Risk and ambiguity in models of business cycles*

David Backus,† Axelle Ferriere,‡ and Stanley Zin§

August 9, 2014

Abstract

We inject aggregate uncertainty — risk and ambiguity — into an otherwise standard business cycle model and describe its consequences. We find that increases in uncertainty generally reduce consumption, but they do not account, in this model, for either the magnitude or the persistence of the most recent recession. We speculate about extensions that might do better along one or both dimensions.

JEL Classification Codes: E32, D81, G12.

Keywords: uncertainty; smooth ambiguity; certainty equivalent; recursive preferences; pricing kernel; asset returns; learning.

*Prepared for the Carnegie-Rochester-NYU Conference, Rochester, April 2014. Comments welcome, including references to related work we inadvertently overlooked. We thank Rudi Bachmann, Anmol Bhandari, Jarda Borovicka, Ana Fostel, Lars Hansen, Cosmin Ilut, Jianjun Miao, Laura Veldkamp, and Venky Vekateswaran for helpful comments. Chase Coleman and Spencer Lyon produced Figures 1-4. The latest version is available at

`http://pages.stern.nyu.edu/~dbackus/BFZ/ms/BFZ_CRN_latest.pdf`.

† Stern School of Business, New York University, and NBER; david.backus@nyu.edu.

‡ Stern School of Business, New York University; axelle.ferriere@nyu.edu.

§ Stern School of Business, New York University, and NBER; stan.zin@nyu.edu.

1 Introduction

In the United States, the most recent recession was the most severe since the 1930s, and it's been followed by an unusually long recovery. We see how this played out in Figures 1 to 3, which are adapted from Bethune, Cooley, and Rupert (2014). The figures illustrate the recession and recovery, relative to the previous peak, in output, consumption, and investment over the last five recessions. Output, for example, hit a cyclical peak in the fourth quarter of 2007, declined by about 4 percent over the next 6 quarters, and only surpassed its previous peak 14 quarters later (Figure 1). The depth and duration are greater than in any of the earlier recessions. The duration reflects both the depth of the recession — it takes longer to dig your way out of a deeper hole — and a slower rate of recovery. We see broadly similar behavior in consumption (Figure 2) and investment (Figure 3). Consumption fell 3 percent and took 12 quarters to return to its previous level. Nonresidential fixed investment fell almost 20 percent and took 24 quarters (6 years!) to recover. Housing, of course, has yet to recover.

These patterns are familiar to all of us. The relative magnitudes are what we tell our students: consumption falls less than output, and investment falls more. The same with comovements: all three variables, and many more besides, declined together. The only somewhat unusual feature here is the slow speed of recovery, particularly in investment.

Since the patterns are familiar, we might expect to account for them with some variant of the Kydland-Prescott (1982) model, in which declines in productivity generate precisely this collection of facts. The unusually large magnitude of the recession would reflect, in this model, an unusually large drop in productivity. The problem with this account is that measured productivity barely fell. We could document this in a number of ways, but the simplest is to measure productivity by the ratio of output to hours worked. We see the result in Figure 4 and in related work by Sprague (2014). Productivity was essentially flat for six quarters until growth resumed, and at its lowest point was only one percent below its peak.

But if productivity wasn't the source of the last recession, what was? Yes, we know, there was a financial crisis. But what shocks — or wedges — do we need to add to the model to reproduce its effects? We have no shortage of candidates, but the leading one right now is "uncertainty." Among the many suggestions to this effect is a wonderful comment from the European Commission (2013, page 7): "Economic theory suggests that uncertainty has a detrimental effect on economic activity by giving agents the incentive to postpone investment, consumption and employment decisions until uncertainty is resolved, and by pushing up the cost of capital through increased risk premia." Bloom (2013), who has written extensively on uncertainty, adds: "The onset of the Great Recession was accompanied by a massive surge in uncertainty. The size of this uncertainty shock was so large it potentially accounted for around one third of the 9% drop in GDP versus trend during 2008-2009." At minimum, we observe strong countercyclical patterns in measures of uncertainty that beg for explanation; see, for example, the evidence reported by Bachmann and Bayer (2013) and Bloom, Floetotto, Jaimovich, Saporta, and Terry (2012).

With this motivation, we study uncertainty in a streamlined business cycle model and ask: Can an increase in uncertainty account for the magnitude and persistence of the last recession in this model? How does uncertainty affect the dynamics of consumption, investment, and output? Does it interfere with traditional business cycle comovements, in which all of these variables move up and down together? Can uncertainty account for differences in the responses of consumption and investment during the latest recovery?

To answer these questions, we inject uncertainty into a streamlined version of the Kydland and Prescott (1982) model with constant labor supply. We add three ingredients: recursive preferences, a unit root in the aggregate productivity process, and several kinds of uncertainty. Recursive preferences are a natural generalization of the additive preferences used in most business cycle models. A large body of work suggests that their extra flexibility is helpful in accounting for asset prices (Bansal and Yaron, 2004, for example) but has little impact on the behavior of macroeconomic quantities (Tallarini, 2000). In this respect they differ markedly from habit-based preferences, which affect both quantities and asset prices. The unit root is essential to delivering realistic risk premiums. Without it, the asset with the largest risk premium is a long-maturity bond (Alvarez and Jermann, 2005, Proposition 2), which is not what we see in the data.

The most important new ingredient is the third one: uncertainty. We specify stochastic processes for the conditional mean and variance of aggregate productivity growth. Preferences play a role here in the representative agent's responses to these sources of uncertainty, whether risk (where the probabilities are well understood) or ambiguity (where they're not). Typically greater aversion to uncertainty leads to stronger responses of consumption to fluctuations in the conditional variance. We consider preferences toward risk, in which the distribution of outcomes is understood by our representative agent, and ambiguity, in which it is not.

We use these ingredients to assess the impact of uncertainty on the magnitude and persistence of economic fluctuations. Consumption and investment decisions are functions of the state, which here includes a state variable representing uncertainty. We find that consumption typically falls when uncertainty rises, although there are parameter configurations in which the reverse is true. The magnitude of the effect depends on both the IES and uncertainty aversion. In this respect, recursive preferences are central to the transmission of uncertainty to the economy. The effects, however, are small: fluctuations in uncertainty play a minor role in the model's business cycle properties.

Does uncertainty interfere with business cycle comovements, the tendency for output, consumption, and investment to move up and down together? Barro and King (1984) showed that shocks to anything but productivity generate counterfactual comovements. Uncertainty is one such shock. In our model, a change in uncertainty drives consumption and investment in opposite directions. If the effect is strong enough, we lose the strong procyclicality of these variables that we see in the data. We find, however, that in models with quasi-realistic parameter values, the productivity shocks dominate. In fact, the modest decline in the correlations of consumption and investment with output brings the model closer to the evidence.

Does uncertainty increase the persistence of economic fluctuations? The answer is no, but it's helpful to describe precisely how uncertainty affects decisions. We compute properties of model economies with accurate numerical procedures, but we get sharper insights from loglinear approximations analogous to Campbell's (1994). We think the approximations give us clarity about how uncertainty works that would be hard to come by otherwise. And despite rumors to the contrary, such loglinear approximations are compatible with uncertainty. These loglinear approximations have two striking properties:

- Tallarini property. Tallarini (2000) showed that in a model with an intertemporal elasticity of substitution (IES) of one, iid productivity growth, and constant uncertainty, the behavior of quantities is the same in models with recursive and additive preferences. We extend his result to a model with arbitrary IES and arbitrary linear dynamics in productivity growth: the behavior of quantities is affected by the IES but not uncertainty aversion. The result applies to loglinear approximations, but we find the approximations hard to distinguish from more accurate solutions.

- Separation property. The impact of uncertainty on dynamics is limited by what we call the *separation property*: the internal dynamics of the capital stock are independent of uncertainty and its properties. More precisely, the response of consumption and next period's capital stock to today's capital stock is independent of the shocks and their properties, including shocks to uncertainty. This is a standard feature of linear-quadratic models. It also applies to our loglinear approximations and, to a close approximation, to accurate numerical solutions of our models. One consequence is that uncertainty cannot account for an unusually slow recovery in this model except through persistence in the shock.

What's the bottom line? The recursive business cycle model provides a mechanism through which fluctuations in aggregate uncertainty affect the dynamics of aggregate quantities.

This mechanism changes the magnitudes of, and correlations between, growth rates of output, consumption, and investment. The effects, however, are small. We also find that uncertainty has essentially no impact on the internal dynamics of the model: an increase in uncertainty produces a more persistent decline in (say) consumption only if uncertainty is itself persistent. We conclude not that uncertainty is irrelevant to business cycles, but that any mechanism that produces a greater impact must operate through other channels.

A few words on notation: We use a number of conventions to keep it as simple as we can. (i) For the most part, Greek letters are parameters and Latin letters are variables or coefficients. (ii) We use t subscripts (x_t, for example) to represent random variables and the same letters without subscripts (x) to represent their means. Or, more commonly, $\log x$ represents the mean of $\log x_t$ rather than the log of the mean of x_t. (iii) We also use t subscripts to denote dependence of a function on the state. Thus $f(x_t)$ might be denoted f_t. (iv) We use variable subscripts to denote derivatives; for example, $f_{xt} = \partial f(x_t)/\partial x_t$. (v) The abbreviation iid means independent and identically distributed and $\text{NID}(a, b)$ means normally and independently distributed with mean a and variance b.

2 Risk

We approach uncertainty from the perspective of decision theory. We use the term *risk* to describe random environments in which the distribution of outcomes is known. We use *ambiguity* to describe environments in which some aspect of the distribution is unknown. *Uncertainty* is an umbrella term that includes both risk and ambiguity. We consider risk here and turn to ambiguity in Section 6.

2.1 Risk preference in static environments

Our treatment of risk is standard in macroeconomics and finance: the distribution over outcomes is known (risk) and equal to the distribution that generates the data (rational expectations).

To make this concrete, consider a static environment with a random state s and consumption $c(s)$ defined over it. Risk is a known nonconstant probability distribution over s, which induces a distribution over c. We summarize attitude toward risk with a certainty equivalent function, which transforms utility back into consumption units. More formally, the certainty equivalent $\mu(c)$ is the level of constant consumption that delivers the same utility. If c is constant, then $\mu(c) = c$. If c is risky, then risk aversion is indicated by $\mu(c) < E(c)$. Some common examples are described in Backus, Routledge, and Zin (2005, Section 3). We refer to the log difference $\log E(c) - \log \mu(c) > 0$ as a *risk adjustment*.

We rely exclusively on the expected utility certainty equivalent,

$$\mu(c) \quad = \quad u^{-1}\left[Eu(c)\right],$$

for some increasing concave function u. The standard example in macroeconomics and finance is power utility, $u(c) = c^\alpha/\alpha$, which implies the certainty equivalent

$$\mu(c) \quad = \quad \left[E(c^\alpha)\right]^{1/\alpha}. \tag{1}$$

Here $\alpha < 1$ and $1-\alpha > 0$ is commonly referred to as the coefficient of relative risk aversion (CRRA).

Most of our models are at least approximately loglinear with normal (Gaussian) uncertainty. Neither is essential, but the combination gives us relatively simple certainty equivalents. Suppose, for example, $\log c = s \sim \mathcal{N}(\kappa_1, \kappa_2)$ (the log of consumption is normal with mean κ_1 and variance κ_2). The log of the moment generating function for s is $\log E(e^{\theta s}) = \log E(c^\theta) = \theta\kappa_1 + \theta^2\kappa_2/2$. Therefore $\log E(c) = \kappa_1 + \kappa_2/2$, $\log E(c^\alpha) = \alpha\kappa_1 + \alpha^2\kappa_2/2$, and $\log \mu(c) = \kappa_1 + \alpha\kappa_2/2$. Risk aversion is implied by the risk adjustment $\log E(c) - \log \mu(c) = (1-\alpha)\kappa_2/2 > 0$.

2.2 Risk preference in dynamic environments

We follow standard practice and extend risk preference to dynamic environments with the recursive technology developed by Kreps and Porteus (1978). Utility U_t from date t on has the form

$$U_t \quad = \quad V[c_t, \mu_t(U_{t+1})]. \tag{2}$$

Time preference is built into the *time aggregator* V and risk preference is built into the certainty equivalent μ_t. The notation is intended to imply that c_t and U_t are functions of s_t, the state at date t. The certainty equivalent is computed from the conditional distribution of future states s_{t+1} given current state s_t.

We assume throughout that the time aggregator V and certainty equivalent μ_t are homogeneous of degree one, which allows us to use them in environments that are stationary in growth rates. We use the constant elasticity time aggregator suggested by Epstein and Zin (1989),

$$V[c_t, \mu_t(U_{t+1})] \quad = \quad \left[(1-\beta)c_t^\rho + \beta\mu_t(U_{t+1})^\rho\right]^{1/\rho}, \tag{3}$$

with $0 < \beta < 1$ and $\rho < 1$. Here $\sigma = 1/(1-\rho)$ is the intertemporal elasticity of substitution or IES: the elasticity of substitution between current consumption and the certainty equivalent of future utility. If we use the power certainty equivalent function (1), the coefficient of relative risk aversion is again $1-\alpha$, but the risk in this case is to future utility.

The time aggregator (3), like a certainty equivalent, expresses utility in consumption units. Consider a constant consumption path $c_{t+j} = c$ for all $j \geq 0$. Then $U_t = U_{t+1} = \mu_t(U_{t+1}) = c$. Differences of U_t from current consumption c_t reflect some combination of timing and uncertainty in the path of future consumption.

3 A recursive business cycle model

We imbed these preferences in a business cycle model, a streamlined version of Kydland and Prescott (1982). The model is conveniently summarized by a Bellman equation. We describe the Bellman equation and show how it can be scaled to take into account growth in productivity. Although risk premiums are not our focus, they provide a useful link to related work in finance. We show how asset returns are related to a pricing kernel, which in this model is the representative agent's intertemporal marginal rate of substitution. We also show how the maximum risk premium can be computed from the entropy of the pricing kernel.

3.1 Model

Our benchmark model starts with recursive preferences: equation (2) with time aggregator V and certainty equivalent function μ_t both homogeneous of degree one. Production uses capital (k_t) and labor (n_t) inputs and leads to the law of motion

$$k_{t+1} \;=\; f(k_t, a_t n_t) - c_t, \tag{4}$$

where f is also homogeneous of degree one and a_t is (labor) productivity. We fix labor supply at one ($n_t = 1$) and use a constant elasticity production function with constant depreciation:

$$f(k_t, a_t n_t) \;=\; [\omega k_t^\nu + (1-\omega)(a_t n_t)^\nu]^{1/\nu} + (1-\delta)k_t \;=\; y_t + (1-\delta)k_t, \tag{5}$$

where $\nu < 1$, $1/(1-\nu)$ is the elasticity of substitution between capital and labor, $0 < \delta \leq 1$ is the depreciation rate, and y_t is output. Investment is $i_t = y_t - c_t$.

