American Economic Review: Papers & Proceedings 2008, 98:2, 74–78
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.2.74

Disasters and the Welfare Cost of Uncertainty

By Ian W. R. Martin*

The combination of power utility and i.i.d. lognormal consumption growth makes for a benchmark model in which asset prices and expected returns can be found in closed form. Introducing the consumption-based model, John H. Cochrane (2005, 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." A message of this paper is that the lognormality assumption can be relaxed without sacrificing tractability.

Working under two assumptions—that there is a representative agent with power utility and that consumption growth is i.i.d.—I introduce, in Section I, a mathematical object (the cumulant-generating function, or CGF) in terms of which four fundamental quantities that are at the heart of consumption-based asset pricing can be simply expressed. Those quantities are the equity premium, riskless rate, consumption-wealth ratio, and mean consumption growth.

The expressions derived relate the fundamentals directly to the cumulants (equivalently, moments) of consumption growth. The lognormal assumption is equivalent to the assumption that all cumulants above the second are zero.

If one is in the business of making up stochastic processes, many suggest themselves most naturally in continuous time. Although there is an obvious discrete-time analogue of Brownian motion—a random walk with Normally distributed increments—it is less natural to map Poisson processes, say, into discrete time, and therefore harder to deal with the possibility of jumps in consumption. In Section II, I show that these results carry over to the continuous-time setting. The i.i.d. growth assumption is replaced by its continuous-time analogue: log consumption is a Lévy process.

* Harvard University, Littauer Center, Cambridge, MA 02138 (e-mail: iwmartin@fas.harvard.edu). I thank John Campbell for his advice and encouragement. I am particularly grateful to Robert Barro, who generously showed me an early draft of the working paper that inspired me to work on this topic.

The first few cumulants of consumption growth can, in principle, be estimated from consumption data, though this approach is not taken in the present paper because, given the sizes of the relevant samples in practice, estimates of higher cumulants have large standard errors. This is especially troubling because the higher cumulants, which are hardest to estimate, are extremely influential for asset prices.

I illustrate this point, and the CGF framework more generally, by investigating a continuous-time model featuring rare disasters as in Thomas A. Rietz (1988) or Robert J. Barro (2006a). Worryingly, the model's predictions are sensitively dependent on the calibration assumed.

As a stark illustration within the i.i.d. framework, suppose that the representative agent has relative risk aversion equal to four. Now imagine adding to the model a certain type of disaster that destroys 90 percent of wealth and strikes, on average, once every 100,000 years. (Barro (2006a) documents that Germany and Greece each suffered a 64 percent fall in per capita real GDP in the course of the Second World War, so such a disaster is not beyond the bounds of possibility.) The introduction of the very rare, very severe disaster drives the riskless rate down by 10 percent—1,000 basis points—and increases the equity premium by 9 percent. Very rare, very severe events exert an extraordinary influence, and we do not expect to estimate their frequency and intensity directly from the data.

We can, however, detect the influence of such events *indirectly* by observing asset prices. I argue, therefore, that the standard approach—which consists of calibrating a particular model and trying to fit the fundamental quantities—is not the way to go. By turning things round—viewing the fundamental quantities as observable and seeing what they imply—it becomes possible to make statements that are robust to the details of the consumption growth process.

In Section III, I take up the question, surveyed by Robert E. Lucas (2003), of the cost of consumption risk. This cost turns out to depend on the time discount rate, ρ, and the risk aversion parameter, γ, and on two observables: mean consumption growth and the consumption-wealth ratio—which, unlike Lucas (2003), I view as observable. The cost does not depend on risk aversion other than 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 captured by the cumulants.

Using plausible parameters, the cost of consumption uncertainty is on the order of 14 percent. This cost is largely attributable to higher cumulants: I estimate that the representative agent would sacrifice only about 1 percent of initial consumption to reduce the standard deviation of consumption growth from 2 percent to 1 percent.

Proofs are provided in an Web Appendix (http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.2.74).

I. Asset Pricing and the CGF

My analysis rests on two assumptions.

