_t, bond purchases B_t, and nominal wage $W_{i,t}$ to maximize utility:

$$\mathrm{E}_t \sum_{s=0}^{\infty} \beta^s d_{t+s} \left[\frac{(C_{t+s} - bC_{t+s-1})^{1-\sigma}}{1-\sigma} - \frac{H_{i,t+s}^{1+\eta}}{1+\eta}\right],$$

where β is a discount factor, b represents consumption habit, σ controls the degree of risk aversion, η controls (the inverse of) the Frisch labor supply elasticity, and $H_{i,t}$ is the number of hours worked.

[14] I specify the productivity process in levels rather than in logs so that it is consistent with the process of the marginal efficiency of investment. During the simulation exercise, the level of productivity never falls below zero. The same remark applies to the preference shock introduced below.

d_t is a preference shock that follows

$$d_t = (1 - \rho_d)d + \rho_d d_{t-1} + \epsilon_{d,t},$$

where $\epsilon_{d,t}$ is i.i.d. distributed from a normal distribution with mean zero and variance σ_d^2.

The household's budget constraint is

$$P_t C_t + P_t I_t + B_t \leq W_{i,t} H_{i,t} + R_t^k K_t + R_{t-1} B_{t-1} + D_t + A_{i,t},$$

where R_t^k is the rental rate of capital, K_t is the stock of capital,[15] R_{t-1} is the gross nominal interest rate from period $t-1$ to t, and D_t is the combined profit of all the intermediate-goods firms distributed equally to each household. I assume that households buy securities, whose payoffs are contingent on whether they can reoptimize their wage.[16] $A_{i,t}$ denotes the net cash inflow from participating in state-contingent security markets at time t.

As in Christiano et al. (2005), I add an investment adjustment cost to the capital accumulation equation described above:

$$K_t = (1 - \delta_t) K_{t-1} + \mu_t \left(1 - S\left(\frac{I_{t-1}}{I_{t-2}}\right)\right) I_{t-1}, \tag{5}$$

where

$$S\left(\frac{I_{t-1}}{I_{t-2}}\right) = \frac{\kappa}{2}\left(\frac{I_{t-1}}{I_{t-2}} - 1\right)^2,$$

with $\kappa > 0$. Other components of the capital accumulation, like the stochastic process of shocks or the informational structure, are exactly the same as described in the previous section.

3.2.4 The Employment Sector and Wage Setting

In each period t, a perfectly competitive representative employment agency hires labor from households to produce an aggregate labor service, H_t, using technology

$$H_t = \left[\int_0^1 H_{i,t}^{\frac{\theta_w - 1}{\theta_w}} di\right]^{\frac{\theta_w}{\theta_w - 1}},$$

where $H_{i,t}$ denotes the time t input of labor service from household i and θ_w controls the price elasticity of demand for each household's labor service. The agency sells the aggregated labor input to the intermediate firms for a nominal price of W_t per unit. The demand function for the

[15] Note that K_t is the capital stock after the period t shock to the capital accumulation equation is realized.

[16] The existence of state-contingent securities ensures that households are homogeneous with respect to consumption and asset holdings, even though they are heterogeneous with respect to the wage rate and hours because of the idiosyncratic nature of the timing of wage reoptimization. See Christiano et al. (2005).

labor service of household i is

$$H_{i,t} = \left(\frac{W_{i,t}}{W_t}\right)^{-\theta_w} H_t,$$

where $W_{i,t}$ denotes the nominal wage rate of the labor service of household i. W_t is related to $W_{i,t}$ via the relationship

$$W_t = \left[\int_0^1 W_{i,t}^{1-\theta_w} di\right]^{\frac{1}{1-\theta_w}}.$$

Households face a Calvo-type wage-setting friction: In each period t, a household can reoptimize its nominal wage with probability $(1 - \xi_w)$. Households that cannot reoptimize index their wage according to the steady-state inflation rate, π.

3.2.5 The Central Bank, Resource Constraint, and Equilibrium

The central bank follows a Taylor rule with interest-rate smoothing:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho_R} \left\{ \left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{\phi_Y} \right\}^{1-\rho_R} \exp(\epsilon_{R,t}),$$

where R is the steady-state level of the nominal interest rate, ρ_R is the persistence of the rule, and ϕ_π and ϕ_Y are the size of the policy response to the deviation of inflation and output growth from their steady states, respectively. $\epsilon_{R,t}$ is a monetary policy shock and is i.i.d. distributed from a normal distribution with mean zero and variance σ_R^2.

Finally, the aggregate resource constraint is $C_t + I_t = Y_t$. I employ a standard sequential market equilibrium concept and hence its formal definition is omitted.

4 Model Solution and Calibration

I follow Fernández-Villaverde et al. (2011) and solve the model using a third-order perturbation method around its deterministic steady state.[17] I use perturbation because the model has many state variables and it is the only method that delivers an accurate solution in a reasonable amount of time (Aruoba et al. 2006). The third-order approximation is necessary because my purpose is to analyze the direct impact of endogenous changes in aggregate uncertainty. In a standard first-order approximation, changes in uncertainty play no role since the decision rules of agents are forced to follow a certainty equivalence principle. In the second-order approximation, changes in uncertainty only appear in the decision rules as cross-product terms with other state variables. Only in the third-order approximation do changes in uncertainty show up as an independent term.

The parameterization of the model is done in two steps. In the first step, I fix several parameter values following micro evidence or estimates found in other papers. In the second step, I choose

[17] The computation is carried out with Dynare (http://www.dynare.org/).

values of the remaining parameters by matching the simulated moments of the model to the data. The first step reduces the number of parameters to be calibrated and thus sharpens the exercise in the second step.

The discount factor, β, is set so that the model steady-state interest rate implied by the Euler equation matches that of the data. The capital share is set to 0.3. $\delta = 0.02$ implies an annual depreciation rate of 8%. The elasticity of goods demand $\theta_p = 21$ and labor demand $\theta_w = 21$ are consistent with previous estimates, for example, Altig et al. (2011).

I set $\sigma = 2$ and the habit persistence parameter is set to $b = 0.75$. The latter value is in line with the estimates found in Smets and Wouters (2007) and Justiniano et al. (2010). As emphasized in Boldrin et al. (2001), a strong habit persistence parameter helps to account for various asset pricing puzzles. Chetty et al. (2011) suggest a Frisch elasticity of labor supply of 0.5 for a macro model that does not distinguish between intensive and extensive margins. This leads to $\eta = 2$.

The Calvo price and wage parameters imply an average duration of one year. As found in Smets and Wouters (2007) and Justiniano et al. (2010), prices and wages need to be sufficiently sticky in order to account for the inflation and wage dynamics in the data. Turning to the monetary policy parameter, I match the steady-state inflation rate to its historical mean. The Taylor rule coefficients feature inertia with a strong response to inflation and a weak response to output growth (Levin et al. 2006, Smets and Wouters 2007, and Justiniano et al. 2010). The persistence coefficients for the technology shock and preference shock are set to $\rho_z = 0.95$ (Cooley and Prescott 1995) and $\rho_d = 0.22$ (Smets and Wouters 2007), respectively.

To determine the values of other parameters, I choose them so that the moments simulated from the model matches the selected moments in the data.[18] There are 9 parameters to calibrate: $\{\kappa, \rho_\mu, \rho_g, \sigma_z, \sigma_d, \sigma_R, \sigma_\mu, \sigma_g, \sigma_\delta\}$. I target the following 9 data moments:

- <u>Macroeconomic variables</u>:
  Standard deviations of output, investment, and consumption.
  Correlations of investment with respect to output.
  Autocorrelation of output, investment, and consumption.