The source of fluctuations in this model is a stochastic process for productivity growth. We refer to its components as "news" and "risk." Productivity growth $g_t = a_t/a_{t-1}$ is tied to a state vector x_t by $\log g_t = \log g + e^\top x_t$, where e is an arbitrary vector of coefficients. The vector x_t has linear dynamics,

$$x_{t+1} \;=\; Ax_t + Bv_t^{1/2} w_{1t+1}, \tag{6}$$

with $\{w_{1t}\} \sim \mathrm{NID}(0, I)$. If $A = [0]$ then the conditional mean of $\log g_t$ is constant. Otherwise Ax_t adds a predictable component of future productivity: in a word, news.

Risk is represented by stochastic volatility. The conditional variance v_t (loosely speaking, "volatility") is a linear first-order autoregression,

$$v_{t+1} \quad = \quad (1-\varphi_v)v + \varphi_v v_t + \tau w_{2t+1}, \tag{7}$$

with $0 < \varphi_v < 1$ and $\{w_{2t}\} \sim \text{NID}(0,1)$ and independent of $\{w_{1t}\}$. This gives positive probability to negative conditional variances, but we choose parameter values that make this probability tiny.

We find a competitive equilibrium as the solution to a planning problem: maximize utility subject to the laws of motion for the state $s_t = (k_t, a_t, x_t, v_t)$. The associated Bellman equation is

$$J(k_t, a_t, x_t, v_t) \quad = \quad \max_{c_t} \; V\big\{c, \mu_t[J(k_{t+1}, a_{t+1}, x_{t+1}, v_{t+1})]\big\}$$

subject to the laws of motion (4), $a_{t+1} = a_t g_{t+1} = a_t \exp(\log g + e^\top x_{t+1})$, (6), and (7).

Since V, μ, and f are all homogeneous of degree one, J is homogeneous of degree one in k_t and a_t. We can, therefore, divide the Bellman equation by a_t and express the problem in terms of scaled variables, $\tilde{k}_t = k_t/a_t$ and $\tilde{c}_t = c_t/a_t$. The scaled Bellman equation is

$$J(\tilde{k}_t, 1, x_t, v_t) \quad = \quad \max_{\tilde{c}_t} \; V\big\{\tilde{c}, \mu_t[g_{t+1} J(\tilde{k}_{t+1}, 1, x_{t+1}, v_{t+1})]\big\} \tag{8}$$

subject to the laws of motion. Similar logic gives us a scaled law of motion for $\tilde{k}_t$,

$$g_{t+1}\tilde{k}_{t+1} \quad = \quad f(\tilde{k}_t, 1) - \tilde{c}_t. \tag{9}$$

From here on, we drop the 1 and write the value function as $J(\tilde{k}_t, x_t, v_t)$.

We find it convenient to work with the log of J. With the constant elasticity time aggregator (3), we can rewrite (8) as

$$\log J(\tilde{k}_t, x_t, v_t) \quad = \quad \max_{\tilde{c}_t} \; \rho^{-1} \log\big\{(1-\beta)\tilde{c}^\rho + \beta\mu_t[g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})]^\rho\big\}. \tag{10}$$

In the limiting case of $\rho = 0$ — and intertemporal elasticity of substitution $\sigma = 1/(1-\rho) = 1$ — we have

$$\log J(\tilde{k}_t, x_t, v_t) \quad = \quad \max_{\tilde{c}_t} \; (1-\beta)\log\tilde{c} + \beta \log \mu_t[g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})].$$

Additive models generally work instead with J^ρ/ρ,

$$J(\tilde{k}_t, x_t, v_t)^\rho/\rho \quad = \quad \max_{\tilde{c}_t} \; (1-\beta)\tilde{c}^\rho/\rho + \beta\mu_t[g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})]^\rho/\rho,$$

which also follows from (8). If we redefine the value function as $\widehat{J}_t = J_t^\rho/\rho$ and set $\alpha = \rho$ (the additive case), the second term becomes $\mu_t(g_{t+1}J_{t+1})^\rho/\rho = E_t(g_{t+1}^\rho \widehat{J}_t)$. We use this in Appendix C.

3.2 Asset pricing fundamentals

The properties of asset returns in our model depend, in large part, on the marginal rate of substitution of the agent. In general no-arbitrage environments, there exists a positive *pricing kernel* m_{t+1} that satisfies the asset pricing relation

$$E_t\big(m_{t+1} r_{t+1}\big) \quad = \quad 1$$

for gross returns r_{t+1} on all assets. In representative agent models such as ours, the pricing kernel is the marginal rate of substitution and the equation is one of the agent's first-order conditions. With recursive preferences, the marginal rate of substitution is

$$m_{t+1} \quad = \quad \beta \left(\frac{c_{t+1}}{c_t}\right)^{\rho-1} \left(\frac{U_{t+1}}{\mu_t(U_{t+1})}\right)^{\alpha-\rho}.$$

See Appendix A. The first term is familiar from additive power utility. The next one is what we might term the Epstein-Zin term; it reflects the non-additivity of preferences over time and across states. It disappears if we set $\rho = \alpha$, which sends us back to the additive case.

A solution of our model includes a stochastic process for consumption that we can use to compute asset prices and, more commonly, risk premiums such as the equity premium. We use the maximum risk premium as an indicator of how risk is priced. The maximum risk premium is related to the entropy of the pricing kernel. The return on a one-period (riskfree) bond is $r^1_{t+1} = 1/E_t(m_{t+1})$. The return r_{t+1} with the highest log expectation $E_t(\log r_{t+1})$ is $r_{t+1} = 1/m_{t+1}$. The (logarithmic) risk premium on an arbitrary asset with return r_{t+1} is $E_t(\log r_{t+1} - \log r^1_{t+1})$. The maximum risk premium follows from the maximum expected return:

$$E_t(\log r_{t+1} - \log r^1_{t+1}) \quad \leq \quad L_t(m_{t+1}),$$

where $L_t(m_{t+1}) = \log E_t(m_{t+1}) - E_t(\log m_{t+1})$ is the conditional entropy of the pricing kernel. See Alvarez and Jermann (2005, Proposition 2) and Backus, Chernov, and Zin (2014, Section I). Taking expectations of both sides gives us the entropy bound:

$$E(\log r_{t+1} - \log r^1_{t+1}) \quad \leq \quad E[L_t(m_{t+1})]. \tag{11}$$

The object on the right is the largest risk premium that can be generated by this pricing kernel.

4 Risk in the recursive Brock-Mirman example

Now that we have a handle on risk, we can explore its impact on the properties of business cycle models. We start with a textbook standard, to which we add a number of bells and whistles. None of the bells and whistles affect the decision rule for consumption, so you might ask why we bothered. We do it to illustrate some general properties in an example we can solve with pen and paper.

4.1 Model and solution

What's come to be known as the Brock-Mirman example shows up in most introductions to dynamic programming for economists, including Ljunqvist and Sargent (2000, Chapter 4, Appendix B). In the notation of the previous section, it consists of additive preferences ($\rho = \alpha$), log utility ($\rho = 0$, corresponding to IES $\sigma = 1$), Cobb-Douglas technology ($f(a,k) = k^{\omega} a^{1-\omega}$ with $0 \leq \omega < 1$), and one hundred percent depreciation ($\delta = 1$). With this structure, the value function and decision rule are loglinear. See, for example, Appendix B. The same is true of the recursive version, where we allow the risk aversion parameter α to take on nonzero values in the certainty equivalent function (1).

The Bellman equation for this recursive Brock-Mirman example is

$$\log J(\tilde{k}_t, x_t, v_t) \quad = \quad \max_{\tilde{c}_t} \, (1-\beta) \log \tilde{c} + \beta \log \mu_t [g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})]$$

subject to (9), (6), and (7). Equation (9) in this case is $g_{t+1} \tilde{k}_{t+1} = f(\tilde{k}_t, 1) - \tilde{c}_t = \tilde{k}_t^{\omega} - \tilde{c}_t$. When $\omega = 0$, we have the Bansal-Yaron (2004) asset pricing model with $\rho = 0$.

We find the solution by guess and verify. We guess the value function is loglinear:

$$\log J(\tilde{k}_t, x_t, v_t) \quad = \quad p_0 + p_k \log \tilde{k}_t + p_x^{\top} x_t + p_v v_t$$

with coefficients (p_0, p_k, p_x) to be determined. Next we substitute the laws of motion into next period's value function and evaluate the certainty equivalent:

$$\begin{aligned} \log[g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})] \quad &= \quad (1-p_k)(\log g + e^{\top} x_{t+1}) + p_0 + p_k \log(\tilde{k}_t^{\omega} - \tilde{c}_t) + p_x^{\top} x_{t+1} + p_v v_{t+1} \\ &= \quad p_0 + (1-p_k) \log g + p_v(1-\varphi_v) v + p_k \log(\tilde{k}_t^{\omega} - \tilde{c}_t) \\ &\quad\quad + [(1-p_k)e + p_x]^{\top} (A x_t + B v_t^{1/2} w_{1t+1}) + p_v[\varphi_v v_t + \tau^{1/2} w_{2t+1}] \\ \log \mu_t [g_{t+1} J(\tilde{k}_{t+1}, x_{t+1}, v_{t+1})] \quad &= \quad p_0 + (1-p_k) \log g + p_v(1-\varphi_v) v + p_k \log(\tilde{k}_t^{\omega} - \tilde{c}_t) \\ &\quad\quad + [(1-p_k)e + p_x]^{\top} A x_t + [\alpha V_x/2 + p_v \varphi_v] v_t + \alpha p_v^2 \tau/2, \end{aligned}$$

where $V_x = [(1-p_k)e + p_x]^{\top} B B^{\top} [(1-p_k)e + p_x]$. Note the risk adjustment $\alpha V_x/2$ in the certainty equivalent, it's central to how risk works in such models and shows how the risk parameter α magnifies its effect.

If we substitute the certainty equivalent into the Bellman equation, the first order condition

$$(1-\beta)/\tilde{c}_t \quad = \quad \beta p_k / (\tilde{k}_t^{\omega} - \tilde{c}_t)$$

gives us the decision rule

$$\tilde{c}_t \;=\; \{(1-\beta)/[\beta p_k + (1-\beta)]\}\tilde{k}_t^\omega.$$

Traditionally we would substitute this back into the Bellman equation and solve for p_k. Here there's a simpler method. The envelope condition for $\tilde{k}_t$ gives us the decision rule

$$\tilde{c}_t \;=\; (1-\beta\omega)\tilde{k}_t^\omega.$$

The controlled law of motion is therefore $g_{t+1}\tilde{k}_{t+1} = \tilde{k}_t^\omega - \tilde{c}_t = \beta\omega\tilde{k}_t^\omega$.

If we substitute the decision rule into the Bellman equation and line up terms, we find

$$\begin{aligned}
p_k &= (1-\beta)\omega/(1-\beta\omega) \\
p_x^\top &= \beta(1-p_k)e^\top A(I-\beta A)^{-1} \\
p_v &= \beta(\alpha/2)V_x/(1-\beta\varphi_v) \\
(1-\beta)p_0 &= (1-\beta)\log(1-\beta\omega) + \beta(1-p_k)\log g + \beta p_v(1-\varphi_v)v + \beta p_k\log(\beta\omega) + \beta\alpha(p_v)^2\tau/2.
\end{aligned}$$

The coefficient p_k is between zero and one, which makes $J(\tilde{k}_t, x_t, v_t)$ increasing and concave in $\tilde{k}_t$. The coefficient p_x captures the predictability of log productivity growth: if $A = 0$, so that productivity growth is unpredictable, then $p_x = 0$ as well. And if $\alpha < 0$, as we'll typically assume, then $p_v < 0$: an increase in v_t lowers utility.

4.2 Qualitative properties

This model has a number of features of general interest. Among them:

- Loglinear solution. If we were searching for a loglinear decision rule, something like

  $$\log\tilde{c}_t \;=\; h_{c0} + h_{ck}\log\tilde{k}_t + h_{cx}^\top x_t + h_{cv}v_t, \tag{12}$$

  then we've found it: $h_{c0} = \log(1-\beta\omega)$ and $h_{cv} = \omega$. Note that $h_{cx} = h_{cv} = 0$: neither news x_t nor risk v_t affects scaled consumption. The controlled law of motion is similar:

  $$\log\tilde{k}_{t+1} \;=\; \log(\beta\omega) + \omega\log\tilde{k}_t - \log g_{t+1}. \tag{13}$$

  Loglinearity allows us to define the steady state in a simple and useful way. We use the term steady state to refer to the mean log of a variable. The steady state value of $\tilde{k}$, for example, is $\log\tilde{k} = E(\log\tilde{k}_t) = \log(\beta\omega/g)/(1-\omega)$. That implies a steady state capital-output ratio of $\log\tilde{k} - \log\tilde{y} = (1-\omega)\log\tilde{k} = \log(\beta\omega/g)$. This last expression contains a result we use later: the discount factor β can be used to adjust the steady state capital-output ratio.

- Tallarini property. It's evident from the decision rule and controlled law of motion that quantities don't depend on risk aversion. This is an illustration of Tallarini's (2000) result: to a loglinear approximation, the decision rules, and therefore the properties of quantities, are approximately the same with additive ($\alpha = \rho$) and recursive (arbitrary α) preferences. Here the result is exact. The risk aversion parameter α affects the value function (see p_0), but not the dynamics of consumption, capital (=investment), or output. In this respect the extension to recursive preferences is a waste of time.

- Separation property. The Brock-Mirman example, like most dynamic programs, includes both endogenous ($\tilde{k}_t$) and exogenous (x_t, v_t) state variables. Anderson, Hansen, McGrattan, and Sargent (1996, Sections 2 and 3) show that in analogous linear-quadratic control problems, the components of the solution separate: the coefficients of the endogenous state variables in the value function and decision rules do not depend on the properties of the exogenous state variables and can be computed separately.

  This problem has a similar feature. The coefficients $p_k = (1-\beta)\omega/(1-\beta\omega)$ in the value function and $h_{ck} = \omega$ in the decision rule do not depend on the parameters governing the dynamics of (x_t, v_t). Similarly, the slope of the controlled law of motion (13) is independent of the shocks. If we plot $\log \tilde{k}_{t+1}$ against $\log \tilde{k}_t$ for given values of (x_t, v_t, g_{t+1}), the separation property tells us that slope is the same for all values of these other variables. That's clearly the case here, where the slope is ω.

  Now think about this in the context of our problem. If we want to generate a slow recovery with risk, we have only two options: make risk more persistent or break the separation property.

