NBER WORKING PAPER SERIES

CONSUMPTION-BASED ASSET PRICING WITH HIGHER CUMULANTS

Ian Martin

Working Paper 16153
http://www.nber.org/papers/w16153

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2010

First draft: 20 August, 2006. I thank Dave Backus, Robert Barro, Emmanuel Farhi, Xavier Gabaix, Simon Gilchrist, Francois Gourio, Greg Mankiw, Anthony Niblett, Jeremy Stein, Adrien Verdelhan, Martin Weitzman and, in particular, John Campbell for their comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Consumption-Based Asset Pricing with Higher Cumulants
Ian Martin
NBER Working Paper No. 16153
July 2010
JEL No. E44,G10

ABSTRACT

I extend the Epstein-Zin-lognormal consumption-based asset-pricing model to allow for general i.i.d. consumption growth. Information about the higher moments--equivalently, cumulants--of consumption growth is encoded in the cumulant-generating function. I apply the framework to economies with rare disasters, and argue that the importance of such disasters is a double-edged sword: parameters that govern the frequency and sizes of rare disasters are critically important for asset pricing, but extremely hard to calibrate. I show how to sidestep this issue by using observable asset prices to make inferences that are robust to the details of the underlying consumption process.

Ian Martin
Graduate School of Business
Stanford University
Stanford, CA 94305
and NBER
ian.martin@gsb.stanford.edu

The combination of power utility and i.i.d. lognormal consumption growth makes for a tractable benchmark model in which asset prices and expected returns can be found in closed form. Introducing the consumption-based model, Cochrane (2005, p. 12) writes, "The combination of lognormal distributions and power utility is one of the basic tricks to getting analytical solutions in this kind of model."

This paper demonstrates that the lognormality assumption can be dropped without sacrificing tractability, thereby allowing for straightforward and flexible analysis of the possibility that, say, consumption is subject to occasional disasters. There has recently been considerable interest in reviving the idea of Rietz (1988) that the presence of rare disasters, or fat tails more generally, can help to explain asset pricing phenomena such as the riskless rate, equity premium and other puzzles (Barro (2006a), Farhi and Gabaix (2008), Gabaix (2008), Jurek (2008)). Here, I take a different line, closer in spirit to Weitzman (2007), and argue that the importance of rare, extreme events is a double-edged sword: those model parameters which are most important for asset prices, such as disaster parameters, are also the hardest to calibrate, precisely because the disasters in question are rare.

Working under two assumptions—that there is a representative agent with Epstein-Zin preferences and that consumption growth is i.i.d.—I exploit, in Section I, a mathematical object (the cumulant-generating function) in terms of which the equity premium, riskless rate, consumption-wealth ratio and mean consumption growth (the "fundamental quantities") can be simply expressed. Cumulant-generating functions crop up elsewhere in the finance literature; the contribution of this paper is to demonstrate how neatly they dovetail with the standard framework used in consumption-based asset-pricing and macroeconomics. Importantly, the framework allows for the possibility of disasters, but is agnostic about whether or not they occur. I present results in both discrete-time and continuous-time settings.

The expressions derived relate the fundamental quantities directly to the cumulants (equivalently, moments) of consumption growth. I show, for example, how the precautionary savings effect which determines the riskless rate in a lognormal model generalizes in the presence of higher cumulants.

I illustrate the framework by investigating a continuous-time model featuring rare disasters, and show that the model's predictions are sensitively dependent on the calibration assumed. As a stark example, take a consumption-based model in which the representative agent has relative risk aversion equal to 4. Now add to the model a certain type of disaster that strikes, on average, once every 1,000 years, and reduces consumption by 64 per cent. (Barro (2006a) documents that Germany and Greece each suffered such a fall in per capita real GDP during the Second World War.) The introduction of this disaster drives the riskless rate down by 5.9 percentage points and increases the equity premium by 3.7 per cent.[1] Very rare, very severe events exert an extraordinary influence on the benchmark model, and we do not expect to estimate their frequency and intensity directly from the data.

The remainder of the paper is devoted to finding ways around this pessimistic conclusion. We can, for example, detect the influence of disaster events *indirectly*, by observing asset prices. I argue, therefore, that the standard approach—calibrating a particular model and trying to fit the fundamental quantities—is not the way to go. I turn things round, viewing the fundamental quantities as observables, and making inferences from them. It then becomes possible to make nonparametric statements that are robust to the details of the consumption growth process.

In this spirit, I derive, in section III, robust restrictions on preference parameters that are valid in *any* Epstein-Zin-i.i.d. model that is consistent with the observed fundamentals. My results restrict the time-preference rate, ρ, and elasticity of intertemporal substitution, ψ, to lie in a certain subset of the positive quadrant. (See Figure 4.) These parameters are of central importance for financial and macroeconomic models. The restrictions depend only on the Epstein-Zin-i.i.d. assumptions and on observed values of the fundamental quantities, and not, for example, on any assumptions about the existence, frequency or size of disasters. They are complementary to econometric or

[1] The effect is smaller with Epstein-Zin preferences if the elasticity of substitution is greater than 1, but even with an elasticity of intertemporal substitution equal to 2, the riskless rate drops by 3.5 per cent.

experimental estimates of ψ and ρ, and are of particular interest because there is little agreement about the value of ψ. (Campbell (2003) summarizes the conflicting evidence.) I also show how good-deal bounds (Cochrane and Saá-Requejo (2000)) can be used to provide upper bounds on risk aversion, based once again on the fundamental quantities, without calibrating a consumption process.

This theme of making inferences from observable fundamentals without making assumptions about the tails of consumption growth recurs in Section IV. I consider the question, surveyed by Lucas (2003), of the cost of consumption risk. This cost turns out to depend on ρ and ψ and on two observables: mean consumption growth and the consumption-wealth ratio. The cost does not depend on risk aversion other than (implicitly) through the consumption-wealth ratio, which summarizes all relevant information about the attitude to risk of the representative agent and the amount of risk in the economy, as perceived by the representative agent.

In the power utility case, these welfare calculations apply to *any* consumption growth process, i.i.d. or not. My results therefore generalize Lucas (1987), Obstfeld (1994) and Barro (2006b). Unlike these authors, I view the consumption-wealth ratio as an observable. Using Barro's preferred preference parameters, I find that the cost of consumption fluctuations is about 14 per cent. I also calculate the welfare gains from a reduction in the variance of consumption growth, and show that almost all the cost of uncertainty can be attributed to the higher cumulants of consumption growth.

Campbell and Cochrane (1999) and Bansal and Yaron (2004) modify the textbook model along different dimensions. This paper explores different features, and implications, of the data, so is complementary to their work. In particular, both Campbell and Cochrane (1999) and Bansal and Yaron (2004) take care to work in a lognormal framework. It would, of course, be interesting to extend these papers by allowing for the possibility of jumps, but doing so would obscure the main point of this paper.

A large body of literature applies Lévy processes to derivative pricing (Carr and Madan (1998), Cont and Tankov (2004)) and portfolio choice (Kallsen (2000), Cvitanić, Polimenis and Zapatero (2005), Aït-Sahalia, Cacho-Diaz and Hurd (2006)). Lustig, Van

Nieuwerburgh and Verdelhan (2008) present estimates of the wealth-consumption ratio. Backus, Foresi and Telmer (2001) derive expressions relating cumulants to risk premia, though their approach is very different from that taken here. An alternative to the approach of this paper is to evaluate disaster models by considering a wider range of asset prices than typically considered in the consumption-based asset pricing literature. In this spirit, Julliard and Ghosh (2008) argue that the cross-section of asset price data is hard to square with disaster explanations of the equity premium, and Backus, Chernov and Martin (2009) explore the evidence in option prices.

I Asset-pricing and the cumulant-generating function

Define $G_t \equiv \log C_t/C_0$ and write $G \equiv G_1$. I make two assumptions.

A1 There is a representative agent with Epstein-Zin preferences, time preference rate ρ, relative risk aversion γ, and elasticity of intertemporal substitution ψ.

A2 The consumption growth, $\log C_t/C_{t-1}$, of the representative agent is i.i.d.,[2] and the moment-generating function of G (defined below) exists on the interval $[-\gamma, 1]$.[3]

Assumption A1 allows risk aversion γ to be disentangled from the elasticity of intertemporal substitution ψ. To keep things simple, those calculations that appear in the main text restrict to the power utility case in which ψ is constrained to equal $1/\gamma$; in this case, the representative agent maximizes

$$\mathbb{E}\sum_{t=0}^{\infty} e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma} \quad \text{if } \gamma \neq 1\,, \quad \text{or} \quad \mathbb{E}\sum_{t=0}^{\infty} e^{-\rho t}\log C_t \quad \text{if } \gamma = 1\,. \tag{1}$$

[2]Really, all that we need is that the representative agent *perceives* himself as having i.i.d. consumption growth and prices assets accordingly; the results of the paper go through without modification.

[3]If not, the consumption-based asset-pricing approach is invalid. This assumption implies, for example, that all moments of G are finite. See Billingsley (1995, Section 21).

Cogley (1990) and Barro (2006b) present evidence in support of A2 in the form of variance-ratio statistics close to one, on average, across nine (Cogley) or 19 (Barro) countries.

For now, I restrict to power utility. We need expected utility to be well defined in that

$$\mathbb{E}\sum_{t=0}^{\infty}\left|e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma}\right| < \infty \quad \text{if } \gamma \neq 1. \tag{2}$$

I discuss this requirement further below.

The Euler equation relates the price of an asset this period to the payoff next period. Expectations are calculated with respect to the measure perceived by the representative agent:

$$P_0 = \mathbb{E}_0\left(e^{-\rho}\left(\frac{C_1}{C_0}\right)^{-\gamma}(D_1+P_1)\right).$$

Iterating forward, we get

$$P_0 = \mathbb{E}\left(\sum_{t=1}^{T} e^{-\rho t}\left(\frac{C_t}{C_0}\right)^{-\gamma} D_t\right) + \mathbb{E}_0 e^{-\rho T}\left(\frac{C_T}{C_0}\right)^{-\gamma} P_T\,.$$

Finally, allowing $T \to \infty$ (and imposing the no-bubble condition that the second term in the above expression tends to zero in the limit) leads to the familiar equation

$$P(D) = \mathbb{E}\left(\sum_{t=1}^{\infty} e^{-\rho t}\left(\frac{C_t}{C_0}\right)^{-\gamma} D_t\right). \tag{3}$$

I start by considering an asset which pays dividend stream $D_t \equiv (C_t)^{\lambda}$ for some constant λ (the λ-asset). The central cases of interest will later be $\lambda = 0$ (the *riskless bond*) and $\lambda = 1$ (the *wealth portfolio* which pays consumption as its dividend), but, as in Campbell (1986) and Abel (1999), it is possible to view values $\lambda > 1$ as a tractable way of modelling levered claims. I write P_λ for the price of this asset at time 0, and D_λ for the dividend at time 0.