ASSUMPTION A1: *There is a representative agent with constant relative risk aversion γ and time discount rate ρ.*

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

Assumption A2 is strong, and it is essential for the calculations of this paper. Timothy 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.

My starting point is the Euler equation

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

Consider an asset that pays dividend stream $D_t \equiv (C_t)^{\lambda}$ for some constant λ. I write P_λ for the price of this asset at time 0, and D_λ for the dividend at time 0. Using A1, A2, and (1),

$$P_\lambda = D_\lambda \sum_{t=1}^{\infty} e^{-\rho t}\left(\mathbb{E}\left(e^{(\lambda-\gamma)G}\right)\right)^t. \tag{2}$$

To make further progress, I now introduce:

DEFINITION 1: *The cumulant-generating function, or CGF, $\boldsymbol{c}(\theta)$ is defined by*

$$\boldsymbol{c}(\theta) \equiv \log \mathbb{E} \exp(\theta G). \tag{3}$$

(Observe in particular that $\boldsymbol{c}(1)$ equals log mean gross consumption growth.)

The CGF can be thought of as capturing information about all moments of G. More precisely, we can expand $\boldsymbol{c}(\theta)$ as a power series, $\boldsymbol{c}(\theta) = \sum_{n=1}^{\infty} \kappa_n \theta^n/n!$, 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 excess kurtosis of log consumption growth. Knowledge of all cumulants implies knowledge of all moments, and vice versa.

Using this definition and setting $d_\lambda/p_\lambda \equiv \log(1 + D_\lambda/P_\lambda) \approx D_\lambda/P_\lambda$, it follows from (2) that $d_\lambda/p_\lambda = \rho - \boldsymbol{c}(\lambda - \gamma)$. If $\lambda = 0$, the asset in question is the riskless bond, whose dividend yield is the riskless rate. If $\lambda = 1$, the asset pays consumption as its dividend, and can therefore be interpreted as aggregate wealth; its dividend yield is then the consumption-wealth ratio.

The gross return on the consumption claim is

$$1 + R_{t+1} = \frac{P_{t+1}}{P_t}\left(1 + \frac{D_{t+1}}{P_{t+1}}\right) = \frac{D_{t+1}}{D_t}\left(e^{\rho - \boldsymbol{c}(1-\gamma)}\right)$$

and thus the expected gross return is $1 + \mathbb{E}R_{t+1} = e^{\rho - \boldsymbol{c}(1-\gamma) + \boldsymbol{c}(1)}$. Once again, it is more convenient to work with log expected gross return, $er \equiv \log(1 + \mathbb{E}R_{t+1}) = \rho + \boldsymbol{c}(1) - \boldsymbol{c}(1-\gamma)$, and to define the risk premium $rp \equiv er - r_f$.

THEOREM 2: *Summarizing the results above, we have*

$$r_f = \rho - \boldsymbol{c}(-\gamma), \tag{4}$$

$$c/w = \rho - \boldsymbol{c}(1-\gamma), \tag{5}$$

$$rp = \boldsymbol{c}(1) + \boldsymbol{c}(-\gamma) - \boldsymbol{c}(1-\gamma). \tag{6}$$

When consumption growth is deterministic, the CGF is linear, and equation (6) shows as expected that 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.

Given a sufficiently long data sample, expressions (4)–(6) could be applied by estimating the cumulants of log consumption directly, without imposing any further structure on the model. In practice, we cannot estimate infinitely many cumulants from a finite dataset; one solution to this is to truncate after the first N cumulants, N being determined by the amount of data available. (The assumption that consumption growth is lognormal is equivalent to truncating at $N = 2$, since when log consumption growth is Normal, all cumulants above the variance are equal to zero.)

II. The Continuous-Time Case

For the purposes of constructing concrete examples, it is often convenient to work in continuous time. The utility function is modified in the obvious manner by replacing a summation with an integral; the assumption that dividend growth is i.i.d. is replaced by an assumption that the log consumption path, G_t, of the representative agent follows a Lévy process.