- Asset prices. This is in Tallarini (2000), too: although quantities do not depend on α, asset prices and risk premiums do. In this case the pricing kernel has these components:

$$
\begin{aligned}
\log(\tilde{c}_{t+1}/\tilde{c}_t) &= \omega \log(\beta\omega) + \omega(\omega-1)\log \tilde{k}_t - \omega \log g_{t+1} \\
\log(c_{t+1}/c_t) &= \log(\tilde{c}_{t+1}/\tilde{c}_t) + \log g_{t+1} \\
&= \omega \log(\beta\omega) + (1-\omega)\log g + \omega(\omega-1)\log \tilde{k}_t \\
&\quad + (1-\omega)e^\top A x_t + (1-\omega)e^\top B v_t^{1/2} w_{1t+1} \\
\log[g_{t+1}J_{t+1}/\mu_t(g_{t+1}J_{t+1})] &= [(1-p_k)e + p_x]^\top B v_t^{1/2} w_{1t+1} + p_v \tau^{1/2} w_{2t+1} \\
&\quad - (\alpha/2) V_x v_t - (\alpha/2)(p_v)^2 \tau.
\end{aligned}
$$

  The pricing kernel is therefore

$$
\begin{aligned}
\log m_{t+1} &= \log \beta - \log(c_{t+1}/c_t) + \alpha \log[g_{t+1}J_{t+1}/\mu_t(g_{t+1}J_{t+1})] \\
&= \log\beta - \omega\log(\beta\omega) - (1-\omega)\log g - \alpha(\alpha/2)(p_v)^2\tau \\
&\quad - \omega(\omega-1)\log\tilde{k}_t - (1-\omega)e^\top A x_t - \alpha(\alpha/2)V_x v_t \\
&\quad + \left\{\alpha[(1-p_k)e + p_x] - (1-\omega)e\right\}^\top B v_t^{1/2} w_{1t+1} + \alpha p_v \tau^{1/2} w_{2t+1}.
\end{aligned}
$$

Conditional on the state at date t, most of this is constant. All the variation comes from the last two terms.

The entropy bound (11) gives us the maximum risk premium:

$$L_t(m_{t,t+1}) \;\; = \;\; (1/2)[V_m v_t + (\alpha p_v)^2 \tau],$$

with

$$V_m \;\; = \;\; \big\{\alpha[(1-p_k)e + p_x] + (1-\omega)e\big\}^\top BB^\top \big\{\alpha[(1-p_k)e + p_x] - (1-\omega)e\big\}.$$

Both terms are affected by α. In a typical case, if we increase risk aversion (meaning larger negative values of α) risk premiums go up, even though the dynamics of consumption are the same. This mirrors Tallarini, where risk premiums are affected by risk aversion, but consumption is not.

Returning to the subject of the paper: introducing stochastic volatility to this economy has no impact on the decision rules for consumption or (implicitly) investment or on the speed at which the capital stock returns to its steady state value. It's not all that helpful, then, in giving us a mechanism through which risk can affect either the magnitude or persistence of macroeconomic fluctuations.

5 Risk in the recursive business cycle model

The Brock-Mirman example is illustrative, but its simplicity is misleading. If we change the technology or allow the IES to differ from one, the role of uncertainty changes and the model isn't nearly as tractable. It is, however, solvable by loglinear approximation methods not much different from Campbell's (1994). It can also be solved, of course, by any number of numerical methods, but a loglinear approximation has the advantage of transparency: we can see exactly how it works and which features determine its properties. We do both: describe the properties of a loglinear approximation and quantify the model's properties with a more accurate numerical solution.

5.1 Model and solution

The recursive business cycle model consists of the Bellman equation (10), the certainty equivalent (1), and the laws of motion (9), (6), and (7). The first-order and envelope conditions are

$$\begin{aligned} 0 &= J_t^{-\rho}\big\{(1-\beta)\tilde{c}_t^{\rho-1} - \beta\mu_t(g_{t+1}J_{t+1})^{\rho-\alpha}E_t[(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}]\big\} \\ J_{kt}/J_t &= J_t^{-\rho}\beta\mu_t(g_{t+1}J_{t+1})^{\rho-\alpha}E_t[(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}]f_{kt}. \end{aligned} \tag{14}$$

Together they imply

$$(1-\beta)\tilde{c}_t^{\rho-1} \quad = \quad J_t^{\rho-1} J_{kt}/f_{kt}. \tag{15}$$

This is similar to what we'd get in the additive model. The resemblance is closer if we transform the value function, as we did at the end of Section 3.1. If we define the tranformed value function by $\widehat{J}_t = J_t^\rho/\rho$, its derivative is $\widehat{J}_{kt} = J_t^{\rho-1} J_{kt}$. We see, then, that the decision rule — the solution of (15) — depends on the derivative of the (transformed) value function but not on the value function itself. That's a general feature of additive dynamic programs with continuous control variables. It's not true of the recursive model, where the first-order and envelope conditions involve the value function as well as its derivative.

We gain some insight into the model from computing loglinear appproximations to the solution. The idea is to take functions that are not loglinear and nevertheless approximate them by loglinear functions. We'll see that this works amazingly well for models of the kind studied here. The approximations involve derivatives at a point, which in our case is the steady state, defined earlier as the mean of the log.

As an illustration, consider an arbitrary positive function f of a positive random variable x_t. A linear approximation in logs around the point $x_t = x$ is

$$\log f(x_t) \quad = \quad \log f + (f_x x/f)(\log x_t - \log x).$$

Typically we ignore the intercept and write this as $\log f(x_t) = (f_x x/f) \log x_t$. Similarly, with two variables we have $\log f(x_t, y_t) = (f_x x/f) \log x_t + (f_y y/f) \log y_t$. An example we use repeatedly is the marginal product of capital $f_{kt} = f_k(\tilde{k}_t, 1)$, which we approximate by

$$\log f_{kt} \quad = \quad (f_{kk}\tilde{k}/f_k) \log \tilde{k}_t \quad = \quad \lambda_r \log \tilde{k}_t. \tag{16}$$

This is the gross return at date t on one unit of capital invested at $t-1$. Another is the law of motion (4), which we approximate by

$$\begin{aligned} \log \tilde{k}_{t+1} \quad &= \quad (f_k/g) \log \tilde{k}_t - (\tilde{c}/\tilde{k}g) \log \tilde{c}_t - e^\top x_{t+1} \\ &= \quad \lambda_k \log \tilde{k}_t - \lambda_c \log \tilde{c}_t - e^\top x_{t+1}. \end{aligned} \tag{17}$$

The notation and approach should be familiar from Campbell (1994).

We apply similar methods to our dynamic programming problem and derive a loglinear approximation to the consumption decision rule and controlled law of motion for capital. In the additive case, we reproduce Campbell's approximation; see Appendix C. We derive an analogous approximation for the recursive case in Appendix D. The result is a decision rule of the form (12) for consumption and a similar approximation to the controlled law of motion,

$$\log \tilde{k}_{t+1} \quad = \quad h_{k0} + h_{kk} \log \tilde{k}_t + h_{kx}^\top x_t + h_{kv} v_t - \log g_{t+1}. \tag{18}$$

The primary difference from (13) is the possibility of nonzero values for h_{kx} and h_{kv}.

5.2 Qualitative properties

We leave the calculations to Appendix D, but summarize the features of the loglinear approximation in a proposition:

Proposition. *If we hold constant the steady state capital-output ratio, the loglinear approximation to the solution is block triangular:*

(a) The coefficients (h_{ck}, h_{kk}) in (12,18) governing dependence of the decision rule and controlled law of motion on the current capital stock are independent of the risk aversion parameter α and of the properties of the shocks (x_t, v_t) to news and risk.

(b) The analogous coefficients (h_{cx}, h_{kx}) of the news shock x_t are independent of the risk parameter α and of the properties of the shock v_t to risk.

To see how the loglinear approximation works, consider the equations for capital. We make the loglinear guess

$$\log(J_t^{\rho-1} J_{k,t}) \quad = \quad q_k \log \tilde{k}_t + q_x^\top x_t + q_v v_t$$

for coefficients (q_k, q_x, q_v) to be determined. Then (15) and (17) give us

$$\begin{aligned} h_{ck} &= -\sigma(q_k - \lambda_r) \\ h_{kk} &= \lambda_k + \sigma\lambda_c(q_k - \lambda_r). \end{aligned}$$

These equations and the combined first-order and envelope condition (15) imply the Riccati-like equation for q_k:

$$q_k \quad = \quad q_k[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)] + \lambda_r.$$

As we note in the proposition, none of this depends on risk aversion or the properties of the shocks.

Thus we have:

- Tallarini property. Without variation in risk — that is, with $\tau = 0$ in equation (7) — the loglinear approximations of the decision rule and controlled law of motion are identical in the additive and recursive models. Recursive preferences are irrelevant here to the behavior of quantities. This generalizes Tallarini's result to economies with arbitrary IES and arbitrary linear dynamics for log productivity growth. Similar examples run throughout Hansen and Sargent (2007).

- Separation property. The internal dynamics of the capital stock — namely, the coefficient h_{kk} in (18) — are independent of the shocks. Similarly, the dynamics induced by news — namely, h_{kx} — are independent of risk v_t and risk aversion α. The relevant equations here are

$$\begin{aligned} h_x^\top &= -\sigma q_x^\top \\ q_x^\top &= -(\sigma^{-1} + q_k) e^\top A \big[(1 - \sigma q_k \lambda_c) I - A \big]^{-1}. \end{aligned}$$

  Not only are recursive preferences irrelevant here, but there's no impact of risk on the decision rule other than the direct one.

Risk does matter in this model, but only through its direct impact on consumption, summarized by the coefficients h_{cv} and h_{kv}. We turn next to the magnitude of this impact.

5.3 Quantitative properties

We can get a sense of the magnitude of the impact of risk if we choose specific parameter values. We compute numerical examples, rather than carefully calibrated models, but most of the parameter values have a history in the literature. We report them in Table 1 without further comment. We then compute properties of the model using a piecewise linear approximation to the value function over a fine grid for the capital stock. Experiments with finer grids and higher-order splines suggest that our calculations are extremely accurate.

So what do we find? We find, first, that the loglinear approximation is very good. If we plot $\log \tilde{k}_{t+1}$ against $\log \tilde{k}_t$ over the range of values generated by the model, the solution produced by our numerical procedure is indistinguishable from the loglinear approximation (18). See Figure 5. The same is true of the Tallarini and separation properties: they're not exact, but they're very good approximations.

Now consider the impact of risk. In Figure 6 we show the responses of (the logarithms of) capital and consumption to a unit increase in the innovation w_{2t} in the conditional variance v_t. The effect is small. A two standard deviation (of v_t) increase would be roughly six times larger. We see that the unit increase reduces consumption by a small amount in the benchmark case; roughly 0.02. If we increase risk aversion to $1 - \alpha = 50$, this rises to 0.1. Multiplying by six, we would get a little more than a half percent decrease in consumption from a two standard deviation increase in risk. This is a noticeable effect, but it's not hard to see that the model needs help to generate an impact as large as we saw in 2008, or even a substantial fraction of it.

We report business cycle properties of various parameter choices in Table 2. The statistics summarized in the table also suggest that the impact of uncertainty fluctuations is small in this model. The first column of the table includes a summary of the evidence taken from Tallarini (2000, Table 6). The statistics are based on (continuously-compounded) growth rates rather than some kind of filtered object, but the properties are familiar: the standard deviation of consumption is smaller, and the standard deviation of investment larger, than that of output. And both consumption and investment are positively correlated with output, although the correlations are smaller than we would see with (say) Hodrick-Prescott filtered variables.

The properties of models are broadly similar to the data, although the differences in standard deviations are smaller in the models and the correlations are larger. More relevant to us is the role of uncertainty shocks. Column (3) summarizes the model with benchmark parameter values. Column (2) shows us that it's not much different from the additive case. Column (5) shows how it changes when we eliminate variation in risk: hardly at all. We get a larger difference when we increase risk aversion to $1-\alpha = 50$; column (4). As we have seen, this increases the impact of fluctuations in uncertainty, but even in this case the effect is modest. Changing the IES has a bigger effect on business cycle statistics, but uncertainty shocks have little to do with this. When we increase the IES to 1.5, as we do in column (6), the standard deviation of consumption growth falls. If this seems counterintuitive, see Kaltenbrunner and Lochstoer (2010, Figure 1B). The long-term effect of a shock on consumption is the same for all values of the IES, but the short-term effect is not.

We also see the impact of risk aversion on risk premiums noted by Tallarini. Here the maximum risk premium is the entropy bound (11), reported at the bottom of the table. When we increase risk aversion from $1 - \alpha = 2$ in column (2) to 50 in column (4), the maximum risk premium increases from 0.02 percent (per quarter) to 16 percent.

6 Ambiguity

We use the term ambiguity to describe situations in which the decision maker does not know some aspect of the distribution of outcomes. Whether it's a parameter or a state variable is in large part a matter of language. The critical aspect of ambiguity is that it's treated differently in preferences than risk. Our treatment of ambiguity is built on the *smooth ambiguity* foundation laid by Hayashi and Miao (2011), Jahan-Parvar and Liu (2012), Ju and Miao (2012), and especially Klibanoff, Marinacci, and Mukerji (2005, 2009). We find it more user-friendly than more popular approaches built on maxmin expected utility, although others may feel that it leaves out something essential.

6.1 Smooth ambiguity in static environments

Consider ambiguity in a static setting with two sources of uncertainty, $s = (s_1, s_2)$. Consumption outcomes are defined over them by $c(s) = c(s_1, s_2)$. The sources of uncertainty are, first, the distribution of s_1 conditional on s_2 and, second, the distribution over s_2. There's no difference between s_1 and s_2 at this level of generality, but in applications s_2 is often a parameter or a hidden state.

We denote the expectations based on these distributions by E_1 and E_2, respectively, and the overall expectation by $E = E_2 E_1$. More explicitly, the various expectations of an arbitrary function $f(s_1, s_2)$ might be expressed by

$$\begin{aligned} E_1[f(s_1,s_2)] &= E[f(s_1,s_2)|s_2] \\ E[f(s_1,s_2)] &= E_2\big(E[f(s_1,s_2)|s_2]\big) = E_2\big(E_1[f(s_1,s_2)]\big). \end{aligned}$$

The second line follows from the law of iterated expectations.

With *smooth ambiguity*, we allow different "smooth" preferences over these two sources of uncertainty. The certainty equivalent has two parts,

$$\mu(c) = \mu_2[\mu_1(c)],$$

where

$$\begin{aligned} \mu_1(c) &= u^{-1}\big[E_1 u(c)\big] \\ \mu_2[\mu_1(c)] &= v^{-1}\big(E_2 v[\mu_1(c)]\big). \end{aligned}$$

These functions exhibit risk aversion if u is concave and ambiguity aversion if $v \circ u^{-1}$ is concave — roughly speaking, if v is more concave than u. The power utility versions are

$$\mu_1(c) = \big[E_1(c^\alpha)\big]^{1/\alpha}, \quad \mu_2[\mu_1(c)] = \big(E_2[\mu_1(c)^\gamma]\big)^{1/\gamma} \tag{19}$$

with parameters $\alpha < 1$ and $\gamma \leq \alpha$. We refer to $1-\alpha$ as risk aversion and $1-\gamma$ as ambiguity aversion. The latter is slightly misleading, given that ambiguity aversion requires $1-\gamma > 1-\alpha$, but it's clearer than the alternatives that cross our minds. If $\gamma = \alpha$, (19) reduces to expected utility:

$$\mu(c) = \mu_2[\mu_1(c)] = \big[E_2 E_1(c^\alpha)\big]^{1/\alpha} = \big[E(c^\alpha)\big]^{1/\alpha}.$$

Alternatively, if we drive γ to minus infinity we get the popular maxmin expected utility. For values of γ between minus infinity and α, the smooth ambiguity model captures the idea of model uncertainty in a user-friendly way.