From (3),

$$
\begin{aligned}
P_\lambda &= \mathbb{E}\left(\sum_{t=1}^{\infty} e^{-\rho t}\left(\frac{C_t}{C_0}\right)^{-\gamma}(C_t)^\lambda\right) \\
&= (C_0)^\lambda \sum_{t=1}^{\infty} e^{-\rho t}\mathbb{E}\left(\left(\frac{C_t}{C_0}\right)^{\lambda-\gamma}\right) \\
&= D_\lambda \sum_{t=1}^{\infty} e^{-\rho t}\mathbb{E}\left(e^{(\lambda-\gamma)G_t}\right) \\
&= D_\lambda \sum_{t=1}^{\infty} e^{-\rho t}\left(\mathbb{E}\left(e^{(\lambda-\gamma)G}\right)\right)^t . \qquad (4)
\end{aligned}
$$

The last equality follows from the assumption that log consumption growth is i.i.d. To make further progress, I now introduce a pair of definitions.

Definition 1. *Given some arbitrary random variable, G, the* moment-generating function $\boldsymbol{m}(\theta)$ *and* cumulant-generating function *or* CGF $\boldsymbol{c}(\theta)$ *are defined by*

$$
\begin{aligned}
\boldsymbol{m}(\theta) &\equiv \mathbb{E}\exp(\theta G) \qquad (5) \\
\boldsymbol{c}(\theta) &\equiv \log \boldsymbol{m}(\theta)\,, \qquad (6)
\end{aligned}
$$

for all θ for which the expectation in (5) is finite.

Here, G is an annual increment of log consumption, $G = \log C_{t+1} - \log C_t$. Notice that $\boldsymbol{c}(0) = 0$ for any growth process and that $\boldsymbol{c}(1)$ is equal to log mean gross consumption growth, so $\boldsymbol{c}(1) \approx 2\%$. The CGF summarizes information about the cumulants (or, equivalently, moments) of G.[4] We can expand $\boldsymbol{c}(\theta)$ as a power series in θ,

$$
\boldsymbol{c}(\theta) = \sum_{n=1}^{\infty} \frac{\kappa_n \theta^n}{n!},
$$

and define κ_n to be the nth *cumulant* of log consumption growth. A small amount of algebra confirms that, for example, $\kappa_1 \equiv \mu$ is the mean, $\kappa_2 \equiv \sigma^2$ the variance, κ_3/σ^3 the skewness and κ_4/σ^4 the kurtosis of log consumption growth. Knowledge of the cumulants of a random variable implies knowledge of the moments, and vice versa.

[4] See Appendix A for further details.

With this definition, (4) becomes

$$\begin{aligned} P_\lambda &= D_\lambda \sum_{t=1}^{\infty} e^{-[\rho - \boldsymbol{c}(\lambda-\gamma)]t} \\ &= D_\lambda \cdot \frac{e^{-[\rho - \boldsymbol{c}(\lambda-\gamma)]}}{1 - e^{-[\rho - \boldsymbol{c}(\lambda-\gamma)]}}. \end{aligned}$$

It is convenient to define the log dividend yield $d_\lambda/p_\lambda \equiv \log(1 + D_\lambda/P_\lambda)$.[5] Then,

$$d_\lambda/p_\lambda = \rho - \boldsymbol{c}(\lambda - \gamma) \tag{7}$$

Two special cases are of particular interest. The first is $\lambda = 0$, in which case the asset in question is the riskless bond, whose dividend yield is the riskless rate. The second is $\lambda = 1$, in which case the asset pays consumption as its dividend, and can therefore be interpreted as aggregate wealth. The dividend yield is then the consumption-wealth ratio.

This calculation also shows that the necessary restriction on consumption growth for the expected utility to be well defined in (2) is that $\rho > \boldsymbol{c}(1 - \gamma)$, or equivalently that the consumption-wealth ratio is positive. When the condition fails, the standard consumption-based asset pricing approach is no longer valid.

The gross return on the λ-asset is (dropping λ subscripts for clarity)

$$\begin{aligned} 1 + R_{t+1} &= \frac{D_{t+1} + P_{t+1}}{P_t} \qquad (8) \\ &= \frac{P_{t+1}}{P_t}\left(1 + \frac{D_{t+1}}{P_{t+1}}\right) \\ &= \frac{D_{t+1}}{D_t}\left(e^{\rho - \boldsymbol{c}(\lambda-\gamma)}\right) \end{aligned}$$

and thus the expected gross return is

$$\begin{aligned} 1 + \mathbb{E}R_{t+1} &= \mathbb{E}\left(\left(\frac{C_{t+1}}{C_t}\right)^{\lambda}\right) \cdot e^{\rho - \boldsymbol{c}(\lambda-\gamma)} \\ &= \mathbb{E}\left(e^{G\lambda}\right) \cdot e^{\rho - \boldsymbol{c}(\lambda-\gamma)} \\ &= e^{\rho - \boldsymbol{c}(\lambda-\gamma) + \boldsymbol{c}(\lambda)} \end{aligned}$$

[5]It is worth emphasizing that log dividend yield, as I have defined it, is a number close to D/P, since $\log(1 + x) \approx x$ for small x. d/p is *not* the same as $d - p$ as used elsewhere in the literature to mean $\log D/P$.

Once again, it turns out to be more convenient to work with log expected gross return, $er_\lambda \equiv \log(1+\mathbb{E}R_{t+1}) = \rho + \boldsymbol{c}(\lambda) - \boldsymbol{c}(\lambda-\gamma)$.

Proposition 1 (Fundamental quantities, power utility case)**.** *The riskless rate, $r_f \equiv \log(1+R_f)$, consumption-wealth ratio, $c/w \equiv \log(1+C/W)$, and risk premium on aggregate wealth, $rp \equiv er_1 - r_f$, are given by*

$$r_f = \rho - \boldsymbol{c}(-\gamma) \tag{9}$$

$$c/w = \rho - \boldsymbol{c}(1-\gamma) \tag{10}$$

$$rp = \boldsymbol{c}(1) + \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\,. \tag{11}$$

Writing these quantities explicitly in terms of the underlying cumulants by expanding $\boldsymbol{c}(\theta)$ in power series form, we obtain

$$r_f = \rho - \sum_{n=1}^{\infty} \frac{\kappa_n(-\gamma)^n}{n!} \tag{12}$$

$$c/w = \rho - \sum_{n=1}^{\infty} \frac{\kappa_n(1-\gamma)^n}{n!} \tag{13}$$

$$rp = \sum_{n=2}^{\infty} \frac{\kappa_n}{n!} \cdot \Big\{1 + (-\gamma)^n - (1-\gamma)^n\Big\}. \tag{14}$$

Writing the first few terms of the series out more explicitly, (12) implies that

$$r_f = \rho + \kappa_1\gamma - \frac{\kappa_2}{2}\gamma^2 + \frac{\kappa_3}{3!}\gamma^3 - \frac{\kappa_4}{4!}\gamma^4 + \text{higher order terms}\,.$$

By definition of the first four cumulants, this can be rewritten as

$$r_f = \rho + \mu\gamma - \frac{1}{2}\sigma^2\gamma^2 + \frac{\text{skewness}}{3!}\sigma^3\gamma^3 - \frac{\text{excess kurtosis}}{4!}\sigma^4\gamma^4 + \text{higher order terms}\,. \tag{15}$$

In the lognormal case, the skewness, excess kurtosis and all higher cumulants are zero, so (15) reduces to the familiar $r_f = \rho + \mu\gamma - \sigma^2\gamma^2/2$. More generally, the riskless rate is low if mean log consumption growth μ is low (an intertemporal substitution effect); if the variance of log consumption growth σ^2 is high (a precautionary savings effect); if there is negative skewness; or if there is a high degree of kurtosis.

Similarly, the consumption-wealth ratio (13) can be rewritten as

$$\begin{aligned} c/w &= \rho + \mu(\gamma - 1) - \frac{1}{2}\sigma^2(\gamma-1)^2 + \frac{\text{skewness}}{3!}\sigma^3(\gamma-1)^3 - \\ &\quad - \frac{\text{excess kurtosis}}{4!}\sigma^4(\gamma-1)^4 + \text{higher order terms}\,. \end{aligned} \tag{16}$$

In the log utility case, $\gamma = 1$, the consumption-wealth ratio is determined only by the rate of time preference: $c/w = \rho$. If $\gamma \neq 1$, the consumption-wealth ratio is low when cumulants of even order are large (high variance, high kurtosis, and so on). The importance of cumulants of odd order depends on whether γ is greater or less than 1. In the empirically more plausible case $\gamma > 1$, the consumption-wealth ratio is low when odd cumulants are low: when mean log consumption growth is low, or when there is negative skewness, for example. If the representative agent is more risk-tolerant than log, the reverse is true: the consumption-wealth ratio is high when mean log consumption growth is low, or when there is negative skewness.

The risk premium (14) becomes

$$\begin{aligned} rp &= \gamma\sigma^2 + \frac{\text{skewness}}{3!}\sigma^3\left(1 - \gamma^3 - (1-\gamma)^3\right) + \\ &\quad + \frac{\text{excess kurtosis}}{4!}\sigma^4\left(1 + \gamma^4 - (1-\gamma)^4\right) + \text{higher order terms}\,. \end{aligned} \tag{17}$$

In the lognormal case, this is just $rp = \gamma\sigma^2$. Since $1 + \gamma^n - (1-\gamma)^n > 0$ for even n, the risk premium is increasing in variance, excess kurtosis and higher cumulants of even order. The effect of skewness and higher cumulants of odd order depends on γ. For odd n, $1 - \gamma^n - (1-\gamma)^n$ is positive if $\gamma < 1$, zero if $\gamma = 1$, and negative if $\gamma > 1$. If $\gamma = 1$, skewness and higher odd-order cumulants have no effect on the risk premium. Otherwise, the risk premium is decreasing in skewness and higher odd cumulants if $\gamma > 1$ and increasing if $\gamma < 1$.

The following result generalizes Proposition 1 to allow for Epstein-Zin preferences.