- <u>Forecast errors from the Survey of Professional Forecasters</u>:
  1st-order autocorrelation and mean size of forecast errors on nominal GDP growth.

Table 1 summarizes the resulting parameter values.

The calibration of the standard deviation of the depreciation shock σ_δ needs further discussion. The parameter is important because it determines the strength of information frictions. When

[18] To simulate the model, I use the pruning procedure as described in Kim et al. (2008) and Den Haan and De Wind (2012). I compute a total of 200 replications of 272 period simulations. I throw away the initial 100 periods, which leaves me with the sample size of the US data (172 periods). For each sample I compute the business cycle moments and then take medians across 200 replications. I checked that the results are not driven by explosive behavior.

σ_δ is very small, the learning problem becomes trivial. When σ_δ is very large, agents learn little about the aggregate state. Thus in both cases, changes in the level of investment have a negligible effect on the level of uncertainty. I discipline the choice of σ_δ by using statistics on forecast errors in the Survey of Professional Forecasters data.[19] The first row in Table 2 reports statistical properties of the one-quarter-ahead median forecast errors on nominal GDP growth rate.[20] The second column shows that the forecast errors of GDP growth are positively autocorrelated. The third column shows the mean size of forecast errors (i.e., forecast precision). I also report the model predictions of the forecast errors for various values of σ_δ.[21] First note that for all values of σ_δ reported, the autocorrelations are positive. This is due to the relatively high persistence parameter of the investment growth shock, ρ_g. The forecast errors are autocorrelated because agents only gradually realize the change in growth rate in response to an innovation to g_t. As σ_δ increases, the autocorrelation decreases because of the additional noise in the filtering problem. On the other hand, the size of the error increases with σ_δ simply because the information friction becomes more severe. I choose $100\sigma_\delta = 0.015$, which matches both the autocorrelation and the size well.

As a preliminary diagnosis of the model's performance, I compare the business cycle moments from the data and the model in Table 3. The model matches the data reasonably well, even for moments that are not explicitly targeted.

5 Results

In this section, I present the results. First, by comparing impulse responses and business cycle moments, I quantify the size of the uncertainty multiplier. Second, I examine the sensitivity of the size of the multiplier to different parameter values for the shock processes. Finally, I highlight the role of real and nominal rigidities by shutting each component one-by-one.

5.1 The Size of the Uncertainty Multiplier

I divide the presentation of the main results into two parts. First, I use impulse responses to explain the basic mechanism of the uncertainty multiplier. Second, I compute business cycle moments and measure the size of the multiplier.

[19] A similar calibration strategy has been used in, for example, Eusepi and Preston (2011) and Görtz and Tsoukalas (2013).

[20] I choose the nominal GDP growth rate because this is the longest forecast series available from the survey. Also, the forecasts do not appear to be biased because the time-series average of the forecast errors is very close to zero.

[21] For the computation of the numbers reported in this table, I only change the value of σ_δ and fix other parameters at the benchmark calibration reported in Table 1.

5.1.1 Impulse Response Analysis

Before examining the impulse responses, I need to consider how to measure the effects of endogenous changes in uncertainty. One potential way is to compare the baseline model with a version of the model without any information friction (i.e., agents know the true value of the shocks). However, this approach is problematic since it confounds the effects of changes in the *variance* of the agents' estimates (which is the main focus of the paper) with the effects of changes in the *mean* of the estimates. Therefore, I consider a version of the model where agents' perception of the variance of the estimates is held constant but they still face information frictions. This way, I can precisely quantify the contribution of fluctuations in uncertainty to business cycle dynamics.

I examine the impulse responses to a one-standard-deviation contractionary monetary policy shock. Recall that from the perspective of the agent at the end of period $t-1$, the capital accumulation equation can be rewritten as follows:

$$K_t = (1-\delta_t)K_{t-1} + (\tilde{\mu}_{t|t-1} + u_t)I_{t-1},$$

where u_t is normally distributed with mean zero and variance $\sigma^2_{u,t}$. In the baseline model featuring the uncertainty multiplier, $\sigma^2_{u,t}$ is given by

$$\sigma^2_{u,t} = \rho^2_\mu \Sigma^{11}_{t-1} + 2\rho_\mu \Sigma^{12}_{t-1} + \Sigma^{22}_{t-1} + \sigma^2_\mu.$$

I shut down the uncertainty multiplier by fixing expectations over $\sigma^2_{u,t}$ at its steady-state level:

$$\sigma^2_{u,t} = \rho^2_\mu \Sigma^{11}_{ss} + 2\rho_\mu \Sigma^{12}_{ss} + \Sigma^{22}_{ss} + \sigma^2_\mu,$$

where Σ^{11}_{ss}, Σ^{12}_{ss}, and Σ^{22}_{ss} are the steady-state levels of Σ^{11}_t, Σ^{12}_t, and Σ^{22}_t. Intuitively, agents act as if ex-ante uncertainty is constant, even though the size of the ex-post forecast error about the marginal efficiency of investment is time-varying.

Figure 2 shows that the output decline in response to a monetary policy shock is deeper when the uncertainty multiplier is present. This is because, as shown in Figure 3, in the baseline model agents perceive an increase in uncertainty about the future realization of effective capital (increase in $\sigma_{u,t}$). The increase in uncertainty is due to a decline in investment originated from a contractionary monetary policy shock. This increase in uncertainty contributes to the additional drop in output compared to the case where the uncertainty multiplier is turned off ($\sigma_{u,t}$ is held constant). Figure 3 also shows that the declines in other real variables are amplified by a similar amount. However, for nominal variables like inflation and the interest rate, the amplification is negligible.

The uncertainty multiplier amplifies the contraction in economic activity for the following reasons. Due to the precautionary motive, an increase in uncertainty induces households to consume less and save more. However, on the saving side, the physical capital becomes a worse hedge for aggregate shocks because the return on capital is subject to more uncertainty. On net, this risk-aversion channel dominates and investment falls as well.

Why, then, do the working hours fall? On the one hand, the fall in consumption induces a desire for households to supply more labor. On the other hand, since aggregate demand is lower, firms demand less labor for a given wage. Since wages are sticky, wages cannot adjust to accommodate more labor and thus equilibrium hours fall. Since prices are sticky, firms increase their price markups and this leads to a further decline in hours. The overall outcome is that output drops substantially.

It is important to stress that in my model, an increase in uncertainty generates a simultaneous fall in output, investment, consumption, and hours. In standard real business cycle models, an increase in uncertainty reduces consumption but also induces a "precautionary labor supply" (Basu and Bundick 2011). As a result, contrary to the data, consumption and hours move in opposite directions. With nominal rigidities, the business cycle comovement is restored through countercyclical movements in markups.

5.1.2 Business Cycle Moments

I measure the size of the uncertainty multiplier by computing the business cycle moments with and without the multiplier. Figure 4 plots the sample path of output from numerical simulations. The uncertainty multiplier amplifies both booms and recessions because uncertainty decreases during booms and increases during recessions. To quantify the magnitude of the amplification, Table 4 compares the standard deviations of output and other variables. The size of the amplification is nontrivial. In particular, the standard deviation of output is 1.18 times larger with the multiplier.[22] Other real variables like investment and hours are amplified by a similar amount.[23] Consistent with the findings from the impulse response analysis, for inflation and the interest rate the amplification is negligible.

Table 5 reports the output uncertainty multiplier for various parameterizations of the standard deviation of the depreciation shock, σ_δ. First, note that the relationship between the size of σ_δ and the multiplier is non-monotonic for the reason discussed in the previous section. Second, for a reasonable range of parameterizations of σ_δ, the size of the multiplier is similar to the baseline value. For example, consider $100\sigma_\delta = 0.050$. While this parameterization implies that the autocorrelation is too low and the forecast errors are too large, the uncertainty multiplier for output is 1.16.