Example (continued). We illustrate the impact of ambiguity aversion in two variants of our example from Section 2.1. (i) Ambiguous mean. We express ambiguity over the mean with a two-part distribution. Part 1: Conditional on s_2, $\log c = s_1 \sim \mathcal{N}(s_2, \kappa_2)$. Part 2: $s_2 \sim \mathcal{N}(\kappa_1, \nu)$. The first certainty equivalent is $\log \mu_1(c) = s_2 + \alpha\kappa_2/2$. The overall certainty equivalent is $\log \mu(c) = \log \mu_2[\mu_1(c)] = \kappa_1 + \gamma\nu/2 + \alpha\kappa_2/2$. The mean satisfies $\log E(c) = \kappa_1 + \kappa_2 + \nu/2$, so the adjustment for risk is $(1-\alpha)\kappa_2/2$ and the adjustment for ambiguity is $(1-\gamma)\nu/2$. (ii) Ambiguous variance. Part 1: conditional on s_2, $\log c = s_1 \sim \mathcal{N}(\kappa_1, s_2)$. Part 2: $s_2 \sim \mathcal{N}(\nu_1, \nu_2)$. The variance is normal and therefore negative with positive probability, which is impossible but analytically convenient. The certainty equivalents are $\log \mu_1(c) = \kappa_1 + \alpha s_2/2$. and $\log \mu_2[\mu_1(c)] = \kappa_1 + \alpha\nu_1/2 + \gamma(\alpha/2)^2\nu_2/2$. Evidently there's no clean separation here between the adjustments for risk and ambiguity.

6.2 Smooth ambiguity in dynamic environments

We follow Hayashi and Miao (2011) and imbed preference toward risk and ambiguity in the traditional recursive utility setup summarized by equation (3).

The action is in the certainty equivalent function μ_t. The main issue in applications is how we distinguish between risk and ambiguity. We associate risk with components of the state we observe (s_{1t}) and ambiguity with components we do not observe (s_{2t}). Aversion to risk and ambiguity are built into the certainty equivalent functions μ_{1t} and μ_{2t}, respectively:

$$\mu_t(U_{t+1}) \quad = \quad \mu_{2t}[\mu_{1t}(U_{t+1})].$$

Here the subscript t refers to the information set at date t, typically the complete history of observable states s_{1t}. The inner certainty equivalent μ_{1t} is computed from the distribution over s_{1t+1} conditional on this history and on the unobserved state s_{2t+1}. The outer certainty equivalent μ_{2t} is the conditional distribution over s_{2t+1} given the same history. By assumption, we have stronger aversion to the latter than the former. We use the same power certainty equivalent functions we used in equation (19).

The marginal rate of substitution now includes an additional term:

$$m_{t+1} \quad = \quad \beta \left(\frac{c_{t+1}}{c_t}\right)^{\rho-1} \left(\frac{U_{t+1}}{\mu_{2t}[\mu_{1t}(U_{t+1})]}\right)^{\alpha-\rho} \left(\frac{\mu_{1t}(U_{t+1})}{\mu_{2t}[\mu_{1t}(U_{t+1})]}\right)^{\gamma-\alpha}. \qquad (20)$$

If $\gamma = \alpha$ the last term drops out, and if $\alpha = \rho$ as well we're back to additive utility.

7 Risk and ambiguity in a business cycle model

We construct three examples of business cycle models with both risk and ambiguity. The models are like the one in Section 5 except for the stochastic process for productivity growth and the certainty equivalent functions. All of them have ambiguity over the conditional mean of log productivity growth, but they differ in whether the amount of ambiguity varies over time. The structure and properties are similar to our model with risk. More specifically, their loglinear approximations have the same structure block triangular structure summarized in the proposition of Section 5.2. Nevertheless, we find ambiguity interesting enough to make it worth working through the details.

7.1 A model with constant volatility

The simplest example has ambiguity over the conditional mean. The process for (log) productivity growth changes from (6,7) to

$$\begin{aligned} \log g_{t+1} &= \log g + x_{t+1} + \tau_1 w_{1t+1} \\ x_{t+1} &= \varphi_x x_t + \tau_2 w_{2t+1}, \end{aligned}$$

with (w_{1t}, w_{2t}) independent iid standard normal random variables. The state here is $(s_{1t}, s_{2t}) = (\log g_t, x_t)$. The agent observes $\log g_t$ but not x_t, making $\tau_1 w_{1t}$ the forecast error of $\log g_t$ conditional on x_t. Therefore, in the first equation the agent treats w_{1t+1} as risky and x_{t+1} as ambiguous. Risk is constant. Ambiguity depends on what we can learn about x_{t+1} from observations of $\log g_t$.

This is a standard filtering problem with "state" x_t and "measurement" $\log g_t$. At each date t, given the history of productivity growth to that date, the agent's distribution over x_t is normal with (say) mean $\widehat{x}_t$ and variance b_t. Starting with arbitrary values $(\widehat{x}_0, b_0)$, we compute them recursively:

$$\begin{aligned} \theta_t &= b_t/(b_t + \tau_1^2) \\ b_{t+1} &= \tau_2^2 + \theta_t \varphi_x^2 \tau_1^2 \\ \widehat{x}_{t+1} &= \varphi_x (1-\theta_t)\widehat{x}_t + \varphi_x \theta_t (\log g_{t+1} - \log g). \end{aligned}$$

At the same date t, the distribution of x_{t+1} conditional on the history to date is normal with mean $\varphi_x \widehat{x}_t$ and variance b_{t+1}. Note, too, that the conditional variance is deterministic. It converges (rapidly in practice) to a constant, which is what we assume here. In short, this example generates ambiguity, reflected in b_{t+1}, but once we've converged the amount of ambiguity is constant.

Now that we've solved the filtering problem, we apply the two-part certainty equivalent. Risk preference applies to τ_1^2, the variance of $\log g_{t+1}$ conditional on (the unobserved) x_{t+1}. Ambiguity applies to b_{t+1}, the standard deviation of x_{t+1} conditional on the history of productivity growth. Other than the initial dynamics of b_t, risk and ambiguity are constant.

We can plug this into our business cycle model and solve as before. The full-information state here is $(\tilde{k}_t, x_t)$, but the effective state is $(\tilde{k}_t, \widehat{x}_t)$. The scaled Bellman equation is (8) with the appropriate modification of the state and certainty equivalent function. The first-order and envelope conditions become

$$\begin{aligned} (1-\beta)\tilde{c}_t^{\rho-1} &= J_t^{\rho-1} J_{kt} \\ J_t^{\rho-1} J_{kt} &= \beta \mu_{2t} (\mu_{1t})^{\rho-\gamma} E_{2t} \left\{ \mu_{1t} (g_{t+1} J_{t+1})^{\gamma-\alpha} E_{1t} \big[(g_{t+1} J_{t+1})^{\alpha-1} J_{kt+1} \big] \right\} f_{kt}. \quad (21) \end{aligned}$$

A loglinear approximation still satisfies the proposition of Section 5.2. And with risk and ambiguity constant, neither affects the dynamics of quantities. See Appendix E.

This example is helpful is showing how we might distinguish between risk and ambiguity, but since both are constant they contribute nothing to business cycles. The next step is to introduce variation in the quantity of ambiguity.

7.2 Models with stochastic volatility

Our next two examples have stochastic variation in the quantity of ambiguity (its conditional variance) and in this respect has a broadly similar flavor to Ilut and Schneider (2014).

In Model I, risk and ambiguity are stochastic. Productivity growth is

$$\begin{aligned} \log g_{t+1} &= \log g + \tau_1 v_t^{1/2} w_{1t+1} + \tau_2 v_t^{1/2} w_{2t+1} \\ v_{t+1} &= (1-\varphi_v) v + \varphi_v v_t + \tau_v w_{3t+1}, \end{aligned}$$

where (w_{1t}, w_{2t}, w_{3t}) are independent standard normal. We also set $\tau_1^2 + \tau_2^2 = 1$, a normalization that makes the unconditional variance of $\log g_t$ equal to v. Relative to the previous model, we have introduced stochastic volatility. We also eliminated dynamics in the conditional mean ($\varphi_x = 0$), so that $x_{t+1} = \tau_2 w_{2t+1}$. The agent observes $\log g_t$ and v_t, but not the components of productivity growth w_{1t} and w_{2t}, so that w_{1t} represents risk and w_{2t} represents ambiguity. The parameters τ_1 and τ_2 control their magnitudes. Since both are multiplied by $v_t^{1/2}$, risk and ambiguity vary together.

The model is similar to the previous one with state $(\tilde{k}_t, v_t)$. The first-order and envelope conditions don't change. We compute numerical solutions by other methods, but a loglinear approximation is again instructive. In particular, the solution still exhibits the separation property: the parameters governing the shocks and uncertainty aversion affect only the loading on volatility v_t. See Appendix E.

In Model II, only ambiguity is stochastic. Productivity growth is

$$\begin{aligned} \log g_{t+1} &= \log g + \tau_1 v^{1/2} w_{1t+1} + \tau_2 v_t^{1/2} w_{2t+1} \\ v_{t+1} &= (1-\varphi_v) v + \varphi_v v_t + \tau_v w_{3t+1}, \end{aligned}$$

where (w_{1t}, w_{2t}, w_{3t}) are independent standard normal and $\tau_1^2 + \tau_2^2 = 1$ (a normalization). The change is that risk is now constant Relative to the previous model, we have eliminated dynamics in the conditional mean ($\varphi_x = 0$) and introduced stochastic volatility. Let us say that the agent observes $\log g_t$ and v_t but not the components of productivity growth w_{1t} and w_{2t}. The idea is that w_{1t} represents risk and w_{2t} represents ambiguity. Since the latter is multiplied by v_t, ambiguity is no longer constant.

The loglinear approximation to the solution has the same form as the proposition. In particular, the coefficient h_{kk} is independent of risk and ambiguity aversion. The coefficient of the conditional variance, for example, is

$$q_v \;=\; q_k h_{kv} + q_v \varphi_v + [(\rho-\gamma)\gamma + (\gamma-1)^2]\tau_2^2/2.$$

We see here that we can make the response to variance shocks solely an artifact of ambiguity aversion through the parameter γ. See Appendix E.

Both of these models generate ambiguity that varies randomly through time. In this respect they're similar to earlier work by Ilut and Schneider (2014) and Jahan-Parvar and Liu (2012). Ilut and Schneider model ambiguity with a dynamic version of maxmin expected utility. They posit exogenous variation in the worst-case probability distribution, just as we posit exogenous variation in the conditional variance. Jahan-Parvar and Liu adopt an environment closer to our first example. They have a world that alternates between two "states" and solve a nonlinear filtering problem to determine the probabilities over the two states for each history. This mechanism generates endogenous variation in the conditional variance as the probabilities change.

We report quantitative properties of these two models in Table 3. Column (1) is again US data. Columns (2) to (5) refer to Model I. We see that the impact of risk and ambiguity on quantities is small regardless. We also see that while ambiguity has little impact on quantities, it has a larger impact on asset prices. The impact on quantities is slightly larger if we increase ambiguity aversion from $1-\gamma = 30$ to 50. Columns (6) and (7) refer to Model II, in which ambiguity varies but risk is constant. The impact is larger when ambiguity plays a greater role [Column (6)]. This reflects, in part, our choice of parameters, in which ambiguity is limited by the amount of the observed variance of productivity growth.

8 Discussion

We have described how aggregate uncertainty affects the dynamics of quantities in a traditional business cycle model. The short answer: The impact is small. A longer answer: Variations in uncertainty in this model have, for reasonable parameter values, a small effect on consumption, investment, and output. They have no effect on the internal dynamics of the capital stock: the speed at which the economy recovers from a temporary increase in uncertainty. Prolonged uncertainty could generate a prolonged recession on its own, but the evidence suggests that this hasn't been the case in the US. As a result, the model cannot account for either the magnitude or the persistence of the most recent recession with uncertainty.

If uncertainty does little in this model, could other mechanisms produce larger, more persistent effects? We review related work and go on to explore alternatives that we think could lead to a more central role of uncertainty in business cycle dynamics.

8.1 Related work

We are hardly the first to explore the role of aggregate uncertainty in business cycle dynamics and the methods used to compute solutions to such models. We summarize some of the most prominent contributions below, organized by topic.

Computation. We've noted the connection to Campbell's (1994) work; Kaltenbrunner and Lochstoer (2010) is similar. Malkhozov (2014) extends loglinear approximation to models with recursive preferences and linear ("affine") conditional variance processes. Dew-Becker (2012) is similar. Caldara, Fernandez-Villaverde, Rubio-Ramirez, and Wen (2012) epitomize a large body of work based on "perturbation methods." These methods, as they are commonly used in macroeconomics, start with first-order approximations in which uncertainty has been turned off. Uncertainty enters when we move on to higher-order approximations. It is possible, of course, to incorporate uncertainty in first-order approximations, as we have done here. The same approach is widely used in finance. We suggest that loglinear approximations can be a useful source of insight even if other methods are used to compute quantitative properties.

Uncertainty and business cycles. A number of papers incorporate time-varying uncertainty in similar models. Most of them use a normal autoregressive process for the log of the conditional variance, which makes the conditional variance lognormal rather than normal. Justiniano and Primiceri (2008) do this for all of their shocks, but decisions in their model do not respond to these measures of uncertainty. Caldara, Fernandez-Villaverde, Rubio-Ramirez, and Wen (2012) compare a number of methods of approximating the decision rules for a business cycle model with stochastic uncertainty. They plot, for their benchmark parameter values, decision rules for consumption and labor supply. They find that consumption falls with uncertainty and labor supply rises (leisure falls). They do not examine further the quantitative impact of uncertainty on business cycles. Bidder and Smith (2013) consider a robust control interpretation of recursive preferences and explore the impact of risk on the agent's implicit ("distorted") probabilities over future events. They do not report the model's business cycle properties.

Liu and Miao (2014), in a paper presented at this conference, consider a model with disappointment aversion in which the conditional mean and variance of productivity growth vary between high- and low-value states. The structure of uncertainty is different but its effects are the same: an increase in the conditional variance leads consumption to rise and investment to fall. It also leads to a decline in the price and return on equity. Overall, variation in uncertainty has only a small impact on the cyclical properties of the model.

Basu and Bundick (2012) add Keynesian features to an otherwise similar model: sticky prices, countercyclical markups, and shocks to current consumption in the time aggregator. They find, as we do, that the effects of uncertainty are small, but they're larger with these additional features. They also find that their model corrects some of the comovement problems associated with fluctuations in uncertainty. We solve the same problem simply by having productivity as the dominant driving process.

Ambiguity. The literature on ambiguity and business cycles includes notable contributions by Hansen and Sargent (2007), Ilut and Schneider (2014), and Jahan-Parvar and Liu (2012). Hansen and Sargent summarize an extensive line of research in which maxmin expected utility is extended to dynamic environments. One issue they address at length is what we've called the Tallarini property: the observational equivalence of additive and recursive preferences for the behavior of quantities. Ilut and Schneider use another version of dynamic maxmin expected utility. In their model, the conditional mean of productivity growth is bounded by an interval $[-a_t, a_t]$ for some $a_t > 0$ that varies with time. The min in maxmin has the effect of having the agent act as if the conditional mean is $-a_t$, a state variable that shows up in decision rules. In their model, variations in a_t account for most of the variation in output.