Proposition 2 (Epstein-Zin case)**.** *Defining* $\vartheta \equiv (1-\gamma)/(1-1/\psi)$*, we have*

$$r_f = \rho - \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\left(\frac{1}{\vartheta} - 1\right) \tag{18}$$

$$c/w = \rho - \boldsymbol{c}(1-\gamma)/\vartheta \tag{19}$$

$$rp = \boldsymbol{c}(1) + \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\,, \tag{20}$$

and the counterparts of (12)–(14) *that result on expanding* (18)–(20) *as power series.*

Proof. See Appendix B. □

Equation (19) reveals that if $\gamma > 1$ and $\psi > 1$—as in the calibration of Bansal and Yaron (2004)—so $\vartheta < 0$, the comparative statics for the consumption-wealth ratio that were discussed above are reversed: a high variance or kurtosis leads to a high consumption-wealth ratio.

Equation (20) shows as expected that when the CGF is linear—that is, when consumption growth is deterministic—there is no risk premium. Roughly speaking, the CGF of the driving consumption process must have a significant amount of convexity over the range $[-\gamma, 1]$ to generate an empirically reasonable risk premium. It also confirms that risk aversion alone influences the risk premium: the elasticity of intertemporal substitution is not a factor.

Expressions (12)–(14), and their analogues in the Epstein-Zin case, can in principle be estimated directly by estimating the cumulants of log consumption, given a sufficiently long data sample, without imposing any further structure on the model. If, say, the high equity premium results from the occasional occurrence of severe disasters, this will show up in the cumulants. Other than (A1) and (A2), no assumption need be made about the arrival rate or distribution of disasters, nor of any other feature of the consumption process.

In practice, of course, we cannot estimate infinitely many cumulants from a finite data set. One solution to this is to impose some particular distribution on log consumption growth, and then to estimate the parameters of the distribution. An alternative approach, more in the spirit of model-independence, is to approximate the equations by truncating after the first N cumulants, N being determined by the amount of data available. (In this context it is worth noting that the assumption that consumption growth is lognormal is equivalent to truncating at $N = 2$, since, as noted above, when log consumption growth is Normal all cumulants above the variance are equal to zero—that is, $\kappa_n = 0$ for n greater than 2.) For reasons given in the Introduction and discussed further in Section II.A below, I do not follow this route.

I.A The Gordon growth model

From equations (18)–(20), we see that

$$c/w = rp + r_f - \boldsymbol{c}(1). \tag{21}$$

This is a version of the traditional Gordon growth model. For example, the last term of (21), $\boldsymbol{c}(1) = \log \mathbb{E} C_{t+1}/C_t$, measures mean consumption growth. Only three of the riskless rate, risk premium, consumption-wealth ratio and mean consumption growth can be independently specified: the fourth is mechanically determined by (21).

This observation, together with equations (18)–(20), provides another way to look at Kocherlakota's (1990) point. In principle, given sufficient asset price and consumption data, we could determine the riskless rate, the risk premium, and CGF $\boldsymbol{c}(\cdot)$ to any desired level of accuracy. Since γ is the only preference parameter that determines the risk premium, it could be calculated from (20), given knowledge of $\boldsymbol{c}(\cdot)$. On the other hand, knowledge of the riskless rate leaves ρ and ψ indeterminate in equation (18), even given knowledge of γ and $\boldsymbol{c}(\cdot)$. That is, the time discount rate and elasticity of intertemporal substitution cannot be disentangled on the basis of the four fundamental quantities alone. On the other hand, as noted in the introduction, the use of Epstein-Zin preferences aids the interpretation of results and allows for a richer set of possible comparative statics.

I.B The asymptotic lognormality of consumption

If G has mean μ and (finite) variance σ^2, the central limit theorem shows that consumption is asymptotically lognormal:[6] as $t \to \infty$

$$\frac{G_t - \mu t}{\sqrt{t}} \xrightarrow{d} N(0, \sigma^2).$$

[6]Informally, $G_t - \mu t$ is typically $O(\sqrt{t})$, so for positive α, $\mathbb{P}(G_t - \mu t \geq \alpha t) \to 0$ as $t \to \infty$, or equivalently, $\mathbb{P}(C_t \geq C_0 e^{(\mu+\alpha)t}) \to 0$. The Cramér-Chernoff theorem tells us how fast this probability decays to zero, and provides an opportunity to mention another context in which the CGF arises. It implies that

$$\frac{1}{t} \log \mathbb{P}\left(C_t \geq C_0 e^{\alpha t}\right) \longrightarrow \inf_{\theta \geq 0} \boldsymbol{c}(\theta) - \alpha\theta$$

It therefore appears that if one measures over very long periods, only the first two cumulants will be needed to capture information about consumption growth. Why, then, does the representative agent care about cumulants of log consumption growth other than mean and variance? To answer this question, it is helpful to define the scale-free cumulants

$$SFC_n \equiv \frac{\kappa_n}{\sigma^n}$$

For example, SFC_3 is skewness and SFC_4 is kurtosis. These scale-free cumulants are normalized to be invariant if the underlying random variable is scaled by some constant factor. Since the (unscaled) cumulants of G_t are linear in t, the nth scale-free cumulant of G_t is proportional to $t \cdot t^{-n/2} = t^{(2-n)/2}$ and so tends to zero for n greater than 2. The asymptotic Normality of $(G_t - \mu t)/\sqrt{t}$ is reflected in the fact that its scale-free cumulants of orders greater than two tend to zero as t tends to infinity. But in terms of the scale-free cumulants, the riskless rate (for example) can be expressed as

$$\begin{aligned} r_f &= \rho - \sum_{n=1}^{\infty} \frac{\kappa_n(-\gamma)^n}{n!} \\ &= \rho - \sum_{n=1}^{\infty} \frac{SFC_n \sigma^n (-\gamma)^n}{n!} \end{aligned} \tag{22}$$

Thus, even though skewness, kurtosis and higher scale-free cumulants tend to zero as the period length is allowed to increase, the relevant asset-pricing equation scales these variables by σ—and this tends to infinity as period length increases, in such a way that higher cumulants remain relevant.

II The continuous-time case

For the purposes of constructing concrete examples, it is convenient to confirm that the simplicity of the above framework carries over to the continuous-time case.

Assumptions A1 and A2 are modified slightly. They become

---

and

$$\frac{1}{t} \log \mathbb{P}\left(C_t \le C_0 e^{\alpha t}\right) \longrightarrow \inf_{\theta \le 0} \boldsymbol{c}(\theta) - \alpha\theta.$$

Van der Vaart (1998) has a proof.

A1c There is a representative agent with constant relative risk aversion γ, who therefore maximizes[7]

$$\mathbb{E}\int_{t=0}^{\infty} e^{-\rho t}\frac{C_t^{1-\gamma}}{1-\gamma} \quad \text{if } \gamma \neq 1\,, \quad \text{or} \quad \mathbb{E}\int_{t=0}^{\infty} e^{-\rho t}\log C_t \quad \text{if } \gamma = 1 \tag{23}$$

A2c The log consumption path, G_t, of the representative agent follows a Lévy process, and $\boldsymbol{m}(\theta)$ exists for θ in $[-\gamma, 1]$.

The analysis is almost identical to that in the discrete-time case, using the fact that

$$\mathbb{E}e^{\theta G_t} = \left(\mathbb{E}e^{\theta G}\right)^t, \tag{24}$$

which follows from Assumption A2c; see Sato (1999) for detailed discussion of Lévy processes.

The following proposition shows that the discrete-time results go through almost unchanged, except that the equations that previously held for log dividend yields, the log riskless rate and the log risk premium now apply to the instantaneous dividend yield, the instantaneous riskless rate and the instantaneous risk premium. The proof, which is very similar to the discrete-time calculations, is omitted.

Proposition 3 (Reprise of earlier results)**.** *The riskless rate, R_f, consumption-wealth ratio, C/W, and risk premium on aggregate wealth, $RP \equiv ER_1 - R_f$, are given by*

$$\begin{aligned} R_f &= \rho - \boldsymbol{c}(-\gamma) \\ C/W &= \rho - \boldsymbol{c}(1-\gamma) \\ RP &= \boldsymbol{c}(1) + \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\,. \end{aligned}$$

The Gordon growth model holds:

$$D_\lambda/P_\lambda = ER_\lambda - \boldsymbol{c}(\lambda).$$

[7]For simplicity, I restrict to the power utility case, although the analysis could be easily generalized to allow for the continuous-time analogue of Epstein-Zin preferences (Duffie and Epstein (1992)).

II.A A concrete example: disasters

To aid intuition, it is helpful to demonstrate the above results in the context of a particular model. In this section, I show how to derive a convenient continuous-time version of Barro (2006a). I use the model to show that i.i.d. disaster models make predictions for the fundamentals that are sensitively dependent on the parameter values assumed. In particular, making disasters more frequent or more severe drives the riskless rate down sharply.

Suppose that log consumption follows the jump-diffusion process

$$G_t = \widetilde{\mu} t + \sigma_B B_t + \sum_{i=1}^{N(t)} Y_i \tag{25}$$

where B_t is a standard Brownian motion, $N(t)$ is a Poisson counting process with parameter ω and Y_i are i.i.d. random variables with some arbitrary distribution. The significance of this example is that any Lévy process can be approximated arbitrarily accurately by a process of the form (25). I will assume that all moments of the disaster size Y_1 are finite, from which it follows that all moments of G are finite.

The CGF is $\boldsymbol{c}(\theta) = \log \boldsymbol{m}(\theta)$, where

$$\begin{aligned} \boldsymbol{m}(\theta) &= \mathbb{E} e^{\theta G_1} \\ &= e^{\widetilde{\mu}\theta} \cdot \mathbb{E} e^{\sigma_B \theta B_1} \cdot \mathbb{E} e^{\theta \sum_{i=1}^{N(1)} Y_i} ; \end{aligned}$$

separating the expectation into two separate products is legitimate since the Poisson jumps and Y_i are independent of the Brownian component B_t. The middle term is the expectation of a lognormal random variable: $\mathbb{E} e^{\theta \sigma_B B_1} = e^{\sigma_B^2 \theta^2 / 2}$. The final term is slightly more complicated, but can be evaluated by conditioning on the number of jumps that take place before $t = 1$:

$$\begin{aligned} \mathbb{E} \exp\left\{ \theta \sum_{i=1}^{N(1)} Y_i \right\} &= \sum_0^\infty \frac{e^{-\omega} \omega^n}{n!} \mathbb{E} \exp\left\{ \theta \sum_1^n Y_i \right\} \\ &= \sum_0^\infty \frac{e^{-\omega} \omega^n}{n!} \left[\mathbb{E} \exp\left\{ \theta Y_1 \right\} \right]^n \\ &= \exp\left\{ \omega \left(\boldsymbol{m}_{Y_1}(\theta) - 1 \right) \right\} , \end{aligned}$$

So, finally, we have

$$\boldsymbol{c}(\theta) = \widetilde{\mu}\theta + \sigma_B^2\theta^2/2 + \omega\left(\boldsymbol{m}_{Y_1}(\theta) - 1\right) . \tag{26}$$