[22] The baseline numbers are derived from the HP-filtered ($\lambda = 1600$) moments. The multiplier is of similar magnitude when other detrending methods are used. For example, when I use linearly detrended moments, the uncertainty multiplier for output is 1.28.

[23] The amplification of consumption is smaller than other real variables because I used a flexible method (HP filter) to detrend the data. When I use linearly detrended moments, the uncertainty multipliers for investment, consumption, and hours are 1.26, 1.26, and 1.28, respectively.

5.2 Changing the Parameters of the Shock Processes

I consider the effects of changing the parameters of the shock processes from the benchmark calibration. The exercise provides additional insights regarding determinants of the size of the uncertainty multiplier.

Table 6 reports the uncertainty multiplier for output under different parameterizations of the standard deviation of the investment shock σ_μ and the growth shock σ_g. I change the ratio of the standard deviations, σ_μ/σ_g, from the benchmark calibration ($\sigma_\mu/\sigma_g = 0.35$) while keeping the standard deviation of output constant. The multiplier is increasing in the relative size of the growth shock. Intuitively, agents respond more to changes in uncertainty about the expected trend growth than to those about the fluctuation around the trend. The uncertainty multiplier is also increasing in the absolute size of the shocks. This can be seen in Table 7, where I scale the standard deviations of shocks ($\sigma_z, \sigma_d, \sigma_R, \sigma_\mu, \sigma_g$, and σ_δ) proportionally from the benchmark calibration. The reason is that the fluctuation in uncertainty becomes more important to agents' decision making as the volatility of shocks becomes larger.

The results in this subsection have an interesting implication for emerging market economies. As shown in Aguiar and Gopinath (2007), these economies feature more volatile business cycles that could be well characterized by fluctuations in expected growth rates.[24] This suggests that the uncertainty multiplier may be much larger in emerging markets than in the U.S.

5.3 The Role of Real and Nominal Rigidities

The benchmark model features several real and nominal rigidities that are absent in a plain vanilla real business cycle model. Table 8 reports the uncertainty multiplier for output under various combinations of frictions.

I highlight three observations. First, nominal rigidities are crucial for generating sizable multipliers. The output uncertainty multiplier is 1.05 without sticky prices and 1.02 without sticky wages. This point connects to Basu and Bundick (2011) and Fernández-Villaverde et al. (2012), who argue that countercyclical markups due to nominal rigidities are important in accounting for the quantitative effects of changes in uncertainty. Second, frictions on the real side of the economy also matter. The real rigidities magnify households' response to changes in uncertainty because (i) they make future adjustments in consumption and investment more costly and (ii) they make investment more persistent and hence make uncertainty more persistent. Third, there are interactions among each set of rigidities. For example, while both the economy with real rigidities only and the economy with nominal rigidities only produce negligible output amplification (1.00 and 1.02, respectively), when the full set of rigidities is present, the amplification is nontrivial (1.18).

[24] See also Boz et al. (2011), who extend Aguiar and Gopinath (2007)'s analysis by incorporating a learning problem.

6 Survey Data Evidence

In this section, I use survey data that directly measure subjective uncertainty and argue that the model is consistent with the data. In particular, I show that the model replicates the VAR relationship between output and uncertainty.

Since uncertainty is an ex-ante concept, its measurement using ex-post realized data is inherently difficult. Probabilistic forecasts reported in the Survey of Professional Forecasters are unique in yielding numeric values on ex-ante uncertainty for a sufficiently long period of time. This survey asks each forecaster for a subjective probability density of the annual percentage change in real GDP. Following the standard in the literature (Zarnowitz and Lambros 1987 and D'Amico and Orphanides 2008), I take the average across the standard deviations of those probability densities for each forecaster and use it as a measure of uncertainty.[25] While the survey data start from 1968:Q4, concerns regarding data consistency and missing data force me to conduct the analysis using the data during the period 1986:Q2–2011:Q4.[26] Finally, since the survey asks for the percentage change in GDP between the previous and current calendar year, there is a seasonality in the forecast horizons. For example, in the first quarter, it is a 4-quarter-ahead forecast. In the second quarter, it is a 3-quarter-ahead forecast. I eliminate this seasonality by applying the Tramo-Seats filter.[27]

I characterize the relationship between real GDP and uncertainty with a generalized impulse response analysis (Pesaran and Lambros 1998) from a bivariate VAR with four lags. The generalized impulse response is appealing in this context because, in contrast to a standard recursive VAR, the results are invariant to the ordering of variables.[28] Both variables are logged and HP-filtered with $\lambda = 1600$. I emphasize that the purpose of this exercise is to look for a statistical relationship between output and uncertainty. Hence, no causal inference is drawn from the impulse responses.

Figure 5 shows that, in the data, an increase in uncertainty is associated with a decline in output that reaches a trough after five quarters. On the other hand, an increase in output is associated with a decline in uncertainty. Hence the VAR responses indicate a clear negative relationship between output and uncertainty. The figure also shows that running a VAR on the artificial data from the model generates impulse responses that are in line with the actual data.[29] In the model,

[25]The survey asks each forecaster to place probabilities in bins spanning a wide range of outcomes for the percentage change in real GDP. To compute the individual standard deviations, I fit a normal distribution to the individual probabilities. For more details, see D'Amico and Orphanides (2008). I have also tried other methods and obtained similar results.

[26]Nevertheless I conducted the analysis using the whole sample period and found similar results.

[27]Since the survey response between the current and the following year is also available, it is possible to construct uncertainty data with different forecast horizons. I have conducted the analysis with different forecast horizons and found similar results.

[28]Nevertheless I also tried a recursive VAR and obtained similar results.

[29]In the model, I define uncertainty as the standard deviation of the density forecast (conditional on the agents' information sets) of the annual percentage change in output: $Std_{t+s|t}(\Delta Y_{t+s})$. The forecast horizon is chosen in a way consistent with the survey data.

the negative relationship between output and uncertainty is due to the endogenous movement in uncertainty and its feedback to real economic activity. Note that the model replicates well the gradual responses of the two variables. This is because uncertainty is driven by investment, which exhibits hump-shaped dynamics, and this uncertainty in turn induces gradual adjustments by households.

7 Conclusion

Much learning about macroeconomic conditions seems to occur through actually undertaking economic activity. This paper formalized the idea in an equilibrium business cycle framework and explored its quantitative implications. Recessions are times of high uncertainty because agents invest less and hence learn less about the state of the economy. The endogenous fluctuations in aggregate uncertainty interact with rigidities and amplify business cycles.

Because the level of learning is tied to the level of investment, changes in uncertainty are large and persistent. As a result, the size of the amplification is nontrivial. Under the benchmark calibration, the uncertainty multiplier amplifies the standard deviation of output by 18%. Other real variables, such as investment and hours, are also amplified by a similar amount.

References

Aguiar, Mark and Gita Gopinath, "Emerging Market Business Cycles: The Cycle Is the Trend," *Journal of Political Economy*, 2007, *115*, 69–102.

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Lindé, "Firm-Specific Capital, Nominal Rigidities and the Business Cycle," *Review of Economic Dynamics*, 2011, *14*, 225–247.

Arellano, Cristina, Yan Bai, and Patrick Kehoe, "Financial Markets and Fluctuations in Volatility," 2012. Manuscript, Federal Reserve Bank of Minneapolis.

Aruoba, S. Borağan, Jesús Fernández-Villaverde, and Juan Rubio-Ramírez, "Comparing Solution Methods for Dynamic Equilibrium Economies," *Journal of Economic Dynamics and Control*, 2006, *30*, 2477–2508.