Jahan-Parvar and Liu use smooth ambiguity. Productivity growth is conditionally log-normal with a conditional mean that follows a two-state Markov chain. The state is not observed, but the agent computes a distribution over states from past realizations. The agent is ambiguous over this uncertainty, leading to a role for ambiguity aversion. Variation in the distribution over states leads to time-varying ambiguity. Their focus is on asset prices, but their reported business cycle properties suggest that quantities are not greatly affected by uncertainty.

8.2 Extensions

Several extensions strike us as having potential for increasing the impact of uncertainty. Among them:

Uncertainty in what? Many applications, including ours, place uncertainty in exogenous variables: in productivity growth, in shocks to current utility, and so on. In the language of traditional business cycle research, this gives us different impulses to the same propagation mechanisms. The effects are likely to be much different if the uncertainty is in an endogenous variable or a parameter. Suppose, for example, the share parameter in the production function is uncertain. The effects would run throughout the model. The separation property disappears, giving us the potential to generate more interesting dynamics than we have described above.

Learning. Several applications of ambiguity to asset pricing involve learning: the models have hidden states whose probabilities are inferred from other variables. This additional uncertainty about the state is a natural source of ambiguity. Several applications do exactly this in models with discrete state spaces, notably Jahan-Parvar and Liu (2012), Ju and Miao (2012), and Klibanoff, Marinacci, and Mukerji (2009). The same idea is easily applied to models with continuous state spaces. For example, a model based on a law of motion like (6) would serve this purpose if some or all of x_t is not observed. If volatility v_t is constant, uncertainty is constant, too. Otherwise it varies through time as in our examples.

The advantage of learning in such models is that it gives us another source of dynamics. If learning is slow enough, it can contribute in a significant way to the dynamics of the model. Collard, Mukerji, Sheppard, and Tallon (2012) and Collin-Dufresne, Johannes, and Lochstoer (2013) generate significant effects on asset prices. A logical next step is to study its effects on business cycles.

8.3 Alternative mechanisms

Endogenous uncertainty. A related line of thought is to make uncertainty endogenous. One version builds in feedback from economic activity to uncertainty. Since uncertainty is no longer exogenous, the model loses the separation property, leaving more room for an impact on aggregate dynamics. Striking examples of this include Fajgelbaum, Schaal, and Taschereau-Dumouchel (2014), Navarro (2014), Van Nieuwerburgh and Veldkamp (2006), and Veldkamp (2005). In all of them, a decline in economic activity produced by other means reduces the amount of information available and increases uncertainty. This reinforces the decline in output, and can prolong deep recessions or even make them (in the words of the first paper) "near permanent."

Fostel and Geanakoplos (2012, 2014) build multi-agent models with incomplete markets in which asset trades are limited by collateral. Several of their examples generate procyclical leverage and endogenous increases in volatility during downturns.

Microeconomic uncertainty. A number of influential papers explore the impact on business cycles of microeconomic uncertainty: of variation in the uncertainty faced by individual firms or households. Bloom, Floetotto, Jaimovich, Saporta, and Terry (2012) imbed an industry model in a business cycle framework. Firms face idiosyncratic shocks to their productivities whose variance follows a two-state Markov chain. Fixed costs reduce firms' willingness to adjust capital and labor inputs as firm-level productivities change. The striking result, also evident in earlier work, is that an increase in microeconomic uncertainty reduces aggregate productivity. The effect can be substantial, perhaps (they argue) three percent of GDP (Figure 6). Declines in labor and investment are larger still (Figure 7). The model does not, however, explain why the recovery has been so slow.

Bachmann and Bayer (2013) argue that the impact is likely to be smaller. They estimate a process for microeconomic uncertainty from German data and find that the variation in firm-level uncertainty is smaller than estimates used elsewhere. They also find that the overall impact of micro uncertainty on business cycle properties is small (Table 5).

Arellano, Bai, and Kehoe (2012) combine microeconomic uncertainty with financial frictions for firms. They note, as we do in Figure 4, that labor productivity was flat during the recession. They propose an explanation of the 2008 recession that operates through financial frictions and employment. In their model, firms must finance their wage bills in advance with noncontingent debt. If they are unable to repay the debt at the end of the period, they default and forego future profits. When uncertainty rises, the probability of default rises with it, and firms hire fewer workers and produce less output. In this way uncertainty generates a negative "labor wedge" that reduces employment and output (Figure 1).

9 Where next?

Even a casual look at the evidence suggests a connection between uncertainty and business cycles. Economic downturns are generally associated with increases in the uncertainty of financial returns and of near-term macroeconomic performance. They are also associated with increases in the dispersion of outcomes experienced by households and firms. The question is what mechanisms connect uncertainty and the economy. Like others before us, we looked at the impact of exogenous increases in uncertainty. Also like others, we've found that the effects are small. Given the evidence, that can't be the end of the story. Perhaps future work will give us more effective ways in which uncertainty influences economic behavior, or even generate such uncertainty as part of a more complete explanation of business cycles than we have now.

A Marginal rates of substitution

We derive the marginal rates of substitution for recursive preferences with risk and ambiguity. Consider an event tree with histories or states $s^t = (s_0, s_1, \ldots, s_t)$. We're interested in the marginal rate of substitution between consumption in state s^t and a succeeding state $s^{t+1} = (s^t, s)$. Since everything starts at s^t, we can ignore it in what follows. We use a finite state-space to simplify the notation.

Preferences are characterized by the time aggregator (3) and the risk and ambiguity certainty equivalents (19). Denote current utility by $U(s^t) = U_t$ and future utility by $U(s^t, s) = U_{t+1}(s) = U_{t+1}$. We divide s into (s_1, s_2) and consider probabilities $\pi(s_1, s_2) = \pi_1(s_1|s_2)\pi_2(s_2)$, all conditional on the current state s^t. The overall certainty equivalent is $\mu_t(U_{t+1}) = \mu_{2t}[\mu_{1t}(U_{t+1})]$. The inner one,

$$\mu_{1t}(U_{t+1}) \quad = \quad \Big[\sum_{s_1} \pi_1(s_1|s_2) U_{t+1}(s_1, s_2)^\alpha \Big]^{1/\alpha},$$

might be expressed $\mu_{1t}(s_2)$, a function of s_2. The outer one is

$$\mu_{2t}[\mu_{1t}(s_2)] \quad = \quad \Big[\sum_{s_2} \pi_2(s_2) \mu_{1t}(s_2)^\gamma \Big]^{1/\gamma}.$$

Marginal utilities follow from repeated application of the chain rule. The relevant derivatives are

$$\begin{aligned}
\partial U_t / \partial c_t &= U_t^{1-\rho}(1-\beta)c_t^{\rho-1} \\
\partial U_t / \partial \mu_{2t}[\mu_{1t}(s_2)] &= U_t^{1-\rho}\beta \mu_{2t}[\mu_{1t}(s_2)]^{\rho-1} \\
\partial \mu_{2t}[\mu_{1t}(s_2)] / \partial \mu_{1t}(s_2) &= \mu_{2t}[\mu_{1t}(s_2)]^{1-\gamma}\pi_2(s_2)\mu_{1t}(s_2)^{\gamma-1} \\
\partial \mu_{1t}(s_2) / \partial U_{t+1}(s_1, s_2) &= \mu_{1t}(s_2)^{1-\alpha}\pi(s_1|s_2)U_{t+1}(s_1, s_2)^{\alpha-1}.
\end{aligned}$$

The marginal rate of substitution is therefore

$$\begin{aligned}
\frac{\partial U_t / \partial c_{t+1}(s_1, s_2)}{\partial U_t / \partial c_t} &= \frac{[\partial U_t/\partial \mu_{2t}][\partial \mu_{2t}/\partial \mu_{1t}][\partial \mu_{1t}/\partial U_{t+1}][\partial U_{t+1}/\partial c_{t+1}]}{\partial U_t/\partial c_t} \\
&= \pi_(s_1|s_2)\pi_2(s_2)\ \beta \left(\frac{c_{t+1}(s_1, s_2)}{c_t}\right)^{\rho-1} \left(\frac{U_{t+1}(s_1, s_2)^{\alpha-\rho}}{\mu_{2t}[\mu_{1t}(U_{t+1})]^{\gamma-\rho}\mu_{1t}(U_{t+1})^{\alpha-\gamma}}\right) \\
&= \pi(s_1, s_2)\ \beta \left(\frac{c_{t+1}(s_1, s_2)}{c_t}\right)^{\rho-1} \left(\frac{U_{t+1}(s_1, s_2)}{\mu_{2t}[\mu_{1t}(U_{t+1})]}\right)^{\alpha-\rho} \left(\frac{\mu_{1t}(U_{t+1})}{\mu_{2t}[\mu_{1t}(U_{t+1})]}\right)^{\gamma-\alpha}.
\end{aligned}$$

B The Brock-Mirman example

Ljungqvist and Sargent (2000, Chapter 4, Appendix B) give the traditional stripped-down version of the Brock-Mirman example. Agents have log utility, which corresponds to our recursive preferences with $\rho = \alpha = \gamma = 0$. The laws of motion for capital and productivity are $k_{t+1} = y_t - c_t = a_t k_t^{\omega} - c_t$ and $\log a_{t+1} \;=\; \varphi_a \log a_t + \tau^{1/2} w_{t+1}$, where $0 \leq \omega < 1$ and $\{w_t\}$ is an iid sequence of standard normal random variables. The Bellman equation is

$$\log J(k_t, a_t) \;=\; \max_{c_t} \,(1-\beta)\log c_t + \beta E_t[\log J(k_{t+1}, a_{t+1})]$$

subject to the laws of motion for k_t and a_t.

We solve by guess and verify. Start with the guess $\log J(k_t, a_t) = p_0 + p_k \log k_t + p_a \log a_t$. Then next period's value function and its expectation are

$$\begin{aligned} \log J(k_{t+1}, a_{t+1}) &= p_0 + p_k \log(a_t k_t^{\omega} - c_t) + p_a(\varphi_a \log a_t + \tau^{1/2} w_{t+1}) \\ E_t \log(J_{t+1}) &= p_0 + p_k \log(a_t k_t^{\omega} - c_t) + p_a \varphi_a \log a_t. \end{aligned}$$

If we substitute into the Bellman equation, the envelope condition for k_t is

$$p_k / k_t \;=\; \beta p_k \omega a_t k_t^{\omega - 1} / (a_t k_t^{\omega} - c_t).$$

That gives us the decision rule

$$c_t \;=\; (1 - \beta\omega) a_t k_t^{\omega},$$

which implies the controlled law of motion $k_{t+1} = \beta\omega a_t k_t^{\omega}$. The Bellman equation is then

$$\begin{aligned} p_0 + p_k \log k_t + p_a \log a_t &= (1-\beta)[\log(1-\beta\omega) + \log a_t + \omega \log k_t] \\ &\quad + \beta\big\{p_0 + p_k[\log(\beta\omega) + \log a_t + \omega \log k_t] + p_a \varphi_a \log a_t\big\}. \end{aligned}$$

Lining up terms gives us the solution:

$$\begin{aligned} p_k &= (1-\beta)\omega/(1-\beta\omega) \;\Rightarrow\; 1 - p_k \;=\; (1-\omega)/(1-\beta\omega) \\ p_a &= [1 - \beta(1-p_k)]/(1-\beta\omega) \;=\; (1-\beta)/[(1-\beta\omega)(1-\beta\varphi_a)] \\ p_0 &= [(1-\beta)\log(1-\beta\omega) + \beta p_k \log(\beta\omega)]/(1-\beta). \end{aligned}$$

Note that $0 \leq p_k < 1$, which makes $J(k_t, a_t)$ (weakly) increasing and concave in k_t.

C Approximating the additive business cycle model

We compute an approximate loglinear solution to a business cycle model with additive preferences and constant volatility and compare it to Campbell's (1994) solution of the same model. The approaches are different, but they deliver the same decision rule.

We approach the problem as a dynamic program. With additive preferences ($\alpha = \rho$) the Bellman equation can be written

$$J(\tilde{k}_t, x_t) \;=\; \max_{\tilde{c}_t} \,(1-\beta)\tilde{c}_t^{\rho}/\rho + \beta E_t[g_{t+1}^{\rho} J(\tilde{k}_{t+1}, x_{t+1})].$$

subject to the laws of motion $\tilde{k}_{t+1} = [f(\tilde{k}_t, 1) - \tilde{c}_t]/g_{t+1}$ and $x_{t+1} = Ax_t + Bw_{t+1}$. Here we've taken equation (8) to the power ρ, divided by ρ, and redefined J_t^ρ/ρ as J_t. The first-order and envelope conditions are

$$\begin{aligned}(1-\beta)\tilde{c}_t^{\rho-1} &= \beta E_t\big(g_{t+1}^{\rho-1} J_{kt+1}\big) \\ J_{kt} &= \beta E_t\big(g_{t+1}^{\rho-1} J_{kt+1}\big) f_{kt}.\end{aligned}$$

Together they imply

$$(1-\beta)\tilde{c}_t^{\rho-1} = J_{kt}/f_{kt}.$$

Evidently a loglinear decision rule requires loglinear approximations of J_{kt} and f_{kt}. This illustrates a point we made earlier: in the additive case, we need the derivative of the value function but not the value function itself.

We guess the derivative of the value function has the form

$$\log J_{kt} = q_k \log \tilde{k}_t + q_x^\top x_t$$

with coefficients (q_k, q_x) to be determined. Putting this together with the loglinear approximations (16,17) gives us

$$\begin{aligned}\log J_{kt} - \log f_{kt} &= (q_k - \lambda_r) \log \tilde{k}_t + q_x^\top x_t \\ \log \tilde{c}_t &= -\sigma(q_k - \lambda_r) \log \tilde{k}_t - \sigma q_x^\top x_t \\ \log \tilde{k}_{t+1} &= \big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big] \log \tilde{k}_t + \sigma\lambda_c q_x^\top x_t - e^\top x_{t+1}.\end{aligned}$$

Ordinarily we would substitute the decision rule for consumption into the Bellman equation and solve for q_k. Here it's sufficient to use the envelope condition, the derivative of the Bellman equation. The right-hand side involves

$$\begin{aligned}\log(g_{t+1}^{\rho-1} J_{kt+1}) &= q_k \log \tilde{k}_{t+1} + [q_x + (\rho-1)e]^\top x_{t+1} \\ &= q_k\big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big] \log \tilde{k}_t \\ &\quad + q_k\sigma\lambda_c q_x^\top x_t + [q_x + (\rho - 1 - q_k)e]^\top (Ax_t + Bw_{t+1}) \\ \log E_t(g_{t+1}^{\rho-1} J_{kt+1}) &= q_k\big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big] \log \tilde{k}_t \\ &\quad + \big\{\sigma\lambda_c q_k q_x + [q_x + (\rho - 1 - q_k)e]^\top A\big\} x_t + V_x/2,\end{aligned}$$

where $V_x = [q_x + (\rho-1)e]^\top BB^\top [q_x + (\rho-1)e]$. The variance term V_x only shows up in the intercept, so we ignore it from here out. The envelope condition then gives us

$$\begin{aligned}q_k \log \tilde{k}_t + q_x^\top x_t &= q_k\big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big] \log \tilde{k}_t \\ &\quad + \big\{q_k\lambda_c\sigma q_x^\top + [q_x + (\rho-1)e]^\top A\big\} x_t + \lambda_r \log \tilde{k}_t.\end{aligned}$$

Equating similar terms, we have

$$\begin{aligned}q_k &= q_k\big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big] + \lambda_r \\ q_x^\top &= \sigma\lambda_c q_k q_x^\top + [q_x + (\rho-1)e]^\top A.\end{aligned}$$

The first equation is quadratic in q_k and has two solutions, one positive and one negative. We take the negative one, which corresponds to a concave value function and a stable controlled law of motion. Note the separation property: the solution for q_k is independent of anything related to the shock x. Given solutions for q_k and q_x, the decision rule for consumption follows immediately.