The cumulants can be read off from (26):

$$\begin{aligned}\kappa_n(G) &= \boldsymbol{c}^{(n)}(0) \\ &= \begin{cases} \widetilde{\mu} + \omega\,\mathbb{E}Y & n = 1 \\ \sigma_B^2 + \omega\,\mathbb{E}Y^2 & n = 2 \\ \omega\,\mathbb{E}Y^n & n \geq 3 \end{cases}\end{aligned}$$

Take the case in which $Y \sim N(-b, s^2)$; b is assumed to be greater than zero, so the jumps represent disasters. The CGF is then

$$\boldsymbol{c}(\theta) = \widetilde{\mu}\theta + \frac{1}{2}\sigma_B^2\theta^2 + \omega\left(e^{-\theta b + \frac{1}{2}\theta^2 s^2} - 1\right) . \tag{27}$$

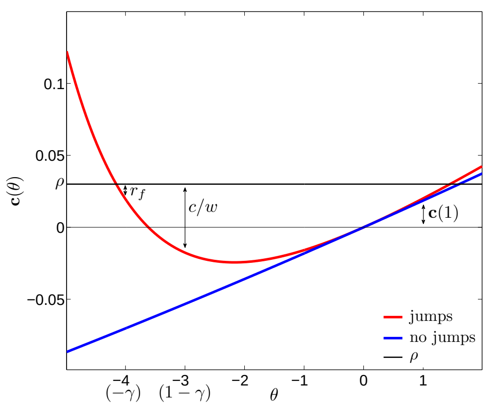

Figure 1: The CGF in equation (27) shown with and without ($\omega = 0$) jumps. The figure assumes that $\gamma = 4$.

Figure 1 shows the CGF of (27) plotted against θ. I set parameters which correspond to Barro's (2006a) baseline calibration—$\gamma = 4, \sigma_B = 0.02, \rho = 0.03, \widetilde{\mu} = 0.025, \omega = 0.017$—and choose $b = 0.39$ and $s = 0.25$ to match the mean and variance of the distribution of jumps used in the same paper. I also plot the CGF that results in the absence of jumps ($\omega = 0$). In the latter case, I adjust the drift of consumption growth to keep mean log consumption growth constant.

The riskless rate, consumption-wealth ratio and mean consumption growth can be read directly off the graph, as indicated by the arrows. The risk premium can be calculated from these three via the Gordon growth formula ($rp = c/w + \boldsymbol{c}(1) - r_f$), or read directly off the graph as follows. Draw one line from $(-\gamma, \boldsymbol{c}(-\gamma))$ to $(1, \boldsymbol{c}(1))$ and another from $(1-\gamma, \boldsymbol{c}(1-\gamma))$ to $(0, 0)$. The midpoint of the first line lies above the midpoint of the second by convexity of the CGF. The risk premium is twice the distance from one midpoint to the other. This procedure is illustrated in Figure 2.

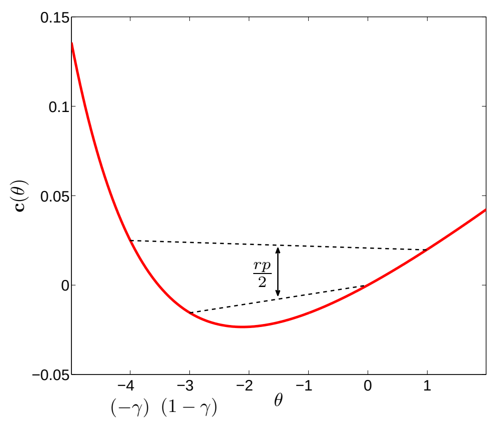

Figure 2: The risk premium. The figure assumes that $\gamma = 4$.

The standard lognormal model predicts a counterfactually high riskless rate; in Figure 1, this is reflected in the fact that the no-jumps CGF lies well below ρ for reasonable values of θ. Similarly, the standard lognormal model predicts a counterfactually low equity premium. In Figure 1, this manifests itself in a no-jump CGF which is practically linear over the relevant range and which is upward-sloping between $-\gamma$ and $1-\gamma$. Conversely, the disaster CGF has a shape which allows it to match observed fundamentals closely.

Zooming out on Figure 1, we obtain Figure 3, which further illustrates the equity premium and riskless rate puzzles. With jumps, the CGF is visible at the right-hand side of the figure; the CGF explodes so quickly as θ declines that it is only visible for θ greater than about -5. The jump-free lognormal CGF has incredibly low curvature. For a realistic riskless rate and equity premium, the model requires a risk aversion above

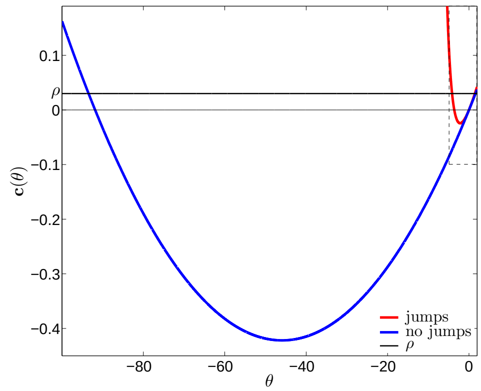

Figure 3: Zooming out to see the equity premium and riskless rate puzzles. The dashed box in the upper right-hand corner is the boundary of the region plotted in Figure 1.

80.

With the explicit expression (27) for the CGF in hand, it is easy to investigate the sensitivity of a disaster model's predictions to the parameter values assumed. Table I shows how changes in the calibration of the distribution of disasters affect the relevant fundamentals and the cost of consumption uncertainty, ϕ. I consider the power utility case, with $\rho = 0.03$ and $\gamma = 4$, and also the Epstein-Zin case, with the same time-preference rate and risk aversion but higher elasticity of intertemporal substitution, $\psi = 1.5$, following Bansal and Yaron (2004). As is evident from the table, small changes in any of the parameters ω, b or s have large effects on the equity premium (and, with power utility, on the riskless rate; this effect is muted in the Epstein-Zin case). Given that these parameters are hard to estimate—disasters happen very rarely—this is problematic.

An important difference between the power utility and Epstein-Zin cases is that increasing risk in the economy leads to lower consumption-wealth ratios in the power utility case, and higher consumption-wealth ratios in the Epstein-Zin case. Bansal and Yaron (2004) emphasize this feature in a model without disasters, and argue that it supports an elasticity of substitution greater than 1.

Table II investigates the consequences of truncating the CGF at the nth cumulant.

	ω	b	s	R_f	C/W	RP	R_f^*	C/W^*	RP^*
Baseline case	**0.017**	**0.39**	**0.25**	**1.0**	**4.8**	**5.7**	**-0.9**	**2.8**	**5.7**
High ω	0.022			-2.4	3.1	7.4	-2.5	3.0	7.4
Low ω	0.012			4.5	6.4	4.1	0.7	2.6	4.1
High b		0.44		-1.9	3.6	7.5	-2.6	2.9	7.5
Low b		0.34		3.5	5.8	4.4	0.4	2.7	4.4
High s			0.30	-2.2	3.8	8.1	-3.1	2.9	8.1
Low s			0.20	3.2	5.5	4.2	0.5	2.7	4.2

Table I: The impact of different assumptions about the distribution of disasters. $\widetilde{\mu} = 0.025$, $\sigma = 0.02$. Unasterisked group assumes power utility, $\rho = 0.03$, $\gamma = 4$. Asterisked group assumes Epstein-Zin preferences, $\rho = 0.03$, $\gamma = 4$, $\psi = 1.5$.

(The risk premium calculation applies in either the power utility or Epstein-Zin case; the riskless rate and consumption-wealth ratio calculations only apply in the power utility case.) When $n = 2$, this is equivalent to making a lognormality assumption, as noted above. With $n = 3$, it can be thought of as an approximation which accounts for the influence of skewness; $n = 4$ also allows for kurtosis. As is clear from the table, however, much of the action is due to cumulants of fifth order and higher. This suggests that one should not expect calculations based on third- or fourth-order approximations to capture fully the influence of disasters.

III Restrictions on preference parameters

Any three of the riskless rate, consumption-wealth ratio, risk premium and expected consumption growth pin down the value of the fourth, via the Gordon growth model $c/w = r_f + rp - \boldsymbol{c}(1)$ given in (21). I now assume that these quantities are *observable*, and take the values given in Table III.

We have seen, too, that the riskless rate, risk premium and consumption-wealth ratio tell us information about the shape of the CGF. I now show how to exploit this observa-

n	R_f	C/W	RP	
1	10.3	8.5	0.0	deterministic
2	7.1	6.7	1.6	lognormal
3	4.7	5.7	3.0	
4	3.0	5.1	4.1	
∞	1.0	4.8	5.7	true model

Table II: The impact of approximating the disaster model by truncating at the nth cumulant. All parameters as in baseline power utility case of Table I.

riskless rate	r_f	0.02
risk premium	rp	0.06
consumption-wealth ratio	c/w	0.06
mean consumption growth	$\boldsymbol{c}(1)$	0.02

Table III: Assumed values of the observables.

tion to find restrictions on preference parameters, in terms of observable fundamentals, that must hold in *any* Epstein-Zin/i.i.d. model, no matter what pattern of (say) rare disasters we allow ourselves to entertain.

Since for example $r_f = \rho - \boldsymbol{c}(-\gamma)$ in the power utility case, observation of the riskless rate tells us something about ρ and something about the value taken by the CGF at $-\gamma$. Similarly, observation of the consumption-wealth ratio tells us something about ρ and something about the value taken by the CGF at $1-\gamma$. Next, $\boldsymbol{c}(1) = \log \mathbb{E}(C_1/C_0)$ is pinned down by mean consumption growth, and $\boldsymbol{c}(0) = 0$ by definition. How, though, can we get control on the enormous range of possible consumption processes? One approach is to exploit the fact that the CGF of any random variable is convex:

Fact 1. *CGFs are convex.*

Proof. Since $\boldsymbol{c}(\theta) = \log \boldsymbol{m}(\theta)$, we have

$$\begin{aligned} \boldsymbol{c}''(\theta) &= \frac{\boldsymbol{m}(\theta) \cdot \boldsymbol{m}''(\theta) - \boldsymbol{m}'(\theta)^2}{\boldsymbol{m}(\theta)^2} \\ &= \frac{\mathbb{E}e^{\theta G}\mathbb{E}G^2 e^{\theta G} - \left(\mathbb{E}Ge^{\theta G}\right)^2}{\boldsymbol{m}(\theta)^2}\,. \end{aligned}$$

The numerator of this expression is positive by a version of the Cauchy-Schwartz inequality which states that $\mathbb{E}X^2 \cdot \mathbb{E}Y^2 \geq \mathbb{E}(|XY|)^2$ for any random variables X and Y. In this case, we need to set $X = e^{\theta G/2}$ and $Y = Ge^{\theta G/2}$. (See Billingsley (1995), for further discussion of CGFs.) □

The following result exploits this convexity to derive bounds on preference parameters, based on the observables, that are valid no matter what is going on in the higher cumulants.