The following result confirms that the simplicity of the framework carries over to the continuous-time case.

THEOREM 3: *Theorem 2 holds in continuous time, except that c/w, r_f, and rp are replaced by the instantaneous consumption-wealth ratio C/W, the instantaneous riskless rate R_f, and the instantaneous risk premium RP.*

A. *A Concrete Example: Disasters*

Suppose that log consumption follows the jump-diffusion process $G_t = \widetilde{\mu} t + \sigma_B B_t + \sum_{i=1}^{N(t)} Y_i$, where B_t is a Brownian motion, $N(t)$ is a Poisson counting process with parameter ω, and Y_i are i.i.d. $N(-b, \psi^2)$ random variables.

A simple calculation reveals that

$$(7) \quad \boldsymbol{c}(\theta) = \widetilde{\mu}\theta + \sigma^2_B\theta^2/2 + \omega(e^{-b\theta + \psi^2\theta^2/2} - 1).$$

With the explicit expression (7) for the CGF in hand, it is easy to investigate the sensitivity of a disaster model's predictions to the parameter values assumed. Table 1 shows how changes in the calibration of the distribution of disasters affect the relevant fundamentals. The baseline mean and variance of disaster sizes are set equal to the mean and variance of the disasters reported in Barro (2006a), and the baseline disaster arrival rate $\omega = 0.017$ is taken from the same paper. As is evident from the table, the predictions of the disaster model are sensitively dependent on the precise calibration. Small changes in the disaster parameters ω and b have large effects on the riskless rate and equity premium. For example, increasing ω (the rate of arrival of disasters, in percent) from 1.7 percent to 2.2 percent drives the riskless rate down by more than 3 per cent. Given that these parameters are hard to estimate—disasters happen rarely—this is a significant difficulty.

Table 1—The Impact of Different Assumptions about the Distribution of Disasters

	ω	b	R_f	C/W	RP
Baseline	**1.7**	**0.39**	**1.0**	**4.8**	**5.7**
High ω	2.2	0.39	−2.4	3.1	7.4
Low ω	1.2	0.39	4.5	6.4	4.1
High b	1.7	0.44	−1.9	3.6	7.5
Low b	1.7	0.34	3.5	5.8	4.4

III. The Cost of Consumption Fluctuations

The discussion above suggests that it is desirable to try to make statements that do not depend on a particular calibration of the disaster process. I illustrate this approach by estimating the cost of consumption fluctuations.

An easy calculation reveals that—assuming $\gamma \neq 1$ for simplicity—expected utility can be expressed in terms of the CGF:

$$(8) \quad U(C_0; \boldsymbol{c}) = \frac{C_0^{1-\gamma}}{1-\gamma}\left(1 + \frac{1}{e^{\rho - \boldsymbol{c}(1-\gamma)} - 1}\right).$$

Expression (8) permits the calculation of expected utility under alternative consumption processes via the corresponding CGFs. The cost of uncertainty of the status quo, relative to some counterfactual summarized by $\widetilde{\boldsymbol{c}}$, is the value of $\widetilde{\phi}$, which solves

$$(9) \quad U[(1 + \widetilde{\phi})C_0; \boldsymbol{c}] = U[C_0; \widetilde{\boldsymbol{c}}].$$

I consider two counterfactuals: (a) a scenario in which all uncertainty is eliminated, and (b) a scenario in which the variance of consumption growth is reduced by α^2 but higher cumulants are unchanged. I hold mean consumption growth constant: $\boldsymbol{c}(1) = \widetilde{\boldsymbol{c}}(1)$.

A. *The Elimination of All Uncertainty*

Since $\mathbb{E}C_1/C_0 = \mathbb{E}e^G = e^{\boldsymbol{c}(1)}$, keeping mean consumption growth constant is equivalent to holding $\boldsymbol{c}(1)$ constant. If, also, log consumption is deterministic, it must follow the trivial Lévy process with CGF $\bar{\boldsymbol{c}}(\theta) = \boldsymbol{c}(1)\cdot\theta$.