Bachmann, Rüdiger and Christian Bayer, "'Wait-and-See' Business Cycles?," *Journal of Monetary Economics*, 2013, *60*, 704–719.

_ **and Giuseppe Moscarini**, "Business Cycles and Endogenous Uncertainty," 2011. Manuscript, RWTH Aachen University.

_ **, Steffen Elstner, and Eric Sims**, "Uncertainty and Economic Activity: Evidence from Business Survey Data," *American Economic Journal: Macroeconomics*, 2013, *5*, 217–249.

Barsky, Robert B. and Eric R. Sims, "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, 2012, *102*, 1343–1377.

Basu, Susanto and Brent Bundick, "Uncertainty Shocks in a Model of Effective Demand," 2011. Manuscript, Boston College.

Beaudry, Paul and Franck Portier, "Exploring Pigou's Theory of Cycles," *Journal of Monetary Economics*, 2004, *51*, 1183–1216.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77*, 623–685.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 2001, *91*, 149–166.

Born, Benjamin and Johannes Pfeifer, "Policy Risk and the Business Cycle," 2012. Manuscript, University of Bonn.

Boz, Emine, Christian Daude, and C. Bora Durdu, "Emerging Market Business Cycles: Learning About the Trend," *Journal of Monetary Economics*, 2011, *58*, 616–631.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review Papers and Proceedings*, 2011, *101*, 471–475.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, 2005, *113*, 1–45.

_ , Roberto Motto, and Massimo Rostagno, "Risk Shocks," *American Economic Review*, forthcoming.

Chugh, Sanjay K., "Firm Risk and Leverage Based Business Cycles," 2012. Manuscript, Boston College.

Cogley, Timothy, Riccardo Colacito, and Thomas J. Sargent, "Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas," *Journal of Money, Credit, and Banking*, 2007, *39*, 67–99.

Cooley, Thomas F. and Edward C. Prescott, "Economic Growth and Business Cycles," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 1–38.

D'Amico, Stefania and Athanasios Orphanides, "Uncertainty and Disagreement in Economic Forecasting," 2008. Manuscript, Board of Governors of the Federal Reserve System.

D'Erasmo, Pablo N and Hernan J Moscoso Boedo, "Intangibles and Endogenous Firm Volatility over the Business Cycle," 2012. Manuscript, University of Maryland.

Eusepi, Stefano and Bruce Preston, "Expectations, Learning, and Business Cycle Fluctuations," *American Economic Review*, 2011, *101*, 2844–2872.

Fajgelbaum, Pablo, Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty Traps," 2013. Manuscript, UCLA.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan Rubio-Ramírez, and Martin Uribe, "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 2011, *101*, 2530–2561.

_ , _ , Keith Kuester, and Juan Rubio-Ramírez, "Fiscal Volatility Shocks and Economic Activity," 2012. Manuscript, University of Pennsylvania.

Fisher, Jonas D. M., "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, 2006, *114*, 413–451.

Gertler, Mark and Peter Karadi, "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 2011, *58*, 17–34.

Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek, "Uncertainty, Financial Frictions, and Investment Dynamics," 2010. Manuscript, Boston University.

Görtz, Christoph and John D. Tsoukalas, "Learning, Capital-Embodied Technology and Aggregate Fluctuations," *Review of Economic Dynamics*, 2013, *16*, 708–723.

Gourio, François, "Disaster Risk and Business Cycles," *American Economic Review*, 2012, *102*, 2734–2766.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman, "Investment, capacity utilization, and the real business cycle," *American Economic Review*, 1988, *78*, 402–417.

Haan, Wouter J. Den and Joris De Wind, "Nonlinear and Stable Perturbation-Based Approximations," 2012. Manuscript, London School of Economics.

Ilut, Cosmin and Martin Schneider, "Ambiguous Business Cycles," 2011. Manuscript, Duke University.

Jaimovich, Nir and Sergio Rebelo, "Can News about the Future Drive the Business Cycle?," *American Economic Review*, 2009, *99*, 1097–1118.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 2010, *57*, 132–145.

_ , _ , and _ , "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 2011, *14*, 102–121.

Kehrig, Matthias, "The Cyclicality of Productivity Dispersion," 2011. Manuscript, University of Texas at Austin.

Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Christopher A. Sims, "Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models," *Journal of Economic Dynamics and Control*, 2008, *32*, 3397–3414.

Levin, Andrew T., Alexei Onatski, John C. Williams, and Noah Williams, "Monetary Policy Under Uncertainty in Micro-Founded Macroeconomic Models," in "NBER Macroeconomics Annual 2005," MIT press, 2006, pp. 229–287.

Liu, Zheng, Daniel F. Waggoner, and Tao Zha, "Sources of Macroeconomic Fluctuations: A Regime-Switching DSGE Approach," *Quantitative Economics*, 2011, *2*, 251–301.

Lorenzoni, Guido, "A Theory of Demand Shocks," *American Economic Review*, 2009, *99*, 2050–2084.

Lucas, Robert E., "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 1972, *4*, 103–124.

Ordoñez, Guillermo, "The Asymmetric Effects of Financial Frictions," 2012. Manuscript, University of Pennsylvania.

Orlik, Anna and Laura Veldkamp, "Understanding Uncertainty Shocks," 2013. Manuscript, New York University.

Pesaran, H. Hashem and Louis A. Lambros, "Generalized Impulse Response Analysis in Linear Multivariate Models," *Economics Letters*, 1998, *58*, 17–29.

Schaal, Edouard, "Uncertainty, Productivity and Unemployment in the Great Recession," 2012. Manuscript, New York University.

Schmitt-Grohe, Stephanie and Martin Uribe, "What's News in Business Cycles," *Econometrica*, 2012, *80*, 2733–2764.

Smets, Frank and Raf Wouters, "Shocks and frictions in US business cycles: a Bayesian DSGE approach," *American Economic Review*, 2007, *97*, 586–606.

Tian, Can, "Riskiness Choice and Endogenous Productivity Dispersion over the Business Cycle," 2012. Manuscript, University of Pennsylvania.

Van Nieuwerburgh, Stijn and Laura Veldkamp, "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 2006, *77*, 753–772.

Veldkamp, Laura, "Slow Boom, Sudden Crash," *Journal of Economic Theory*, 2005, *124*, 230–257.

Zarnowitz, Victor and Louis A. Lambros, "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy*, 1987, *95*, 591–621.

Appendix

A Data Source

The data set spans the period 1969Q1 to 2011Q4.[30] Whenever the data set is provided in monthly frequencies, I simply take the average to transform it into quarterly frequencies.

Data from the National Income and Product Accounts are downloaded from the Bureau of Economic Analysis website. Nominal GDP, nominal consumption (defined as the sum of personal consumption expenditures on nondurables and services), and nominal investment (defined as the sum of gross private domestic investment and personal consumption expenditures on durables) are divided by the civilian noninstitutional population,[31] downloaded from the Bureau of Labor Statistics (BLS hereafter) website, to convert the variables into per capita terms. I then divide them by the GDP deflator to convert them into real terms.

Working hours are measured by nonfarm business hours (available on the BLS website) divided by the population. Real wages are measured by hourly compensation in nonfarm business sectors (available on the BLS website) divided by the GDP deflator. Inflation rates are measured by changes in the GDP deflator. I use the effective federal funds rates (downloaded from the Federal Reserve Board website) to measure the nominal interest rates.

To compute the forecast error statistics, I use the median forecast of nominal GDP growth rate, downloaded from the FRB Philadelphia website. The one-period-ahead forecast error is defined as the one-period-ahead nominal GDP growth rate forecast minus the realized nominal GDP growth rate.