Campbell derives the same decision rule by another route. He starts with the "Euler equation,"

$$E_t\big[\beta(c_{t+1}/c_t)^{\rho-1} f_{kt+1}\big] \;=\; 1,$$

a consequence of the first-order and envelope conditions. Ignoring risk (which is constant, in any case), the loglinear version is

$$E_t(\log c_{t+1} - \log c_t) \;=\; -\sigma E_t(\log f_{kt+1}) \;=\; \sigma\lambda_r E_t(\log \tilde{k}_{t+1}). \tag{22}$$

Where we use a guess for J_{kt}, he uses a guess for the consumption decision rule,

$$\log \tilde{c}_t \;=\; h_{ck} \log \tilde{k}_t + h_{cx}^\top x_t,$$

with coefficients (h_{ck}, h_{cx}) to be determined. Using (16,17), the (controlled) law of motion for capital is

$$\log \tilde{k}_{t+1} \;=\; (\lambda_k - \lambda_c h_{ck}) \log \tilde{k}_t - \lambda_c h_{cx}^\top x_t - e^\top x_{t+1}.$$

The right side of the Euler equation (22) then becomes

$$\sigma\lambda_r E_t(\log \tilde{k}_{t+1}) \;=\; \sigma\lambda_r(\lambda_k - \lambda_c h_{ck}) \log \tilde{k}_t - \sigma\lambda_r\big(\lambda_c h_{cx}^\top + e^\top A\big) x_t.$$

The left side becomes

$$\begin{aligned} E_t(\log c_{t+1} - \log c_t) &= h_{ck}(\log \tilde{k}_{t+1} - \log \tilde{k}_t) + h_{cx}^\top(x_{t+1} - x_t) + e^\top x_{t+1} \\ &= [h_{ck}(\lambda_k - \lambda_c h_{ck}) - h_{ck}] \log \tilde{k}_t + [h_{cx}^\top(A - I) + e^\top A + h_{ck}\lambda_c h_{cx}^\top] x_t. \end{aligned}$$

Equating the two gives us

$$\begin{aligned} h_{ck}(\lambda_k - \lambda_c h_{ck}) - h_{ck} &= \sigma\lambda_r(\lambda_k - \lambda_c h_{ck}) \\ h_{cx}^\top(A - I) + e^\top A + h_{ck}\lambda_c h_{cx}^\top &= -\sigma\lambda_r\big(\lambda_c h_{cx}^\top + e^\top A\big). \end{aligned}$$

The first equation is Campbell's equation (24) in slightly different notation. Kaltenbrunner and Lochstoer (2010) do the same. It's tedious but direct to show that the two solutions are identical: use the relation $h_{ck} = -\sigma(q_k - \lambda_r)$ to convert the quadratic in h_{ck} into one for q_k.

D Approximating the risky business cycle model

We take two approaches to loglinear approximation of the model in Section 5. The first is based on a loglinear approximation to the Bellman equation suggested by Hansen, Heaton, and Li (2008, Section III). The second is based on a loglinear approximation of the envelope condition, the derivative of the Bellman equation. We find the second more helpful, but in principle the two should give similar answers.

Value function approximation. Relative to the additive case of Appendix C, we need loglinear approximations of both the value function and its derivative. Consider a loglinear approximation to J_{kt}:

$$\begin{aligned} \log J_{kt} &= p_1 + \log p_k + (p_k - 1)\log \tilde{k}_t + p_x^\top x_t + p_v v_t \\ \Rightarrow\ J_{kt} &= p_k \tilde{k}_t^{p_k - 1} \exp(p_1 + p_x^\top x_t + p_v v_t). \end{aligned}$$

If we integrate with respect to $\tilde{k}_t$, we get the value function

$$J_t = p_0 + \tilde{k}_t^{p_k} \exp(p_1 + p_x^\top x_t + p_v v_t)$$

It's unfortunate that J_t isn't loglinear unless $p_0 = 0$, but we use a loglinear approximation,

$$\log J_t = d(p_k \log \tilde{k}_t + p_x^\top x_t + p_v v_t)$$

with $d = (J - p_0)/J$. The combined term

$$\log(J_t^{\rho-1} J_{kt}) = \{[1 + (\rho - 1)d]p_k - 1\} \log \tilde{k}_t + [1 + (\rho - 1)d](p_x^\top x_t + p_v v_t).$$

appears in the combined first-order and envelope condition (15).

Hansen-Heaton-Li approximation. The idea is to approximate (10), the log of the Bellman equation:

$$\begin{aligned} \log J_t &= \rho^{-1} \log \left[(1-\beta) e^{\rho \log \tilde{c}_t} + \beta e^{\rho \log \mu_t(g_{t+1} J_{t+1})}\right] \\ &\cong b_0 + (1 - b_1) \log \tilde{c}_t + b_1 \log \mu_t(g_{t+1} J_{t+1}). \end{aligned} \tag{23}$$

This is exact if $\rho = 0$, in which case $b_0 = 0$ and $b_1 = \beta$.

We need three things to put this to work: the value function J_t, consumption $\tilde{c}_t$, and the certainty equivalent of future utility $\mu_t(g_{t+1} J_{t+1})$. The first one we've done. Condition (15) then gives us the decision rule

$$\log \tilde{c}_t = [(d - \sigma)p_k + \sigma(1 + \lambda_r)] \log \tilde{k}_t + (d - \sigma)(p_x^\top x_t + p_v v_t),$$

the second component of the approximate Bellman equation. The final component is the certainty equivalent of future utility. For that, we need the controlled law of motion (17),

$$\begin{aligned} \log \tilde{k}_{t+1} &= \lambda_k \log \tilde{k}_t - \lambda_c \log \tilde{c}_t - e^\top x_{t+1} \\ &= \left\{\lambda_k - \lambda_c[(d - \sigma)p_k + \sigma(1 + \lambda_r)]\right\} \log \tilde{k}_t - \lambda_c(d - \sigma)(p_x^\top x_t + p_v v_t) - e^\top x_{t+1}. \end{aligned}$$

Future utility and its certainty equivalent are then

$$\begin{aligned}
\log(g_{t+1}J_{t+1}) &= \log g_{t+1} + d(p_k \log \tilde{k}_{t+1} + p_x^\top x_{t+1} + p_v v_{t+1}) \\
&= dp_k\big\{\lambda_k - \lambda_c[(d-\sigma)p_k + \sigma(1+\lambda_r)]\big\}\log \tilde{k}_t \\
&\quad + \big\{[(1-dp_k)e^\top + dp_x^\top]A - (\lambda_c dp_k(d-\sigma)p_x^\top\big\}x_t \\
&\quad + dp_v[\varphi_v - \lambda_c dp_k(d-\sigma)]v_t \\
&\quad + [(1-dp_k)e + dp_x]^\top v_t^{1/2}Bw_{1t+1} + dp_v\tau w_{2t+1} \\
\log \mu_t(g_{t+1}J_{t+1}) &= dp_k\big\{\lambda_k - \lambda_c[(d-\sigma)p_k + \sigma(1+\lambda_r)]\big\}\log \tilde{k}_t \\
&\quad + \big\{[(1-dp_k)e^\top + dp_x^\top]A - \lambda_c dp_k(d-\sigma)p_x^\top\big\}x_t \\
&\quad + \big\{dp_v[\varphi_v - \lambda_c dp_k(d-\sigma)] + (\alpha/2)V_x\big\}v_t,
\end{aligned}$$

where $V_x = [dp_x + (1-dp_k)e]^\top BB^\top[dp_x + (1-dp_k)e]$ is the contribution of x_{t+1} to the conditional variance of $\log(g_{t+1}J_{t+1})$.

Now we plug this into the approximate Bellman equation and line up coefficients:

$$\begin{aligned}
p_k &= (1-b_1)\big[\sigma + (d-\sigma)p_k + \sigma\lambda_r\big] + b_1 dp_k\big\{\lambda_k - \lambda_c[\sigma + (d-\sigma)p_k + \sigma\lambda_r]\big\} \\
p_x^\top &= (1-b_1)(d-\sigma)p_x^\top + b_1\big\{[(1-dp_k)e^\top + dp_x^\top]A - \lambda_c dp_k(d-\sigma)p_x^\top\big\} \\
p_v &= (1-b_1)(d-\sigma)p_v + b_1\big\{dp_v[\varphi_v - \lambda_c dp_k(d-\sigma)] + (\alpha/2)V_x\big\}.
\end{aligned}$$

The first equation is quadratic in p_k. Given a solution for p_k, the equations for p_x and p_v are linear.

Envelope condition approximation. Our second method is based on the envelope condition and mirrors the approach we took to the additive model in Appendix C. The envelope condition for the recursive model is

$$J_t^{\rho-1}J_{kt} = \beta\mu_t(g_{t+1}J_{t+1})^{\rho-\alpha}E_t[(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}]f_{kt}. \tag{24}$$

For the left side we reconstitute our earlier expression,

$$\begin{aligned}
\log(J_t^{\rho-1}J_{kt}) &= \big\{[1+(\rho-1)d]p_k - 1\big\}\log \tilde{k}_t + [1+(\rho-1)d](p_x^\top x_t + p_v v_t) \\
&= q_k \log \tilde{k}_t + q_x^\top x_t + q_v v_t.
\end{aligned}$$

The substitution of coefficients (q_k, q_x, q_v) for (p_k, p_x, p_v) is helpful because the same terms reappear elsewhere.

We need to evaluate $g_{t+1}J_{t+1}$, which requires the decision rule and controlled law of motion. The decision rule in this notation is

$$\log \tilde{c}_t = -\sigma(q_k - \lambda_r)\log \tilde{k}_t - \sigma(q_x^\top x_t + q_v v_t).$$

That gives us the controlled law of motion

$$\log \tilde{k}_{t+1} = \big[\lambda_k + \sigma\lambda_c(q_k - \lambda_r)\big]\log \tilde{k}_t + \sigma\lambda_c(q_x^\top x_t + q_v v_t) - e^\top x_{t+1}.$$

Now back to the envelope condition (24). The critical components are

$$
\begin{aligned}
\log(g_{t+1}J_{t+1}) &= dp_k \log \tilde{k}_{t+1} + (e+dp_x)^\top x_{t+1} + dp_v v_{t+1} \\
&= dp_k \big[\lambda_k + \sigma\lambda_c(q_k-\lambda_r)\big] \log \tilde{k}_t + dp_k\lambda_c\sigma q_x^\top x_t + dp_k\lambda_c\sigma q_v v_t \\
&\quad + [(1-dp_k)e + dp_x]^\top (Ax_t + v_t^{1/2}Bw_{1t+1}) + dp_v[(1-\varphi_v)v + \varphi_v v_t + \tau w_{2t+1}] \\
\log J_{kt+1} &= (p_k-1)\log \tilde{k}_{t+1} + p_x^\top x_{t+1} + p_v v_{t+1} \\
&= (p_k-1)\big[\lambda_k + \sigma\lambda_c(q_k-\lambda_r)\big] \log \tilde{k}_t + (p_k-1)\lambda_c\sigma q_x^\top x_t + (p_k-1)\lambda_c\sigma q_v v_t \\
&\quad + [(1-p_k)e + p_x]^\top (Ax_t + v_t^{1/2}Bw_{1t+1}) + p_v[(1-\varphi_v)v + \varphi_v v_t + \tau w_{2t+1}],
\end{aligned}
$$

which show up in $\mu_t(g_{t+1}J_{t+1})$ and $E_t[(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}]$. We collect similar terms and find:

- Capital. These terms show up only in the mean. If we work through the right side of the envelope condition, we see that the mean terms are multiplied by $(\rho-\alpha)+(\alpha-1)=\rho-1$. The $\log \tilde{k}_t$ terms are therefore

  $$
  \begin{aligned}
  q_k &= (\rho-1)dp_k[\lambda_k + \sigma\lambda_c(q_k-\lambda_r)] + (p_k-1)[\lambda_k + \sigma\lambda_c(q_k-\lambda_r)] + \lambda_r \\
  &= q_k[\lambda_k + \sigma\lambda_c(q_k-\lambda_r)] + \lambda_r.
  \end{aligned}
  $$

  In the first line, the first term on the right comes from $\log(g_{t+1}J_{t+1})$ and the second from $\log J_{kt+1}$. Through some piece of luck, they combine nicely. We are left with a quadratic in the coefficient q_k, in fact the same one we derived in Appendix C for the additive case.

  As a result, the coefficients h_{ck} of the approximate decision rule (12) and h_{kk} of the controlled law of motion (18) are the same in the additive and recursive cases. As in Tallarini (2000), the risk aversion parameter plays no role in what we call the "internal dynamics" of the capital stock. Note, too, that the behavior of the exogenous state variables x_t and v_t have no impact on the solution. This is what we call the separation property: the capital coefficients (q_k, h_{ck}, h_{kk}) are independent of the rest of the model.

- News. Again, x_t shows up only in the mean terms, so we find the coefficient q_x in much the same way. The relevant terms are

  $$
  q_x^\top = q_k\lambda_c\sigma q_x^\top + [(\rho-1)e + q_x]^\top A.
  $$

  This is, again, the same as the additive case we solved earlier. It's also independent of risk (v_t) and risk aversion (α): capital dynamics enter through q_k, but the properties of risk don't affect the response to news. When $A=0$, there's no persistence in x_t and $q_x = 0$.

- Risk. This one's more involved, it incorporates risk and recursive preferences in a fundamental way. The coefficients of v_t might be collected in two groups:

  $$
  q_v = \text{mean terms} + \text{variance terms}.
  $$

The former are similar to what we've done:

$$\text{mean terms} \;\; = \;\; (q_k \lambda_c \sigma) q_v.$$

The latter involve the variances of terms containing $v_t^{1/2}$ in the envelope condition (24) in both the certainty equivalent $\mu_t(g_{t+1}J_{t+1})$ and the expectation $E_t[(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}]$. Adding them in order, we have

$$\text{variance terms} \;\; = \;\; (\rho-\alpha)\alpha V_1/2 + V_2/2$$

where

$$\begin{aligned} V_1 &= [(1-dp_k)e + dp_x]^\top BB^\top [(1-dp_k)e + dp_x] \\ V_2 &= \big\{(\alpha-1)[(1-dp_k)e + dp_x] + [(1-p_k)e + p_x]\big\}^\top BB^\top \\ &\quad\; \big\{(\alpha-1)[(1-dp_k)e + dp_x] + [(1-p_k)e + p_x]\big\}. \end{aligned}$$

We need a value for d to implement this, but we can get a sense of where the sign comes from. Both V_1 and V_2 are positive, so the sign depends on their relative magnitudes and the sign of $\rho-\alpha$.