Proposition 4. *In the power utility case, we have*

$$r_f - c/w \leq \frac{c/w - \rho}{\gamma - 1} \leq rp + r_f - c/w \tag{28}$$

In the Epstein-Zin case, we have

$$r_f - c/w \leq \frac{c/w - \rho}{1/\psi - 1} \leq rp + r_f - c/w \tag{29}$$

Proof. From equation (19) we have, in the Epstein-Zin case,

$$\frac{c/w - \rho}{1/\psi - 1} = \frac{\boldsymbol{c}(1-\gamma)}{1-\gamma}\,.$$

The convexity of $\boldsymbol{c}(\theta)$ and the fact that $\boldsymbol{c}(0) = 0$ imply that

$$\frac{\boldsymbol{c}(-\gamma)}{-\gamma} \leq \frac{\boldsymbol{c}(1-\gamma)}{1-\gamma} \leq \boldsymbol{c}(1)\,;$$

to see this, note that if $f(\theta)$ is a convex function passing through zero, then $f(\theta)/\theta$ is increasing. Putting the two facts together, we have

$$\frac{\boldsymbol{c}(-\gamma)}{-\gamma} \leq \frac{c/w - \rho}{1/\psi - 1} \leq \boldsymbol{c}(1)\,.$$

After some rearrangement of the left-hand inequality using (18) and (19), and substituting out for $\boldsymbol{c}(1)$ using the Gordon growth model (21), we get (29). Equation (28) follows since $\gamma = 1/\psi$ in the power utility case. □

The intuition for the result is that as ψ approaches one, the consumption-wealth ratio approaches ρ. Therefore, when the consumption-wealth ratio is far from ρ, ψ must be far from one. Using the empirically reasonable values $rp = 6\%, r_f = 2\%, c/w = 6\%$, we have the restriction that $-0.04 \leq (0.06 - \rho)/(1/\psi - 1) \leq 0.02$, or equivalently

$$4 - \frac{1}{\psi} \leq 50\rho \leq 1 + \frac{2}{\psi} \quad \text{if} \quad \psi \leq 1$$
$$1 + \frac{2}{\psi} \leq 50\rho \leq 4 - \frac{1}{\psi} \quad \text{if} \quad \psi \geq 1 \, .$$

Figures 4a and 4b illustrate these constraints. If ψ is greater than one, ρ is constrained to lie between 0.02 and 0.08; if also ψ is less than two, ρ must lie between 0.04 and 0.07.

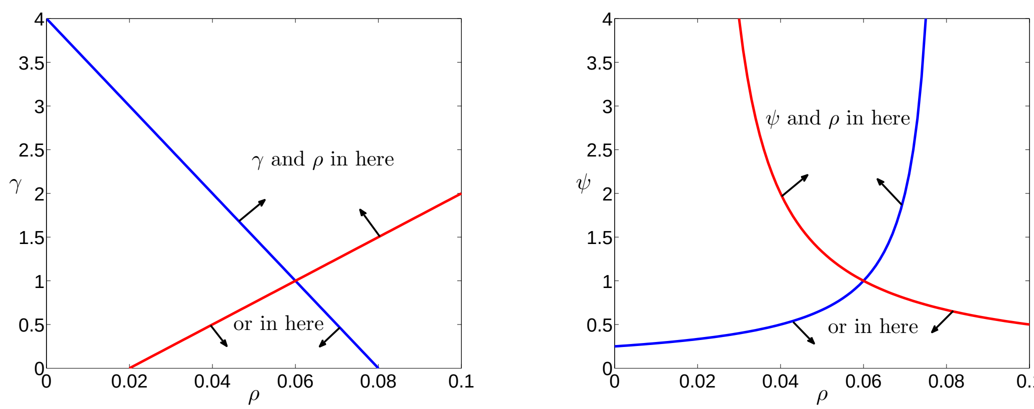

(a) Power utility case: γ and ρ (b) Epstein-Zin case: ψ and ρ

Figure 4: Parameter restrictions for i.i.d. models with $rp = 6\%$, $r_f = 2\%$ and log expected consumption growth of 2%.

III.A Hansen-Jagannathan and good-deal bounds

The restrictions in Proposition 4 are complementary to the bound derived by Hansen and Jagannathan (1991), which relates the standard deviation and mean of the stochastic discount factor, M, to the Sharpe ratio on an arbitrary asset, SR:

$$SR \leq \frac{\sigma(M)}{\mathbb{E}M}\,. \tag{30}$$

In the Epstein-Zin-i.i.d. setting, the right-hand side of (30) becomes

$$\begin{aligned}\frac{\sigma(M)}{\mathbb{E}M} &= \sqrt{\frac{\mathbb{E}M^2}{(\mathbb{E}M)^2}-1} \\ &= \sqrt{e^{\boldsymbol{c}(-2\gamma)-2\boldsymbol{c}(-\gamma)}-1}\,; \end{aligned} \tag{31}$$

combining (30) and (31), we obtain a Hansen-Jagannathan bound in CGF notation:

$$\log\left(1+SR^2\right) \leq \boldsymbol{c}(-2\gamma)-2\boldsymbol{c}(-\gamma)\,. \tag{32}$$

Cochrane and Saá-Requejo (2000) observe that inequality (30) suggests a natural way to restrict asset-pricing models. Suppose $\sigma(M)/\mathbb{E}M \leq h$; then (30) implies that the maximal Sharpe ratio is less than h. In CGF notation, the good-deal bound is that

$$\boldsymbol{c}(-2\gamma)-2\boldsymbol{c}(-\gamma) \leq \log\left(1+h^2\right) \tag{33}$$

Suppose, for example, that we wish to impose the restriction that Sharpe ratios above 100% are too good a deal to be available. Then the good-deal bound is $\boldsymbol{c}(-2\gamma) - 2\boldsymbol{c}(-\gamma) \leq \log 2$. This expression can be evaluated under particular parametric assumptions about the consumption process. In the case in which consumption growth is lognormal, with volatility of log consumption equal to σ, it supplies an upper bound on risk aversion: $\gamma \leq \sqrt{\log 2}/\sigma$ (which is about 42 if $\sigma = 0.02$). However, this upper bound is rather weak, and in any case the postulated consumption process is inconsistent with observed features of asset markets such as the high equity premium and low riskless rate.

Alternatively, one might model the consumption process as subject to disasters in the sense of Section II.A. In this case, the good-deal bound implies tighter restrictions on γ, but these restrictions are sensitively dependent on the disaster parameters.

In order to progress from (33) to a bound on γ and ρ which does not require parametrization of the consumption process, we want to relate $\boldsymbol{c}(-2\gamma)-2\boldsymbol{c}(-\gamma)$ to quantities which can be directly observed. For example, the Hansen-Jagannathan bound (32) improves on a conclusion which follows from the convexity of the CGF, namely, that

$$0 \leq \boldsymbol{c}(-2\gamma) - 2\boldsymbol{c}(-\gamma)\,. \tag{34}$$

This trivial inequality follows by considering the value of the CGF at the three points $\boldsymbol{c}(0)$, $\boldsymbol{c}(-\gamma)$, and $\boldsymbol{c}(-2\gamma)$. Convexity implies that the average slope of the CGF is more negative (or less positive) between -2γ and $-\gamma$ than it is between $-\gamma$ and 0. To be precise, it implies that

$$\frac{\boldsymbol{c}(-\gamma) - \boldsymbol{c}(-2\gamma)}{\gamma} \leq \frac{\boldsymbol{c}(0) - \boldsymbol{c}(-\gamma)}{\gamma} \tag{35}$$

from which (34) follows immediately, given that $\boldsymbol{c}(0) = 0$. Combining (33) and (34), we obtain the (underwhelming!) result that

$$0 \leq \log\left(1 + h^2\right)\,.$$

However, we can sharpen (34) by comparing the slope of the CGF between -2γ and $-\gamma$ to the slope between $-\gamma$ and $1-\gamma$ (as opposed to that between $-\gamma$ and 0). Making this formal, we have by convexity of the CGF that

$$\frac{\boldsymbol{c}(-\gamma) - \boldsymbol{c}(-2\gamma)}{\gamma} \leq \frac{\boldsymbol{c}(1-\gamma) - \boldsymbol{c}(-\gamma)}{1}\,,$$

from which it follows that

$$\begin{aligned}\boldsymbol{c}(-2\gamma) - 2\boldsymbol{c}(-\gamma) &\geq (\gamma-1)\boldsymbol{c}(-\gamma) - \gamma\boldsymbol{c}(1-\gamma) \\ &= (\gamma-1)(c/w - r_f) + \vartheta(c/w - \rho)\end{aligned}$$

or equivalently

$$\frac{\sigma(M)}{\mathbb{E}M} \geq \sqrt{e^{(\gamma-1)(c/w-r_f)+\vartheta(c/w-\rho)} - 1}\,. \tag{36}$$

Proposition 5. *If the maximal Sharpe ratio is less than or equal to h, then we must have*

$$(\gamma-1)(c/w - r_f) + \vartheta(c/w - \rho) \leq \log\left(1 + h^2\right)\,. \tag{37}$$

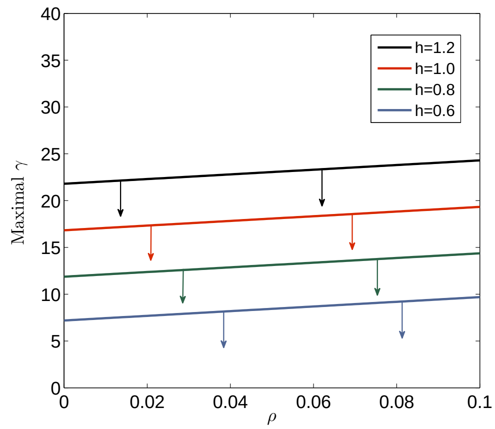

Figure 5: Restrictions on γ and ρ implied by good-deal bounds in the power utility case with $c/w = 0.06, r_f = 0.02$.

Working with the power utility case for simplicity ($\vartheta = 1$) and setting $c/w = 0.06, r_f = 0.02$, Figure 5 shows the upper bounds on γ that result for various different h. Lower values of h imply tighter restrictions. When $h = 1$—ruling out Sharpe ratios above 100%—we have $\gamma \leq 16.8 + 25\rho$. So if $\rho = 0.03$, $\gamma < 17.6$.