B Countercyclical Uncertainty: Full Derivation

I restate agents' Kalman-filtering problem below:

$$\begin{bmatrix} \mu_t \\ g_t \end{bmatrix} = \begin{bmatrix} (1-\rho_\mu)\mu \\ 0 \end{bmatrix} + \begin{bmatrix} \rho_\mu & 1 \\ 0 & \rho_g \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ g_{t-1} \end{bmatrix} + \begin{bmatrix} \epsilon_{\mu,t} \\ \epsilon_{g,t} \end{bmatrix},$$

$$K_t - (1-\delta)K_{t-1} = [I_{t-1} \quad 0] \begin{bmatrix} \mu_t \\ g_t \end{bmatrix} + K_{t-1}\epsilon_{\delta,t}.$$

[30] I pick this starting date because the Survey of Professional Forecasters began around that time.

[31] Since raw population data display occasional breaks due to changes in population controls, I use an HP-filtered ($\lambda = 1600$) trend instead.

At the end of period $t-1$, agents forecast the values of $\{\mu_t, g_t\}$:

$$\begin{aligned}
\tilde{\mu}_{t|t-1} &= (1-\rho_\mu)\mu + \rho_\mu \tilde{\mu}_{t-1} + \tilde{g}_{t-1}, \\
\tilde{g}_{t|t-1} &= \rho_g g_{t-1}.
\end{aligned}$$

The elements of the associated forecasting error covariance matrix, $\Sigma_{t|t-1}$, are

$$\begin{aligned}
\Sigma^{11}_{t|t-1} &= \rho_\mu^2 \Sigma^{11}_{t-1} + 2\rho_\mu \Sigma^{12}_{t-1} + \Sigma^{22}_{t-1} + \sigma_\mu^2, \\
\Sigma^{12}_{t|t-1} &= \rho_\mu \rho_g \Sigma^{12}_{t-1} + \rho_g \Sigma^{22}_{t-1}, \\
\Sigma^{21}_{t|t-1} &= \Sigma^{12}_{t|t-1}, \\
\Sigma^{22}_{t|t-1} &= \rho_g^2 \Sigma^{22}_{t-1} + \sigma_g^2.
\end{aligned}$$

After observing period t realization of capital, K_t, agents update their belief according to

$$\begin{aligned}
\tilde{\mu}_t &= \tilde{\mu}_{t|t-1} + \frac{I_{t-1}\Sigma^{11}_{t|t-1}}{I^2_{t-1}\Sigma^{11}_{t|t-1} + K^2_{t-1}\sigma^2_\delta} \cdot \{K_t - (1-\delta)K_{t-1} - I_{t-1}\tilde{\mu}_{t-1}\}, \\
\tilde{g}_t &= \tilde{g}_{t|t-1} + \frac{I_{t-1}\Sigma^{12}_{t|t-1}}{I^2_{t-1}\Sigma^{11}_{t|t-1} + K^2_{t-1}\sigma^2_\delta} \cdot \{K_t - (1-\delta)K_{t-1} - I_{t-1}\tilde{\mu}_{t-1}\}.
\end{aligned}$$

The elements of the forecasting error covariance matrix are given by

$$\begin{aligned}
\Sigma^{11}_t &= \left[1 - \frac{I^2_{t-1}\Sigma^{11}_{t|t-1}}{I^2_{t-1}\Sigma^{11}_{t|t-1} + K^2_{t-1}\sigma^2_\delta}\right]\Sigma^{11}_{t|t-1}, \\
\Sigma^{12}_t &= \left[1 - \frac{I^2_{t-1}\Sigma^{11}_{t|t-1}}{I^2_{t-1}\Sigma^{11}_{t|t-1} + K^2_{t-1}\sigma^2_\delta}\right]\Sigma^{12}_{t|t-1}, \\
\Sigma^{21}_t &= \Sigma^{12}_t, \\
\Sigma^{22}_t &= \Sigma^{22}_{t|t-1} - \frac{I^2_{t-1}(\Sigma^{12}_{t|t-1})^2}{I^2_{t-1}\Sigma^{11}_{t|t-1} + K^2_{t-1}\sigma^2_\delta}.
\end{aligned}$$

Thus, the elements of Σ_t are decreasing in $\frac{I_{t-1}}{K_{t-1}}$.

C Investment Problem in a Decentralized Equilibrium

To implement the investment problem in a decentralized competitive equilibrium, consider perfectly competitive capital producers owned by the households. At the end of each period t, they purchase investment goods I_t and capital K_t from households. The price of investment goods relative to consumption goods is unity and the price of capital is $\tilde{Q}_t$. In period $t+1$, they build new capital K_{t+1} using the technology (5). The capital producers can observe the path of capital stock and investment but cannot observe the underlying shocks. The new capital is sold at price Q_{t+1}. The profits are transferred back in a lump-sum manner each period.

The capital producers choose the inputs I_t and K_t to maximize their discounted sum of profits:

$$\max_{I_t, K_t} -\lambda_t(I_t + \tilde{Q}_t K_t) + \mathrm{E}_t \sum_{s=0}^{\infty} \beta^{s+1} \lambda_{t+s+1} \Bigg[Q_{t+s+1} \Bigg\{ (1-\delta_{t+s+1}) K_{t+s} + \mu_{t+s+1} \Bigg(1 - S\Bigg(\frac{I_{t+s}}{I_{t+s-1}} \Bigg) \Bigg) I_{t+s} \Bigg\} - I_{t+s+1} - \tilde{Q}_{t+s+1} K_{t+s+1} \Bigg].$$

The first-order conditions of this profit maximization problem yield an evolution for the expected price of capital:

$$\lambda_t = \beta \mathrm{E_t} \left[\lambda_{t+1} Q_{t+1} \mu_{t+1} \left\{ 1 - S\left(\frac{I_t}{I_{t-1}} \right) - S'\left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} \right\} + \beta \lambda_{t+2} Q_{t+2} \mu_{t+2} S'\left(\frac{I_{t+1}}{I_t} \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right].$$

They also provide an expression for the "rental" rate of capital:

$$\lambda_t \tilde{Q}_t = \beta \mathrm{E}_t[(1-\delta_t)\lambda_{t+1} Q_{t+1}],$$

which says that the price of capital at the end of period takes into account the discounting and depreciation that occur at the beginning of the next period.

Finally, relative prices like the price of capital, Q_t, do not reveal additional information about the unobservable shocks. This is simply because capital producers also face the filtering problem described in Section 2 in the main text.

D Procyclical Signal-to-Noise Ratio Arising From Aggregation

The procyclical signal-to-noise ratio arises from aggregation of investment units with common and idiosyncratic shocks. The argument does not require depreciation shocks and closely follows the discussion made in Van Nieuwerburgh and Veldkamp (2006).

Consider an economy with many investment units, where each unit has a technology that transforms investment goods into efficiency units of capital. The capital production is increasing in the number of investment units operating. Denote N_t as the number of investment units operating at time t. The output of each unit is the product of its own efficiency, which has a common component μ_t and idiosyncratic component η_t^i, and its input normalized to $i_t^i = 1$. Then, aggregate net investment is the sum of output of all the investment units:

$$K_t - (1-\delta)K_{t-1} = \sum_{i=1}^{N_t} (\mu_t + \eta_t^i) i_t^i = \mu_t N_t + \sum_{i=1}^{N_t} \eta_t^i$$

Define the signal-to-noise ratio as $\mathrm{Var}(\mu_t N_t)/\mathrm{Var}(\sum_{i=1}^{N_t} \eta_t^i)$. Then, as long as the correlation of idiosyncratic shocks across investment units is less than one, the signal-to-noise ratio is increasing in the number of units operating, N_t. A more detailed proof of this argument can be found in Van Nieuwerburgh and Veldkamp (2006).