E Approximating the ambiguous business cycle model

Constant volatility. The ambiguous business cycle model has the same structure and properties as the model with risk. In the notation of Appendix D, let

$$\begin{aligned} \log(J_t^{\rho-1}J_{kt}) &= [((\rho-1)d+1)p_k - 1]\log \tilde{k}_t + [((\rho-1)d+1)p_x]\hat{x}_t \\ &= q_k \log \tilde{k}_t + q_x \hat{x}_t. \end{aligned}$$

As before, the calculation the right-hand side of (21) follows from loglinear approximations of the relevant functions and the certainty equivalent formula for lognormal random variables. Consider first the term inside the square brackets. We will need the expected value of the log (the constant variances will ultimately enter the intercept terms which we ignore):

$$\begin{aligned} & E_{1t}[(\alpha-1)(\log g_{t+1} + \log J_{t+1}) + \log J_{kt+1}] \\ =\;& E_{1t}[(\rho-1)\log J_{t+1} + \log J_{kt+1}] + (\rho-1)E_{1t}\log g_{t+1} + (\alpha-\rho)E_{1t}[\log g_{t+1} + \log J_{t+1}] \\ =\;& E_{1t}[q_k \log \tilde{k}_{t+1} + q_x \widehat{x}_{t+1}] + (\rho-1)E_{1t}\log g_{t+1} + (\alpha-\rho)E_{1t}[\log g_{t+1} + \log J_{t+1}]. \end{aligned}$$

The last expected value in this expression, $E_{1t}[\log g_{t+1} + \log J_{t+1}]$, is precisely the expectation in μ_{1t}, which affords some convenient cancellation of terms, such that the expectation inside the square brackets on the right-hand side of (21) is simply

$$\begin{aligned} \log E_{2t} \log \left[\mu_{1t}^{\gamma-\alpha} E_{1t}(g_{t+1}J_{t+1})^{\alpha-1}J_{kt+1}\right] &= E_{2t}E_{1t}[q_k \log \tilde{k}_{t+1} + q_x \widehat{x}_{t+1} + (\rho-1)\log g_{t+1}] \\ &\quad + (\gamma-\rho)E_{2t}E_{1t}[\log g_{t+1} + \log J_{t+1}]. \end{aligned}$$

Once again, the last expected value in this expression is precisely the expectation that appears in the term $\mu_{2t}(\mu_{1t})^{\rho-\gamma}$, but with the opposite sign, which once again affords some convenient cancellation of terms resulting in greatly simplified version of equation (21):

$$
\begin{aligned}
(q_k - \lambda_r)\log \tilde{k}_t + q_x \widehat{x}_t &= E_{2t}E_{1t}[q_k \log \tilde{k}_{t+1} + q_x \widehat{x}_{t+1} + (\rho - 1)\log g_{t+1}] \\
&= q_k h_{kk} \log \tilde{k}_t + [q_k h_{kx} + q_x \varphi_x (1-\theta)]\widehat{x}_t \\
&\quad + [\rho - 1 - q_k + q_x \varphi_x \theta] E_{2t}E_{1t} \log g_{t+1} \\
&= q_k h_{kk} \log \tilde{k}_t + [q_k(h_{kx} - 1) + q_x \varphi_x + \rho - 1]\widehat{x}_t.
\end{aligned}
$$

(Recall: $h_{kk} = \lambda_k + \sigma\lambda_c(q_k - \lambda_r)$ and $h_{kx} = \sigma\lambda_c q_x$.) This equation gives us identical loglinear solutions as the earlier models. Neither the ambiguity aversion parameter γ nor the risk aversion parameter α affects the dynamics of the model. (Both parameters appear only in the intercept terms.)

References

Alvarez, Fernando, and Urban Jermann, 2005, "Using asset prices to measure the persistence of the marginal utility of wealth," *Econometrica* 73, 1977-2016.

Anderson, Evan, Lars Hansen, Ellen McGrattan, and Thomas Sargent, 1996, "Mechanics of forming and estimating dynamic linear economies," *Handbook of Computational Economics*, ed. H.M. Amman, D.A. Kendrick, and J. Rust, Elsevier Science.

Arellano, Cristina, Yan Bai, and Patrick Kehoe, 2012, "Financial frictions and fluctuations in volatility," manuscript, Federal Reserve Bank of Minneapolis, July.

Bachmann, Rudiger, and Christian Bayer, 2013, "Wait-and-see business cycles?," *Journal of Monetary Economics* 60, 704-719.

Backus, David, Bryan Routledge, and Stanley Zin, 2005, "Exotic preferences for macroeconomists," in Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual 2004*, MIT Press.

Backus, David, Mikhail Chernov, and Stanley Zin, 2014, "Sources of entropy in representative agent models," *Journal of Finance* 69, 51-99.

Bansal, Ravi, and Amir Yaron, 2004, "Risks for the long run: A potential resolution of asset pricing puzzles," *Journal of Finance* 59, 1481-1509.

Barro, Robert J., and Robert G. King, 1984, "Time-separable preferences and intertemporal-substitution models of business cycles, *Quarterly Journal of Economics* 99, 817-839.

Basu, Susanto, and Brent Bundick, 2012, "Uncertainty shocks in a model of effective demand," manuscript, Boston College, September.

Bethune, Zach, Thomas Cooley and Peter Rupert, 2014, "GDP Report," US Economic Snapshot blog, March 31.

Bidder, Rhys, and Matthew Smith, 2012, "Robust animal spirits," *Journal of Monetary Economics* 59, 738-750.

Bloom, Nicholas, 2013, "Fluctuations in uncertainty," NBER Working Paper 19714, December.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry, 2012, "Really uncertain business cycles," NBER Working Paper 18245, July.

Caldara, Dario, Jesus Fernandez-Villaverde, Juan F. Rubio-Ramirez, and Yao Wen, 2012, "Computing DSGE models with recursive preferences and stochastic volatility," *Review of Economic Dynamics* 15, 188-206.

Campbell, John, 1994, "Inspecting the mechanism: An analytical approach to the real business cycle model," *Journal of Monetary Economics* 33, 463-506.

Collard, Fabrice, Sujoy Mukerji, Kevin Sheppard, and Jean-Marc Tallon, 2012, "Ambiguity and the historical equity premium," manuscript, August.

Collin-Dufresne, Pierre, Michael Johannes, and Lars Lochstoer, 2013, "Parameter learning in general equilibrium: The asset pricing implications," manuscript, May.

Dew-Becker, Ian, 2012, "Essentially affine approximations for economic models," manuscript, June.

Epstein, Larry G., and Stanley E. Zin, 1989, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica* 57, 937-969.

European Commission, 2013, *Quarterly Report on the Euro Area* 12, European Union.

Fajgelbaum, Pablo, Edouard Schaal, and Mathieu Taschereau-Dumouchel, 2014, "Uncertainty traps," manuscript, March.

Fostel, Ana, and John Geankoplos, 2012, "Why does bad news increase volatility and decrease leverage?" *Journal of Economic Theory* 147, 501-525.

Fostel, Ana, and John Geankoplos, 2014, "Endogenous collateral constraints and the leverage cycle," *Annual Review of Economics* 6, forthcoming.

Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, "Consumption strikes back? Measuring long-run risk," *Journal of Political Economy* 116, 260-302.

Hansen, Lars Peter, and Thomas J. Sargent, 2007, *Robustness*, Princeton: Princeton University Press.

Hayashi, Takashi, and Jianjun Miao, 2011, "Intertemporal substitution and recursive smooth ambiguity preferences," *Theoretical Economics* 6, 423-472.

Ilut, Cosmin, and Martin Schneider, 2014, "Ambiguous business cycles," *American Economic Review*, forthcoming.

Jahan-Parvar, Mohammad, and Hening Liu, 2012, "Ambiguity aversion and asset prices in production economies," manuscript.

Ju, Nengjiu, and Jianjun Miao, 2012, "Ambiguity, learning, and asset returns," *Econometrica* 80, 559-591.

Justiniano, Alejandro, and Giorgio Primiceri, 2008, "The time-varying volatility of macroeconomic fluctuations," *American Economic Review* 98, 604-681.

Kaltenbrunner, Georg, and Lars Lochstoer, 2010, "Long-run risk through consumption smoothing," *Review of Financial Studies* 23, 3190-3224.

Klibanoff, Peter, Massimo Marinacci, Sujoy Mukerji, 2005, "A smooth model of decision making under ambiguity," *Econometrica* 73, 1849-1892.

Klibanoff, Peter, Massimo Marinacci, Sujoy Mukerji, 2009, "Recursive smooth ambiguity preferences," *Journal of Economic Theory* 144, 930-976.

Kreps, David M., and Evan L. Porteus, 1978, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica* 46, 185-200.

Kydland, Finn, and Edward Prescott, 1982, "Time to build and aggregate fluctuations," *Econometrica* 50, 1345-1370.

Liu, Hening, and Jianjun Miao, 2014, "Growth uncertainty, generalized disappointment aversion, and production-based asset pricing," manuscript prepared for the April 2014 Carnegie-Rochester-NYU conference.

Ljungqvist, Lars, and Thomas Sargent, 2000, *Recursive Macroeconomic Theory*, Cambridge, MA: MIT Press.

Malkhozov, Aytek, 2014, "Asset prices in affine real business cycle models," *Review of Economic Dynamics*, forthcoming.

Navarro, Gaston, 2014, "Financial crises and endogenous volatility," manuscript, New York University.

Sprague, Shawn, 2014, "What can labor productivity tell us about the US economy?" *Beyond the Numbers* 12, Bureau of Labor Statistics.

Tallarini, Thomas D., 2000, "Risk-sensitive real business cycles," *Journal of Monetary Economics* 45, 507-532.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2006, "Learning asymmetries in real business cycles," *Journal of Monetary Economics* 53, 753-772.

Veldkamp, Laura, 2005, "Slow boom, sudden crash," *Journal of Economic Theory* 124, 230-257.

Figure 1
Output over the last five recessions

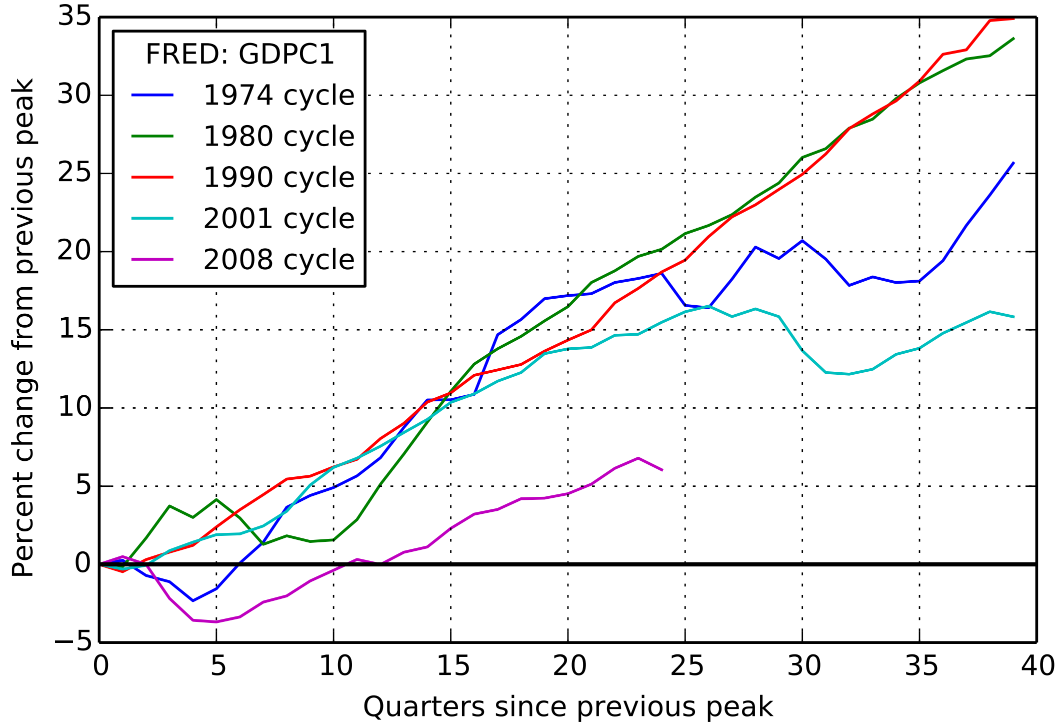

Notes: The figure plots the last five recessions starting from the previous peak. Adapted from Bethune, Cooley, and Rupert (2014).

Figure 2
Consumption over the last five recessions

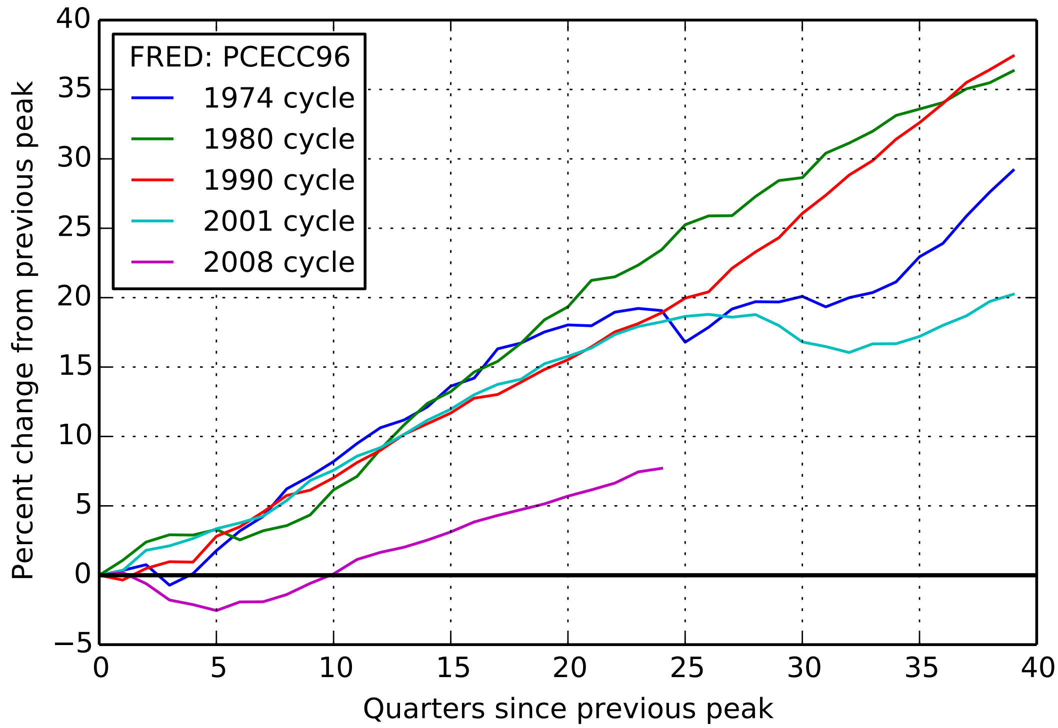

Notes: The figure plots the last five recessions starting from the previous peak. Adapted from Bethune, Cooley, and Rupert (2014).