Alternatively, we could take the approach suggested at the end of the previous section, by setting $\rho = c/w$. In the general (Epstein-Zin) case, equation (37) then implies the restriction

$$\gamma \leq 1 + \frac{\log\left(1 + h^2\right)}{c/w - r_f}\,. \tag{38}$$

(To avoid unnecessary complication I have imposed the empirically relevant case $c/w \geq r_f$.) Setting $c/w = 0.06$, $r_f = 0.02$, and $h = 1$, this implies that $\gamma < 18.4$.

The important feature of the bounds (37) and (38) is that by exploiting the observable consumption-wealth ratio and riskless rate, they do not require one to take a stand on the hard-to-estimate higher cumulants of consumption growth.

IV The cost of consumption fluctuations

Continuing with the theme of extracting information from observable fundamentals, I now explore the implications of the consumption-wealth ratio for estimates of the cost of consumption fluctuations in the style of Lucas (1987), Obstfeld (1994) or Barro (2006b).

A starting point is the close correspondence between expected utility and the price of the consumption claim (that is, wealth):

$$U(\gamma) \equiv \mathbb{E}\left[\sum_{t=0}^{\infty} e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma}\right] \longleftrightarrow \mathbb{E}\left[\sum_{t=1}^{\infty} e^{-\rho t} \left(\frac{C_t}{C_0}\right)^{1-\gamma}\right] = \frac{W_0}{C_0}\,.$$

In fact we have

$$U(\gamma) = \frac{C_0^{1-\gamma}}{1-\gamma} \cdot \left(1 + \frac{W_0}{C_0}\right)\,. \tag{39}$$

This correspondence between expected utility and the consumption-wealth ratio, and hence (39), does not have a meaningful analogue in the log utility case. In a sense, the consumption-wealth ratio is less informative in the log utility case since it is pinned down by the time discount rate, $C/W = e^{\rho} - 1$.

Expected utility can also be expressed in terms of the CGF:

$$U(\gamma) = \frac{C_0^{1-\gamma}}{1-\gamma} \left(1 + \frac{1}{e^{\rho - \boldsymbol{c}(1-\gamma)} - 1}\right)\,, \qquad \gamma \neq 1\,. \tag{40}$$

When $\gamma < 1$ the representative agent prefers large values of $\boldsymbol{c}(1-\gamma)$ and when $\gamma > 1$ the representative agent prefers small values of $\boldsymbol{c}(1-\gamma)$. When $\gamma > 1$, the representative agent likes positive mean and positive skew and positive cumulants of odd orders but dislikes large values of variance, kurtosis and cumulants of even orders; when $\gamma < 1$ the representative agent likes large means, large variances, large skewness, large kurtosis—large positive values of cumulants of *all* orders.[8]

Equation (39) gives expected utility under the status quo; expression (40) permits the calculation of expected utility under alternative consumption processes with their corresponding CGFs. I compare two quantities: expected utility with initial consumption $(1+\phi)C_0$ and the status quo consumption growth process,[9] and expected utility with initial consumption C_0 and the alternative consumption growth process. The *cost*

[8] As always, these cumulants are the cumulants of *log* consumption growth. This explains the result that risk-averse agents with $\gamma < 1$ prefer large variances, which may initially seem counterintuitive.

[9] Since the consumption growth process is unchanged, the consumption-wealth ratio remains constant. The increase in initial consumption therefore corresponds to an increase in initial *wealth* by proportion ϕ.

of uncertainty is the value of ϕ which equates the two. This definition follows the lead of Lucas (1987) and Obstfeld (1994) and Section V of Alvarez and Jermann (2004).

The following sections consider two possible counterfactuals: (i) a scenario in which all uncertainty is eliminated, and (ii) a scenario in which the variance of consumption growth is reduced by α^2 but higher cumulants are unchanged. In each case, mean consumption growth $\mathbb{E}C_{t+1}/C_t$ is held constant.

IV.A The elimination of all uncertainty

Since $\mathbb{E}\left(C_1/C_0\right) = e^{\boldsymbol{c}(1)}$, keeping mean consumption growth constant is equivalent to holding $\boldsymbol{c}(1) = \log \mathbb{E}(C_1/C_0)$ constant. If all uncertainty is also to be eliminated, log consumption follows the trivial Lévy process $\overline{G}_t$ whose CGF is $\boldsymbol{c}_{\overline{G}}(\theta) = \boldsymbol{c}(1) \cdot \theta$ for all θ.

From (39) and (40), ϕ solves the equation

$$\frac{[(1+\phi)C_0]^{1-\gamma}}{1-\gamma} \cdot \left(1 + \frac{W_0}{C_0}\right) = \frac{C_0^{1-\gamma}}{1-\gamma} \cdot \frac{e^{\rho - \boldsymbol{c}(1)\cdot(1-\gamma)}}{e^{\rho - \boldsymbol{c}(1)\cdot(1-\gamma)} - 1} \,. \tag{41}$$

Simplifying, we have

$$\phi = \left(1 + \frac{W_0}{C_0}\right)^{\frac{1}{\gamma-1}} \left\{1 - e^{-\rho} \left[\mathbb{E}\left(\frac{C_1}{C_0}\right)\right]^{1-\gamma}\right\}^{\frac{1}{\gamma-1}} - 1 \,. \tag{42}$$

What assumptions are required to derive (42)? The left-hand side of (41) relies on the correspondence between expected utility and the consumption-wealth ratio that was noted at the beginning of section IV. This correspondence follows directly from Lucas's (1978) Euler equation with power utility: the assumption that real-world consumption growth is i.i.d. is not required. The cost of *all* uncertainty given in (42) depends only on the power utility assumption. The counterfactual case of deterministic growth is trivially i.i.d., so it is convenient to work with a CGF, though not necessary. (Below, I calculate the benefit associated with a reduction in the variance of consumption growth, while higher moments remain constant. In this case, the i.i.d. assumption is required and CGFs are central to my calculations.)

In the Epstein-Zin case it is also necessary to rely on the i.i.d. assumption. It turns out that (42) is misleading in that the γ terms that appear in it are capturing not risk aversion but the elasticity of intertemporal substitution, as the following proposition shows.

Proposition 6. *In the Epstein-Zin case with elasticity of intertemporal substitution* ψ*, the cost of uncertainty,* ϕ*, satisfies*

$$\phi = \left(1 + \frac{W_0}{C_0}\right)^{\frac{1}{1/\psi - 1}} \left\{1 - e^{-\rho} \left[\mathbb{E}\left(\frac{C_1}{C_0}\right)\right]^{1-\frac{1}{\psi}}\right\}^{\frac{1}{1/\psi - 1}} - 1\,. \tag{43}$$

With power utility and $\gamma \neq 1$*, the above equation holds, even in the absence of the i.i.d. assumption, with* $1/\psi$ *replaced by* γ*.*

With log utility we do require the i.i.d. assumption, and have

$$\begin{aligned} \phi &= \exp\left[(\boldsymbol{c}(1) - \mu) \,/\, (e^{\rho} - 1)\right] - 1 \\ &= \exp\left[(\boldsymbol{c}(1) - \mu)\,\frac{W_0}{C_0}\right] - 1\,. \end{aligned}$$

Proof. See appendix B for the Epstein-Zin calculations. □

Proposition 6 shows that if the mean consumption growth rate in levels, consumption-wealth ratio and preference parameters ρ and ψ can be estimated accurately, then the gains notionally available from eliminating all uncertainty can be estimated without needing to make assumptions about the particular stochastic process followed by consumption. In particular, in the Epstein-Zin case, ϕ is not—directly—dependent on γ, nor on estimates of the variance (and higher cumulants) of consumption growth. The consumption-wealth ratio encodes all relevant information about the amount of risk (that is, the cumulants κ_n, $n \geq 2$) and the representative agent's attitude to risk (γ).

In the power utility case in particular, this result is rather general. It applies to arbitrary consumption processes and so nests results obtained by Lucas (1987, 2003), Obstfeld (1994) and Barro (2006b).[10] The important feature is that I treat the consumption-wealth ratio as an observable. Lucas, Obstfeld and Barro postulate some particular

[10]Lucas (1987, 2003) assumes that current consumption C_0 is not known in the risky case. I follow Alvarez and Jermann (2004) in assuming that C_0 is known. The distinction is quantitatively insignificant in practice.

consumption process and, implicitly or explicitly, calculate the consumption-wealth ratio implied by that consumption process. For these authors, a change in γ is accompanied by a change in C/W; I, on the other hand, hold C/W constant and view it as containing information about the underlying consumption process.

IV.A.1 The cost of uncertainty with power utility

As before, suppose that $c/w = 0.06$ and $\boldsymbol{c}(1) = 0.02$, and that $\rho = 0.03$ and $\gamma = 4$. Substituting these values into (42) gives $\phi \approx 14\%$. This cost estimate is roughly two orders of magnitude higher than that obtained by Lucas (1987, 2003), even allowing for the higher risk aversion assumed in this paper. Although Lucas's calculations do not make use of the observable consumption-wealth ratio, it is possible to calculate the consumption-wealth ratio implied by his assumptions on the consumption process and my assumptions on ρ and γ; the result is an implied consumption-wealth ratio $c/w = 0.0896$. Substituting this value back into (42), we recover the far lower cost estimate, $\phi \approx 0.14\%$. Once one considers the consumption-wealth ratio as an observable, the cost of uncertainty appears to be considerably higher.

	ρ	γ	$\boldsymbol{c}(1)$	c/w	ϕ
Baseline case	**0.03**	**4**	**0.02**	**0.06**	**14%**
High ρ	0.04				18%
Low ρ	0.02				10%
High γ		5			16%
Low γ		3			7.7%
High growth			0.025		20%
Low growth			0.015		7.5%
High c/w				0.07	8.4%
Low c/w				0.05	21%

Table IV: The cost of consumption fluctuations with power utility.

Table IV shows how different assumptions on preference parameters and on mean consumption growth and the consumption-wealth ratio affect the estimate of the cost of uncertainty. Apart from the last two lines of the table, the consumption-wealth ratio c/w is held constant in the calculations.

The cost of uncertainty is *higher* when agents are more impatient (high ρ). When ρ is low, the (relatively) high consumption-wealth ratio signals that there is not too much risk in the economy. When ρ is high, the (relatively) low consumption-wealth ratio signals that there is considerable risk in the economy, or that risk aversion is high.

The case in which γ varies is somewhat more complicated. Suppose, first, that ρ is low relative to c/w, as in the above table. If we imagine holding the level of risk constant, then increasing γ from a low level will lead, first, to an increase in c/w because the representative agent is less inclined to substitute consumption intertemporally. Ultimately, however, increasing γ must lead to a decrease in c/w, once the precautionary saving motive starts to dominate. (These statements are most easily understood if one keeps Figure 1 in mind.) Turning the logic around, if γ increases but c/w remains constant, the level of risk in the economy must first be increasing and then declining. It follows that we may expect increases in γ to have ambiguous effects on the cost of uncertainty, holding c/w constant. In Table IV, the former effect dominates.