Figure 1: Timing of events

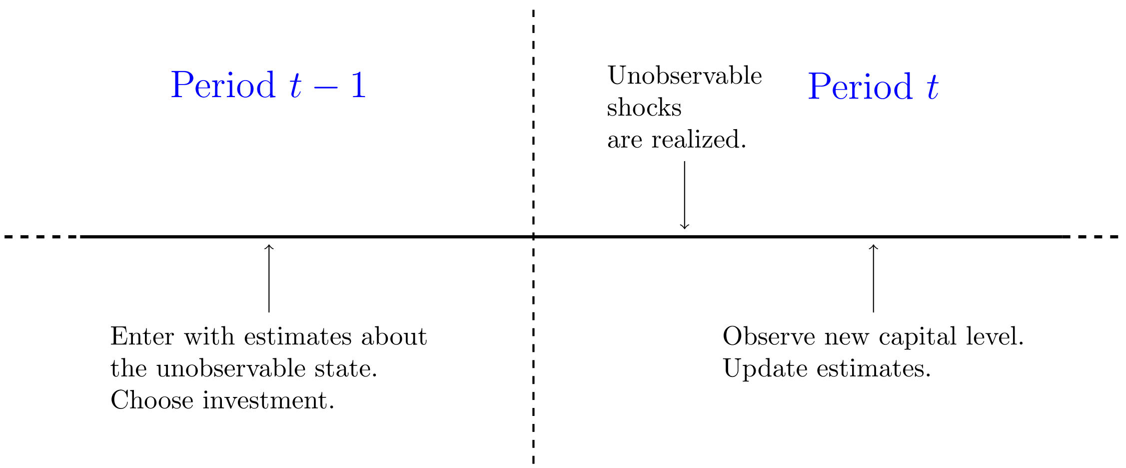

Figure 2: The uncertainty multiplier amplifies output response to a monetary policy shock

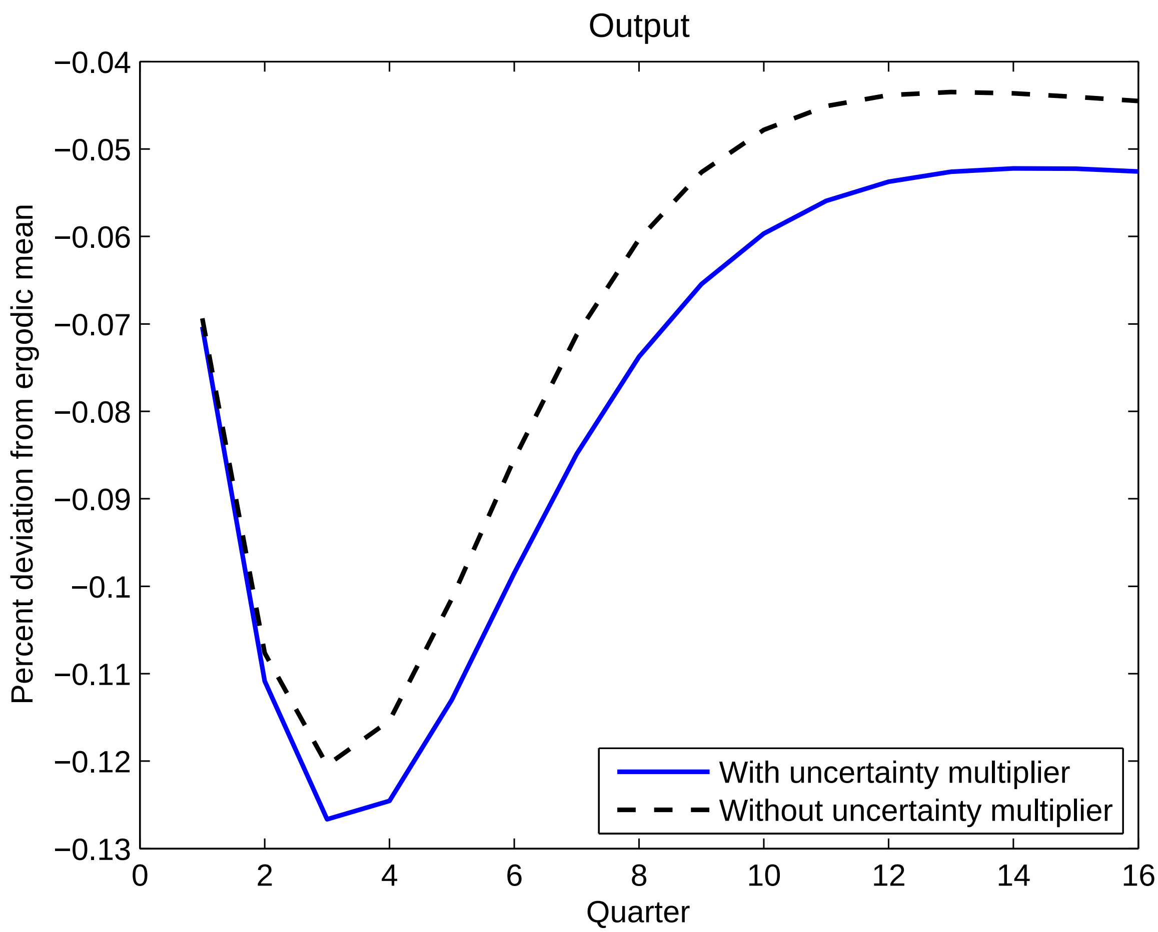

Notes: Since third-order approximations move the ergodic distribution of endogenous variables away from the steady state (Fernández-Villaverde et al. 2011), I report the impulse responses in terms of percent deviation from the ergodic mean.