Figure 3
Investment over the last five recessions

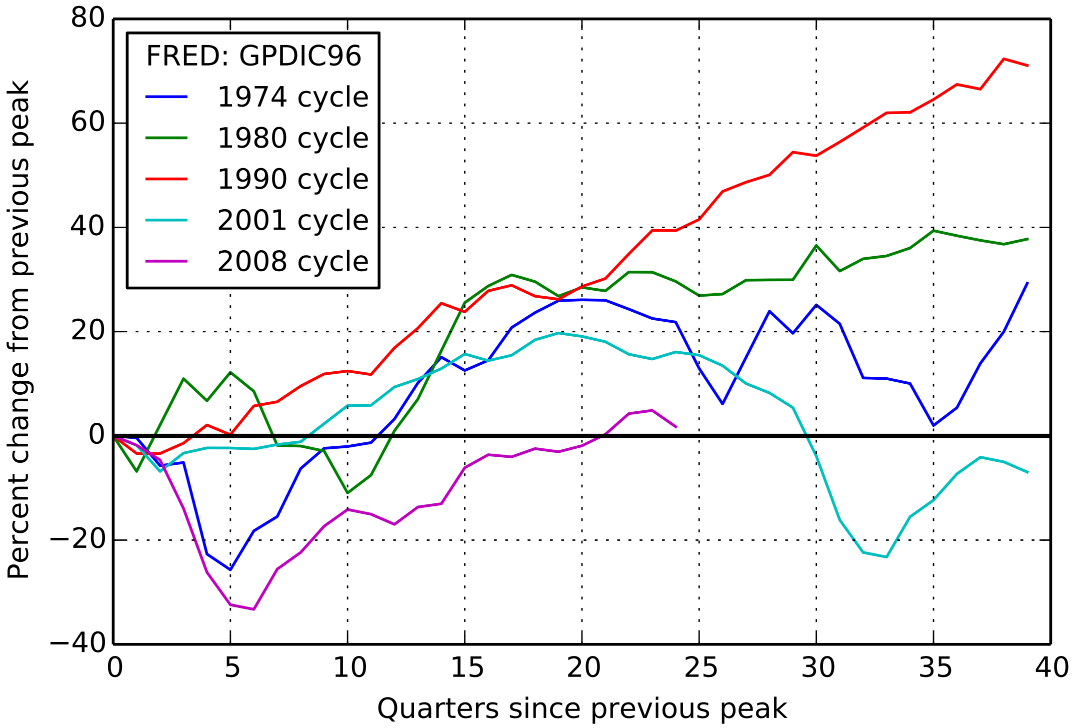

Notes: The figure plots the last five recessions starting from the previous peak. Adapted from Bethune, Cooley, and Rupert (2014).

Figure 4
Labor productivity over the last five recessions

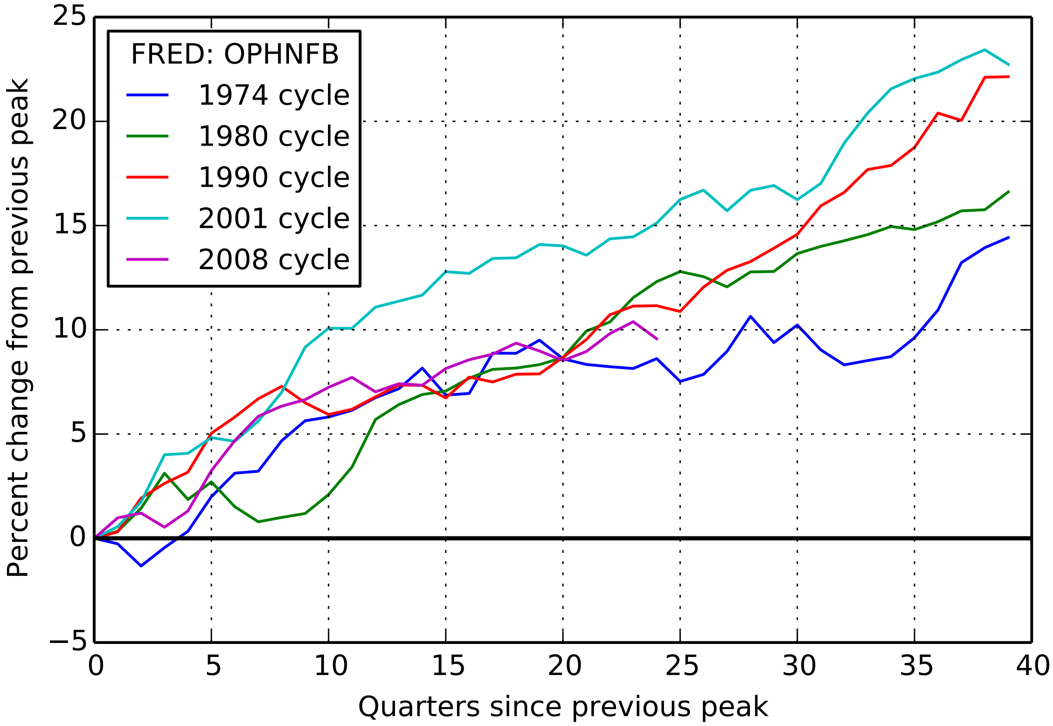

Notes: The figure plots the last five recessions starting from the previous peak. Adapted from Bethune, Cooley, and Rupert (2014).

Figure 5
Slope of controlled law of motion: loglinear approximation and numerical solution

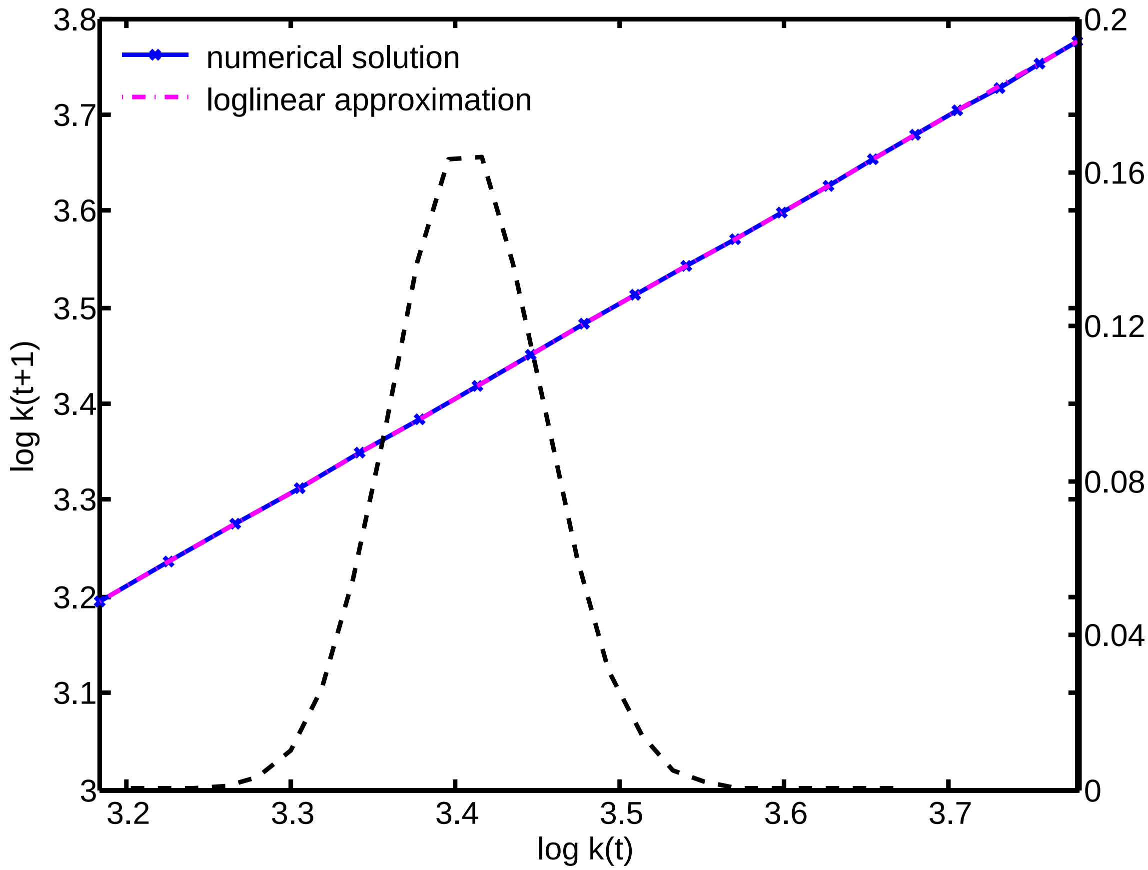

Figure 6
Response to a one standard deviation increase in conditional variance

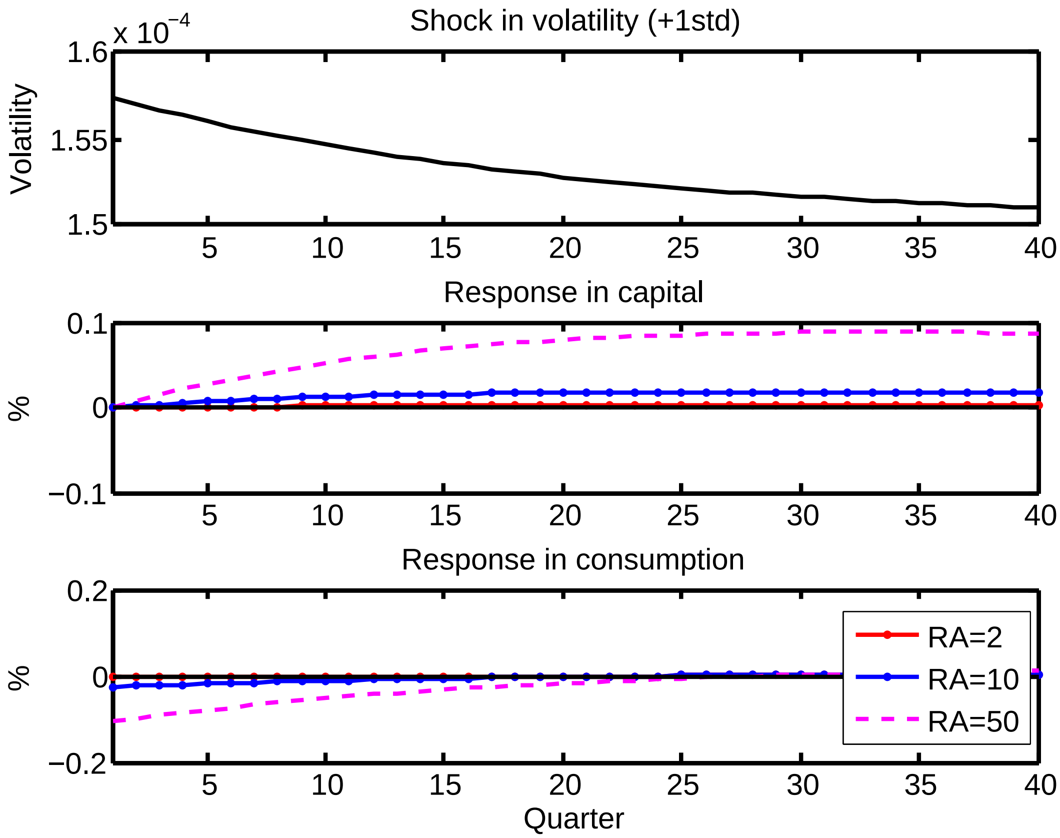

Table 1
Benchmark parameter values

Parameter	Value	Comment
(a) Preferences		
ρ	-1	IES $= \sigma = 1/(1-\rho) = 1/2$
α	-9	RA $= 1-\alpha = 10$, Bansal & Yaron (2004, Table II)
β	—	chosen to hit $k/y = 10$ (quarterly)
(b) Technology		
ν	0	Cobb-Douglas
ω	1/3	Kydland & Prescott (1982, Table I), rounded off
δ	0.025	Kydland & Prescott (1982, Table I)
(c) Productivity growth		
$\log g$	0.004	Tallarini (2000, Table 4)
e	1	normalization
A	0	no predictable component (news)
B	1	normalization
$v^{1/2}$	0.015	Tallarini (2000, Table 4), rounded off
φ_v	0.95	arbitrary
τ	0.74×10^{-5}	makes v three standard deviations from zero

Table 2
Properties of business cycle models with risk

Variable	US Data (1)	Additive (2)	Recursive (3)	Recursive (4)	Recursive (5)	Recursive (6)
(a) Parameter values						
Changes from benchmark		$1-\alpha=2$		$1-\alpha=50$	$\tau=0$	$\sigma=1.5$
(b) Standard deviations (%)						
Output growth	1.04	0.82	0.82	0.82	0.82	0.82
Consumption growth	0.55	0.75	0.75	0.76	0.75	0.47
Investment growth	2.79	1.03	1.04	1.06	1.02	1.82
(c) Standard deviations relative to output growth						
Consumption growth	0.53	0.91	0.91	0.93	0.91	0.58
Investment growth	2.69	1.26	1.27	1.30	1.24	2.23
(d) Correlations with output growth						
Consumption growth	0.52	0.99	0.99	0.97	0.99	0.91
Investment growth	0.65	0.98	0.97	0.93	0.98	0.96
(e) Asset Pricing						
Entropy bound (%)		0.02	0.58	16.33	0.57	0.60

Notes. US data from Tallarini (2000, Table 6). Parameter values from Table 1 except where noted.

Table 3
Properties of business cycle models with ambiguity

Variable	US Data (1)	Model I (2)	Model I (3)	Model I (4)	Model I (5)	Model II (6)	Model II (7)
(a) Parameter values							
Ambiguity aversion $1-\gamma$		30	30	50	50	50	50
Ambiguity τ_2^2		0.9	0.1	0.9	0.1	0.9	0.1
(b) Standard deviations (%)							
Output growth	1.04	0.82	0.82	0.82	0.82	0.82	0.82
Consumption growth	0.55	0.71	0.71	0.71	0.71	0.72	0.70
Investment growth	2.79	1.09	1.09	1.16	1.11	1.14	1.15
(c) Standard deviations relative to output growth							
Consumption growth	0.53	0.87	0.87	0.87	0.86	0.88	0.85
Investment growth	2.69	1.34	1.33	1.41	1.36	1.39	1.41
(d) Correlations with output growth							
Consumption growth	0.52	0.99	0.99	0.98	0.99	0.98	0.99
Investment growth	0.65	0.98	0.99	0.96	0.99	0.96	0.99
(e) Asset Pricing							
Entropy bound (%)		5.15	1.07	13.76	2.06	13.93	1.97

Notes. US data from Tallarini (2000, Table 6). Parameter values from Table 1 except where noted.