When, on the other hand, ρ is large relative to c/w, the CGF must have significant curvature—look at Figure 1. It follows that there is considerable risk in the economy; in this case, for γ to increase while c/w remains constant, it can only be that the level of risk is declining. Thus we expect to see that for low values of ρ, the cost of uncertainty is first increasing and then decreasing in γ, while for larger values of ρ, the cost is declining in γ.

These observations are borne out by Figure 6. When $\rho = 0.03$, the cost of uncertainty is first increasing and then decreasing in γ. When $\rho = 0.06$ or 0.09, the cost of uncertainty is decreasing in γ.

Finally, when ρ equals 0.03, γ must be at least 2.5 to be consistent with the assumed mean consumption growth and consumption-wealth ratio. In Figure 6, the black line hits zero at $\gamma = 2.5$ because the only possibility consistent with $\rho = 0.03, \gamma = 2.5, \boldsymbol{c}(1) =$

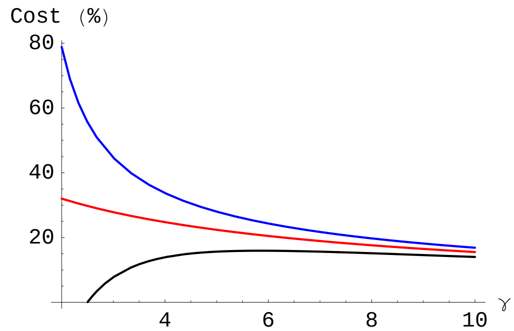

Figure 6: The cost of consumption uncertainty plotted against risk aversion, γ, when $\rho = 0.03$ (in black), $\rho = 0.06$ (in red) and $\rho = 0.09$ (in blue). The cost of uncertainty ultimately declines as γ increases: for very high values of γ, c/w can only equal 0.06 if there is relatively little risk in consumption growth.

$0.02, c/w = 0.06$ is that consumption is deterministic.

IV.A.2 The cost of uncertainty with Epstein-Zin preferences

Figure 7a illustrates the effects of changes in ρ and ψ on the cost of uncertainty. When ρ is high, the cost is high, for the same reasons as above. It is not possible to set ψ and ρ arbitrarily while retaining consistency with observed values of the consumption-wealth ratio. In Figure 7a, we see that we cannot have ψ between 0.4 and 1 if $\rho = 0.03$. However, if $\rho = c/w$, then ψ can take any value. Figure 7b therefore sets $\rho = c/w$ and shows that the cost of uncertainty increases in ψ. When ψ is around one, the implied cost of uncertainty is high, at about 40% of current wealth.

IV.B A reduction in the variance of consumption growth

The preceding section showed that there are significant costs due to uncertainty. This section investigates the utility benefit of a reduction in variance of α^2, holding all higher cumulants fixed. (It is possible to adjust variance alone, leaving higher cumulants unchanged, because the Brownian component of log consumption growth only affects the second cumulant.)

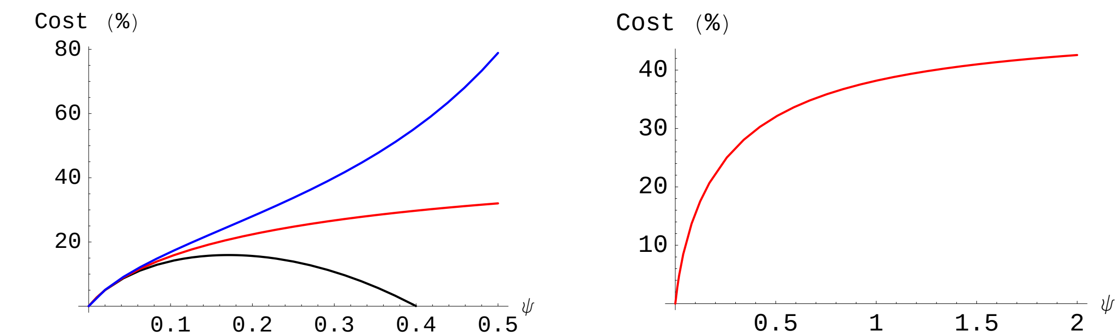

(a) Against ψ, with $\rho = 0.03$ (in black), $\rho = 0.06$ (in red) and $\rho = 0.09$ (in blue).

(b) Against ψ, with $\rho = c/w = 0.06$.

Figure 7: The cost of uncertainty with Epstein-Zin preferences.

Under the new reduced-volatility process, the CGF is

$$\widetilde{\boldsymbol{c}}(\theta) = \boldsymbol{c}(\theta) + \alpha^2\theta/2 - \alpha^2\theta^2/2 \,. \tag{44}$$

The term of order θ^2 decreases the variance of log consumption growth by α^2. The term of order θ adjusts the drift of log consumption growth to hold mean consumption growth constant in levels, that is, to ensure that $\widetilde{\boldsymbol{c}}(1) = \boldsymbol{c}(1)$.

The cost of uncertainty, ϕ_α, solves

$$\frac{[(1+\phi_\alpha)C_0]^{1-\gamma}}{1-\gamma} \cdot \left(1 + \frac{W_0}{C_0}\right) = \frac{C_0^{1-\gamma}}{1-\gamma} \cdot \frac{e^{\rho - \widetilde{\boldsymbol{c}}(1-\gamma)}}{e^{\rho - \widetilde{\boldsymbol{c}}(1-\gamma)} - 1} \,.$$

Substituting in from (44), and replacing $\rho - \boldsymbol{c}(1-\gamma)$ with the observable $c/w = \log(1 + C/W)$, we obtain after some simplification

$$\phi_\alpha = \left\{1 + \frac{W_0}{C_0}\left[1 - e^{-\frac{1}{2}\alpha^2\gamma(\gamma-1)}\right]\right\}^{1/(\gamma-1)} - 1 \,. \tag{45}$$

Carrying out similar calculations in the Epstein-Zin case, we find

Proposition 7. *In the Epstein-Zin case with elasticity of intertemporal substitution ψ, a reduction in consumption variance of α^2 is equivalent in utility terms to a proportional increase in current consumption of ϕ_α, where*

$$\phi_\alpha = \left\{1 + \frac{W_0}{C_0}\left[1 - e^{-\frac{1}{2}\alpha^2\gamma\left(\frac{1}{\psi}-1\right)}\right]\right\}^{\frac{1}{1/\psi - 1}} - 1 \,. \tag{46}$$

In the power utility case, the above equation holds with $1/\psi$ replaced by γ.

With log utility, we have

$$\begin{aligned}\phi_\alpha &= \exp\left[\frac{1}{2}\alpha^2/\left(e^\rho - 1\right)\right] - 1 \\ &= \exp\left[\frac{1}{2}\alpha^2 \frac{W_0}{C_0}\right] - 1\,.\end{aligned}$$

In all cases, we have the first-order approximation for small α^2

$$\phi_\alpha \approx \frac{W_0}{C_0}\frac{\gamma\alpha^2}{2}\,. \tag{47}$$

Proof. See Appendix B for the Epstein-Zin calculations. □

Obstfeld (1994) observes that (47) holds in the power utility case with i.i.d. lognormal consumption growth, but does not show that it holds in the Epstein-Zin case or for general i.i.d. consumption processes.

With $\gamma = 4$, and setting $c/w = 0.06$ as usual, it follows from (47) that a reduction in variance of 0.0003—as would be associated with a decline in the standard deviation of log consumption growth from 2% to 1%—is equivalent in welfare terms to an increase in current consumption (or equivalently wealth) of 1.0%. These calculations suggest, by comparison with the calculations of the previous subsection, that most of the cost of uncertainty can be attributed to higher-order cumulants.

V Conclusion

As pointed out by Rietz (1988), Barro (2006a) and Weitzman (2007), the tails of the distribution of consumption growth exert an enormous influence on asset prices. In this paper, I have taken an agnostic approach to the existence and importance of disasters by introducing a framework that handles general i.i.d. consumption growth processes. I showed that the predictions of disaster models are sensitively dependent on the assumptions made about the parameters governing the size and frequency of disasters.

This is problematic, because these parameters cannot be accurately estimated given the available data. I sidestepped this problem by deriving results that are valid no matter what is going on in the tails of consumption growth. In particular, these results do not depend on any assumptions about the size—or indeed existence—of disasters.

First, I derived bounds on the time preference rate and elasticity of intertemporal substitution, which are of central importance in an wide range of economic models. The bounds depend on observed values of the riskless rate, risk premium, and consumption-wealth ratio. Second, I showed that good-deal bounds can be combined with the consumption-wealth ratio and riskless rate to provide bounds on risk aversion for given time preference rate and elasticity of intertemporal substitution. Third, I showed, under assumptions more general than those made by Lucas (1987), Obstfeld (1994) or Barro (2006b), that it is possible to use the observed consumption-wealth ratio to estimate the welfare cost of uncertainty without specifying a consumption process. I estimate that the cost of uncertainty is on the order of 14%, and that almost all of this cost can be attributed to higher cumulants.

VI References

Aït-Sahalia, Y., Cacho-Diaz, J. and T. R. Hurd, 2006, Portfolio Choice with Jumps: A Closed Form Solution, preprint.

Alvarez, F. and U. J. Jermann, 2004, Using Asset Prices to Measure the Costs of Business Cycles, *Journal of Political Economy*, 112:6:1223–1256.

Backus, D. K., Chernov, M. and I. W. R. Martin, 2009, Extreme Events Implied by Equity Index Options, preprint.

Backus, D. K., Foresi, S. and C. I. Telmer, 2001, Affine Term Structure Models and the Forward Premium Anomaly, *Journal of Finance*, 56:1:279–304.

Barro, R. J., 2006a, Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics*, 121:3:823–866.

Barro, R. J., 2006b, On the Welfare Costs of Consumption Uncertainty, preprint.

Billingsley, P., 1995, *Probability and Measure*, 3rd edition (John Wiley & Sons, New York, NY).

Campbell, J. Y., 1986, Bond and Stock Returns in a Simple Exchange Model, *Quarterly Journal of Economics*, 101:4:785–804.

Campbell, J. Y., 2003, Consumption-Based Asset Pricing, Chapter 13 in George Constantinides, Milton Harris and Rene Stulz eds., *Handbook of the Economics of Finance vol. IB* (North-Holland, Amsterdam), 803–887.