Figure 3: Responses of other real variables are also amplified

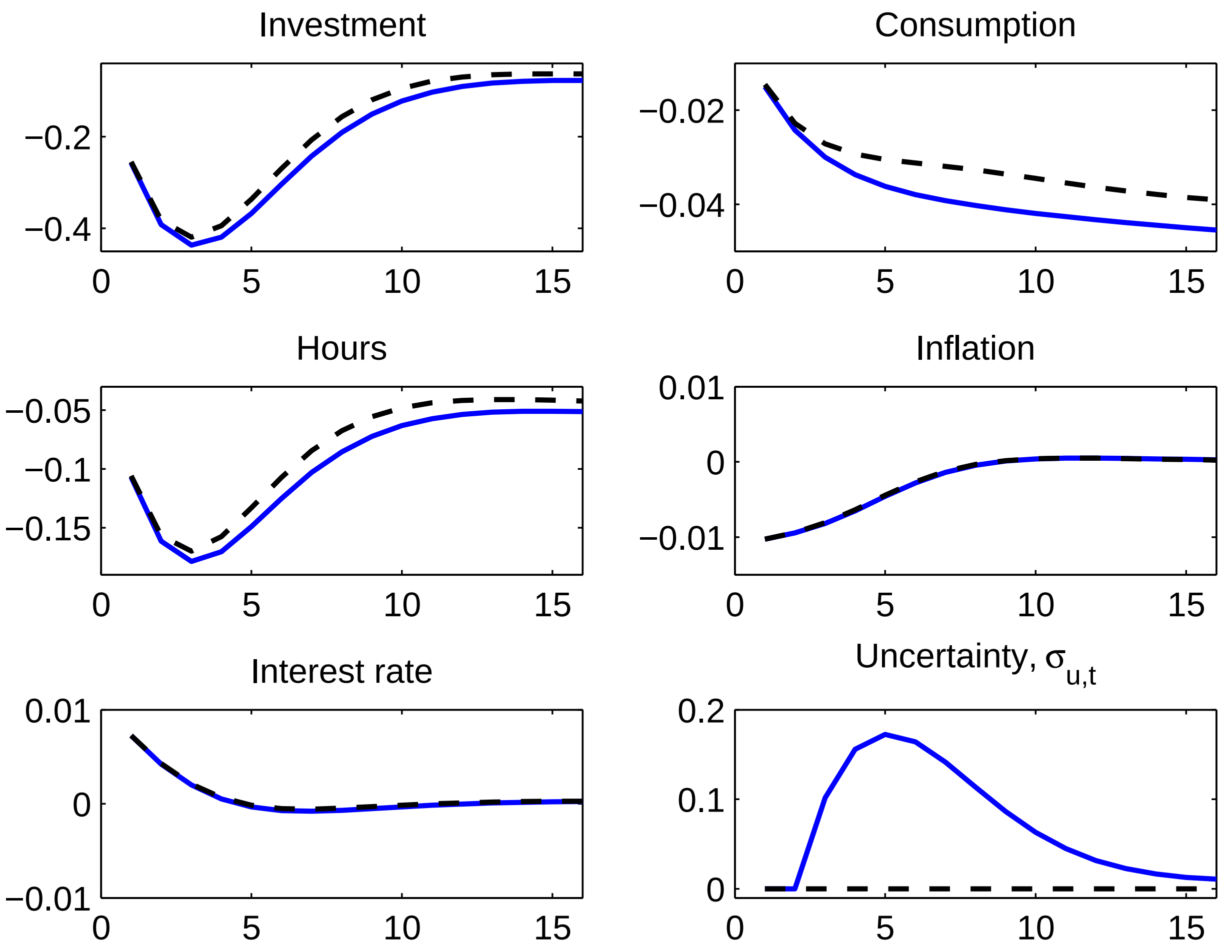

Notes: See the notes for Figure 2.

Figure 4: The uncertainty multiplier amplifies business cycles

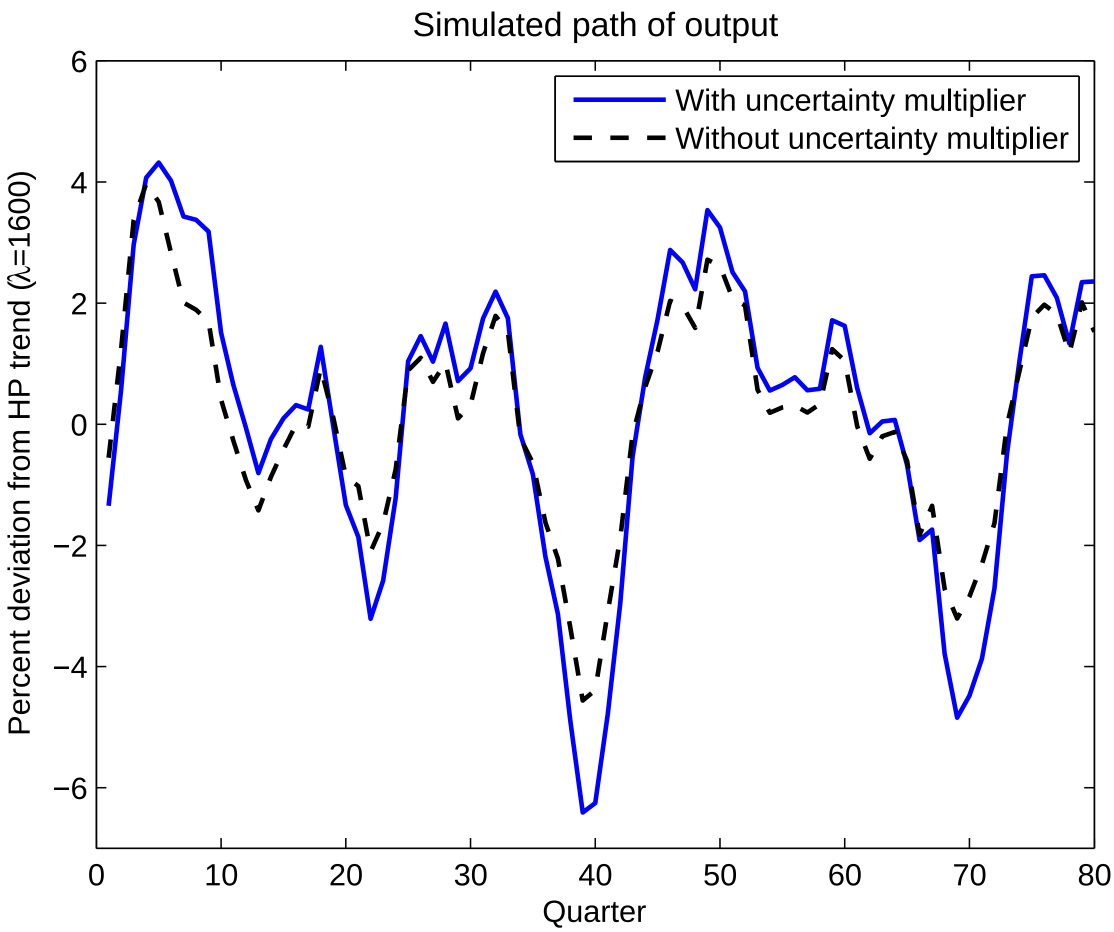

Figure 5: Impulse responses in a bivariate VAR

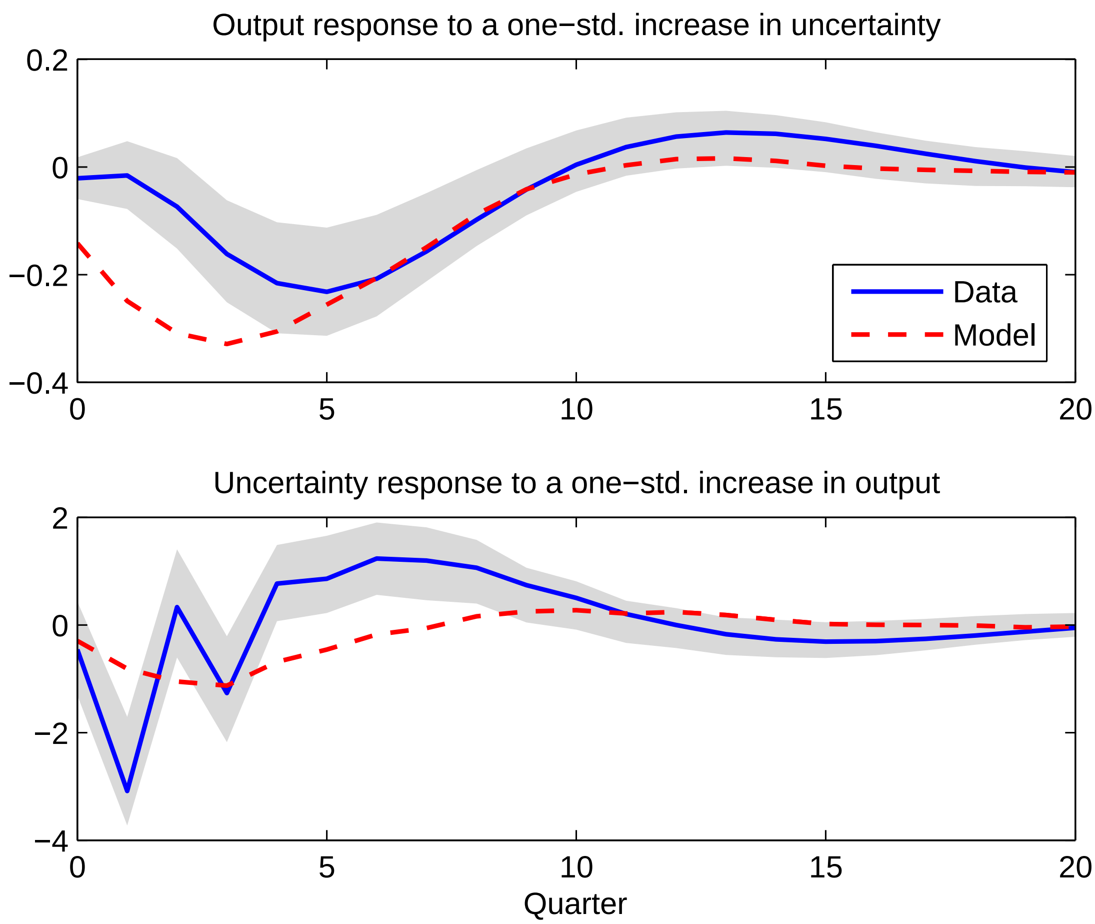

Notes: The shaded area represents ± one-standard-deviation bootstrap confidence bands.