Carr, P. and D. Madan, 1998, Option Valuation Using the Fast Fourier Transform, *Journal of Computational Finance*, 2:61–73.

Cochrane, J. H., 2005, *Asset Pricing*, revised edition (Princeton University Press, Princeton, NJ).

Cochrane, J. H. and J. Saá-Requejo, 2000, Beyond Arbitrage: Good Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy*, 108:1:79–119.

Cogley, T., 1990, International Evidence on the Size of the Random Walk in Output, *Journal of Political Economy*, 98:3:501–518.

Cont, R. and P. Tankov, 2004, *Financial Modelling with Jump Processes* (Chapman & Hall/CRC, Boca Raton, FL).

Cvitanić, J., Polimenis, V. and F. Zapatero, 2005, Optimal Portfolio Allocation with Higher Moments, preprint.

Duffie, D. and L. J. Epstein, 1992, Asset Pricing with Stochastic Differential Utility, *Review of Financial Studies*, 5:3:411–436.

Hansen, L. P. and R. Jagannathan, 1991, Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy*, 99:2:225–262.

Julliard, C. and A. Ghosh, 2008, Can Rare Events Explain the Equity Premium Puzzle?, preprint.

Jurek, J. W., 2008, Crash-Neutral Currency Carry Trades, preprint.

Kallsen, J., 2000, Optimal Portfolios for Exponential Lévy Processes, *Mathematical Methods of Operations Research*, 51:357–374.

Kocherlakota, N. R., 1990, Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result, *Journal of Finance*, 45:1:175–190.

Lucas, R. E., 1987, *Models of Business Cycles* (Basil Blackwell, Oxford, UK).

Lucas, R. E., 2003, Macroeconomic Priorities, *American Economic Review*, 93:1–14.

Lustig, H., Van Nieuwerburgh, S. and A. Verdelhan, 2008, The Wealth-Consumption Ratio, working paper.

Martin, I. W. R., 2008, The Lucas Orchard, working paper, Harvard University.

Obstfeld, M., 1994, Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability, *European Economic Review*, 38:1471–1486.

Rietz, T. A., 1988, The Equity Premium: A Solution, *Journal of Monetary Economics*, 22:117–131.

Sato, K., 1999, *Lévy Processes and Infinitely Divisible Distributions*, 1st edition, (Cambridge University Press, Cambridge, UK).

van der Vaart, A. W., 1998, *Asymptotic Statistics*, (Cambridge University Press, Cambridge, UK).

Weitzman, M. L., 2007, Subjective Expectations and Asset-Return Puzzles, *American Economic Review*, 97:4:1102–1130.

A Cumulants and cumulant-generating functions

This section lays out some important properties of cumulant-generating functions. It turns out that $\boldsymbol{c}(\theta)$ can be thought of as a power series in θ that encodes the cumulants (equivalently, moments) of consumption growth. To preview the main result, we have

$$\boldsymbol{c}(\theta) = \mu \cdot \frac{\theta}{1!} + \frac{\sigma^2\theta^2}{2!} + \text{skewness} \cdot \frac{\sigma^3\theta^3}{3!} + \text{kurtosis} \cdot \frac{\sigma^4\theta^4}{4!} + \dots$$

μ and σ denote the unconditional mean and standard deviation of log consumption growth.

Definition 2. *The* cumulants *of G are the coefficients κ_n in the power series expansion* *of the CGF* $\boldsymbol{c}(\theta)$*:*

$$\boldsymbol{c}(\theta) = \sum_{n=1}^{\infty} \frac{\kappa_n(G)\theta^n}{n!}\,. \tag{48}$$

Proposition 8. *We have the following properties.*

1. $\mathbb{E}G = \kappa_1$; $\operatorname{var} G = \kappa_2 \equiv \sigma^2$; *skewness* $(G) = \kappa_3/\sigma^3$; *excess kurtosis* $(G) = \kappa_4/\sigma^4$.

2. *For any two independent random variables* G *and* H*,* $\kappa_n(G+H) = \kappa_n(G)+\kappa_n(H)$ *and* $\boldsymbol{c}_{G+H}(\theta) = \boldsymbol{c}_G(\theta) + \boldsymbol{c}_H(\theta)$*.*

3. $\kappa_1(G) = \boldsymbol{c}'_G(0)$; $\kappa_2(G) = \boldsymbol{c}''_G(0)$; $\kappa_n(G) = \boldsymbol{c}^{(n)}_G(0)$*.*

4. κ_n *is a polynomial in the first* n *moments of* G*(and the* n*th moment of* G *is a polynomial in the first* n *cumulants of* G*).*

Proof. See Billingsley (1995, section 9). □

B Calculations with Epstein-Zin preferences

The Epstein-Zin first-order condition leads to the pricing formula

$$P = \mathbb{E}\sum_{1}^{\infty} e^{-\rho\vartheta t}\left(\frac{C_t}{C_0}\right)^{-\vartheta/\psi}\left(1 + R_{m,0\to t}\right)^{\vartheta-1}\left(C_t\right)^{\lambda},$$

where $\vartheta = (1-\gamma)/(1-1/\psi)$ and $R_{m,0\to t}$ is the cumulative return on the wealth portfolio from period 0 to period t. I assume that $\psi \neq 1$ for convenience.

Now,

$$\begin{aligned} 1 + R_{m,s-1\to s} &= \frac{C_s + W_s}{W_{s-1}} \\ &= \frac{C_s}{C_{s-1}}\left(\frac{C_{s-1}}{W_{s-1}} + \frac{W_s}{C_s}\frac{C_{s-1}}{W_{s-1}}\right) \\ &= \frac{C_s}{C_{s-1}}e^{\nu}, \end{aligned}$$

where the last equality follows by making the assumption—provisional for the time being, but subsequently shown to be correct—that the consumption-wealth ratio is constant. I have defined $1 + C/W \equiv e^{\nu}$. It follows, then, that

$$1 + R_{m,0\to t} = \frac{C_t}{C_0}e^{\nu t},$$

and hence that

$$\begin{aligned} P &= (C_0)^\lambda \cdot \mathbb{E}\sum_1^\infty e^{-\rho\vartheta t}\left(\frac{C_t}{C_0}\right)^{\lambda-\vartheta/\psi}\left(\frac{C_t}{C_0}\right)^{\vartheta-1} e^{\nu(\vartheta-1)t} \\ &= (C_0)^\lambda \cdot \sum_1^\infty e^{-[\rho\vartheta+\nu(1-\vartheta)-\boldsymbol{c}(\lambda-\gamma)]t} \\ &= \frac{(C_0)^\lambda}{e^{\rho\vartheta+\nu(1-\vartheta)-\boldsymbol{c}(\lambda-\gamma)}-1}\,, \end{aligned}$$

and so, finally, that

$$\frac{D}{P} = e^{\rho\vartheta+\nu(1-\vartheta)-\boldsymbol{c}(\lambda-\gamma)} - 1\,.$$

Defining d/p as usual,

$$d/p = \rho\vartheta + \nu(1-\vartheta) - \boldsymbol{c}(\lambda-\gamma)\,. \tag{49}$$

Setting $\lambda = 1$, we get an expression for $c/w \equiv \nu$ which can be solved for ν:

$$\nu = c/w = \rho\vartheta + \nu(1-\vartheta) - \boldsymbol{c}(1-\gamma)\,,$$

from which it follows that

$$\nu = \rho - \boldsymbol{c}(1-\gamma)\cdot\frac{1-\psi}{\psi(\gamma-1)}\,.$$

Note that this exercise confirms the provisional assumption made above that ν is constant.

Substituting back into (49), we have

$$dp = \rho - \frac{1-\psi\gamma}{\psi(\gamma-1)}\boldsymbol{c}(1-\gamma) - \boldsymbol{c}(\lambda-\gamma)\,.$$

We also have, as before, that

$$1 + R_{t+1} = \frac{D_{t+1}}{D_t}\left(e^{\rho\vartheta+\nu(1-\vartheta)-\boldsymbol{c}(\lambda-\gamma)}\right)\,,$$

so

$$er = \rho\vartheta + \nu(1-\vartheta) + \boldsymbol{c}(\lambda) - \boldsymbol{c}(\lambda-\gamma)\,.$$

To summarize, we have

$$
\begin{aligned}
r_f &= \rho - \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\left(\frac{1}{\vartheta} - 1\right) \\
c/w &= \rho - \boldsymbol{c}(1-\gamma)/\vartheta \\
rp &= \boldsymbol{c}(1) + \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma)\,.
\end{aligned}
$$

The objective function at time 0 satisfies

$$
(U_0)^{(1-\gamma)/\vartheta} = \left(1 - e^{-\rho}\right)(C_0)^{(1-\gamma)/\vartheta} + e^{-\rho}\left(\mathbb{E}(U_1)^{1-\gamma}\right)^{1/\vartheta}
$$

or

$$
a_0^{(1-\gamma)/\vartheta} = 1 - e^{-\rho} + e^{-\rho}\mathbb{E}\left[\left(\frac{C_1}{C_0}\right)^{1-\gamma} a_1^{1-\gamma}\right]^{1/\vartheta}, \tag{50}
$$

where I have defined $a_i \equiv U_i/C_i$.

I now conjecture that $a_i = a$, some constant, solves (50). If so,

$$
a^{(1-\gamma)/\vartheta} = 1 - e^{-\rho} + e^{-\rho} a^{(1-\gamma)/\vartheta} e^{\boldsymbol{c}(1-\gamma)/\vartheta},
$$

from which it follows that

$$
a = \left(\frac{1 - e^{-\rho}}{1 - e^{-\rho + \boldsymbol{c}(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)},
$$

which confirms the conjecture that a was constant. Hence,

$$
U_0 = C_0 \cdot \left(\frac{e^{\rho} - 1}{e^{\rho} - e^{\boldsymbol{c}(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)}.
$$

The cost of all uncertainty, ϕ, solves the equation

$$
(1+\phi)\,C_0 \cdot \left(\frac{e^{\rho} - 1}{e^{\rho} - e^{\boldsymbol{c}(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)} = C_0\left(\frac{e^{\rho} - 1}{e^{\rho} - e^{\boldsymbol{c}(1)\cdot(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)},
$$

from which (43) follows.

Similarly, ϕ_α solves

$$
(1+\phi_\alpha)C_0 \cdot \left(\frac{e^{\rho} - 1}{e^{\rho} - e^{\boldsymbol{c}(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)} = C_0 \cdot \left(\frac{e^{\rho} - 1}{e^{\rho} - e^{\widetilde{\boldsymbol{c}}(1-\gamma)/\vartheta}}\right)^{\vartheta/(1-\gamma)},
$$

and after substituting in for $\widetilde{\boldsymbol{c}}(\theta)$ from equation (44), we obtain the expression (46).