Table 1: Parameters and targets

	Description	Value	Comments/Targets
Technology and preference			
β	Discount factor	0.9948	Historical mean of interest rate
θ_p	Goods demand elasticity	21	5% price markup (Altig et al. 2011)
θ_w	Labor demand elasticity	21	5% wage markup (Altig et al. 2011)
α	Capital share	0.3	Standard choice
δ	Depreciation rate	0.02	8% annual depreciation
σ	Preference parameter	2	Standard choice
η	Inverse Frisch elasticity	2	Frisch elasticity = 0.5 (Chetty et al. 2011)
b	Habit persistence	0.75	Smets and Wouters (2007)
κ	Investment adj. cost	0.34	*Calibrated*
ξ_p	Calvo price	0.75	Duration of price 4 quarters
ξ_w	Calvo wage	0.75	Duration of wage 4 quarters
Monetary policy			
π	SS inflation rate	1.0095	Historical mean of inflation rate
ρ_R	Taylor rule smoothing	0.9	Standard choice
ϕ_π	Taylor rule inflation	2	Standard choice
ϕ_Y	Taylor rule output growth	0.1	Standard choice
Shock process			
ρ_z	Technology	0.95	Cooley and Prescott (1995)
ρ_d	Preference	0.22	Smets and Wouters (2007)
ρ_μ	Investment level	0.91	*Calibrated*
ρ_g	Investment growth	0.855	*Calibrated*
$100\sigma_z$	Technology	0.2	*Calibrated*
$100\sigma_d$	Preference	4.2	*Calibrated*
$100\sigma_R$	Monetary policy	0.01	*Calibrated*
$100\sigma_\mu$	Investment level	0.25	*Calibrated*
$100\sigma_g$	Investment growth	0.72	*Calibrated*
$100\sigma_\delta$	Depreciation	0.015	*Calibrated*

Table 2: Identification of σ_δ from survey data moments

	$Corr(FE_t^{1Q}, FE_{t-1}^{1Q})$	$Mean(\|FE_t^{1Q}\|)$
Data	0.17	0.55
Model		
$100\sigma_\delta = 0.005$	0.18	0.49
$100\sigma_\delta = 0.010$	0.19	0.53
$100\sigma_\delta = 0.015$	0.17	0.58
$100\sigma_\delta = 0.025$	0.14	0.69
$100\sigma_\delta = 0.050$	0.13	0.92

Notes: The forecast errors are multiplied by 100 to express them in percentage terms. The data statistics are calculated using the final data vintage. As a robustness check, I calculated the statistics using alternative data vintages and found that they are similar. For example, $(Corr(FE_t^{1Q}, FE_{t-1}^{1Q}), Mean(|FE_t^{1Q}|))$ for the first, the third, and the fifth vintages are (0.23, 0.47), (0.18, 0.52), and (0.20, 0.54), respectively.

Table 3: Business cycle moments

	Std.	$Corr(Y_t, X_t)$	AR(1)
Data			
Output	1.61	1.00	0.87
Investment	6.31	0.94	0.87
Consumption	0.93	0.84	0.87
Hours	1.99	0.88	0.92
Real wage	0.84	0.07	0.76
Inflation	0.29	0.18	0.48
Interest rate	0.41	0.34	0.75
Model			
Output	1.61	1.00	0.89
Investment	6.39	0.90	0.92
Consumption	0.92	0.59	0.81
Hours	2.26	0.98	0.87
Real wage	1.05	-0.36	0.89
Inflation	0.46	0.12	0.68
Interest rate	0.26	-0.29	0.93

Notes: Both data and model moments are in logs, HP-filtered ($\lambda = 1600$), and multiplied by 100 to express them in percentage terms.

Table 4: The size of the amplification

	Amplification $\sigma_{\text{With multiplier}}/\sigma_{\text{Without multiplier}}$
Output	1.18
Investment	1.25
Consumption	1.05
Hours	1.22
Real wage	1.10
Inflation	1.03
Interest rate	1.04

Notes: Both data and model moments are in logs and HP-filtered ($\lambda = 1600$).

Table 5: The uncertainty multiplier for different values of σ_δ

	$Corr(FE_t^{1Q}, FE_{t-1}^{1Q})$	$Mean(\lvert FE_t^{1Q}\rvert)$	Output amplification
Data	0.17	0.55	
Model			
$100\sigma_\delta = 0.005$	0.18	0.49	1.10
$100\sigma_\delta = 0.010$	0.19	0.53	1.17
$100\sigma_\delta = 0.015$	0.17	0.58	1.18
$100\sigma_\delta = 0.025$	0.14	0.69	1.23
$100\sigma_\delta = 0.050$	0.13	0.92	1.16

Notes: Both data and model moments are in logs and HP-filtered ($\lambda = 1600$). The forecast errors are multiplied by 100 to express them in percentage terms.

Table 6: The uncertainty multiplier is increasing in the relative size of the growth shock

	$Corr(FE_t^{1Q}, FE_{t-1}^{1Q})$	$Mean(\|FE_t^{1Q}\|)$	Std. of investment	Output amplification
Data	0.17	0.55	6.31	
Model				
$\sigma_\mu/\sigma_g = 1.00$	0.13	0.69	6.37	1.16
$\sigma_\mu/\sigma_g = 0.35$	0.17	0.58	6.39	1.18
$\sigma_\mu/\sigma_g = 0.00$	0.17	0.56	6.46	1.21

Notes: Both data and model moments are in logs and HP-filtered ($\lambda = 1600$). σ_μ and σ_g are scaled so that the standard deviation of output is the same as that in the benchmark specification ($\sigma_Y = 1.61$).

Table 7: The uncertainty multiplier is increasing in the size of shocks

	Output standard dev.	Output amplification
Data	1.61	
Model		
$\sigma \times 0.9$	1.30	1.14
$\sigma \times 1.0$	1.61	1.18
$\sigma \times 1.1$	2.39	1.32

Notes: Both data and model moments are in logs and HP-filtered ($\lambda = 1600$). I define $\sigma \equiv (\sigma_z, \sigma_d, \sigma_R, \sigma_\mu, \sigma_g, \sigma_\delta)$.

Table 8: The role of real and nominal rigidities

Consump. habit	Investment adj. cost	Sticky price	Sticky wage	Output amplification
✓	✓	✓	✓	1.18
✓	✓		✓	1.05
✓	✓	✓		1.02
✓	✓			1.00
	✓	✓	✓	1.06
✓		✓	✓	1.03
		✓	✓	1.02
				1.00

Notes: Both data and model moments are in logs and HP-filtered ($\lambda = 1600$). For each specification, I scale $(\sigma_z, \sigma_d, \sigma_R, \sigma_\mu, \sigma_g, \sigma_\delta)$ proportionally to generate the standard deviation of output as in the benchmark specification ($\sigma_Y = 1.61$).