From equations (8) and (9), and replacing $\rho - \boldsymbol{c}(1-\gamma)$ with c/w, we find

$$\text{(10)} \qquad \bar{\phi} = \left(1 + \frac{W_0}{C_0}\right)^{1/(\gamma-1)} \times \left\{1 - e^{-\rho}\left[\mathbb{E}\left(\frac{C_1}{C_0}\right)\right]^{1-\gamma}\right\}^{1/(\gamma-1)} - 1.$$

Equation (10) shows that if the mean consumption growth rate in levels, consumption-wealth ratio, and preference parameters ρ and γ 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, $\bar{\phi}$ is not—directly—dependent on estimates of the variance (or higher cumulants) of consumption growth: the relevant information is encoded in the consumption-wealth ratio.

This result applies to arbitrary consumption processes and so nests results obtained by Lucas (2003) and Barro (2006b). Unlike these authors, I treat the consumption-wealth ratio as an observable that encodes information about the underlying consumption process.

To make this concrete, I will impose the baseline parameters $c/w = 0.06$, $\boldsymbol{c}(1) = 0.02$, $\rho = 0.03$, $\gamma = 4$. Substituting into (10) gives $\bar{\phi} \approx 14$ percent. This cost estimate is roughly two orders of magnitude higher than that obtained by Lucas (2003), even allowing for the higher risk aversion assumed in this paper.

Figure 1 shows how the cost of uncertainty varies with c/w. The maximum possible value of c/w, 0.09, is achieved if there is no uncertainty; then, of course, the cost of uncertainty is zero. As c/w decreases from this maximum possible level, the implied cost of uncertainty increases rapidly. Although Lucas's calculations do not take c/w as observable, his assumptions on the consumption process, combined with my assumptions on ρ and γ, imply $c/w = 0.0896$. Substituting into (10), we recover the far lower cost estimate $\bar{\phi} \approx 0.14$ percent.[1]

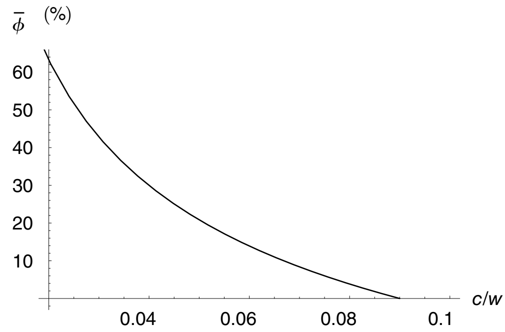

FIGURE 1. COST OF UNCERTAINTY AGAINST *c/w*

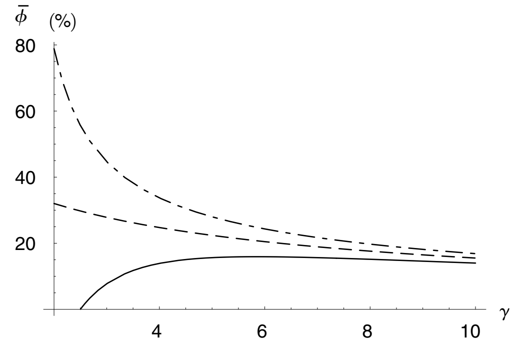

FIGURE 2. COST OF UNCERTAINTY AGAINST γ
$\rho = 0.03$ (*solid*), 0.06 (*dashed*), 0.09 (*dot-dashed*)

Figure 2 shows how the cost of uncertainty depends on γ. As γ becomes very large, the cost of uncertainty ultimately *declines*: if γ is extremely large, it must be the case, given that c/w is held fixed, that there is very little risk in the economy. Similarly, when $\rho = 0.03$,

[1] Lucas's calculations do not assume i.i.d. consumption growth. In fact, though, expression (10) does not require any assumptions on the consumption process: it follows directly from the Euler equation. When I consider the cost of variance uncertainty below, the i.i.d. assumption is required.

the line hits zero at $\gamma = 2.5$ because the only possibility consistent with $\rho = 0.03$, $\gamma = 2.5$, $\boldsymbol{c}(1) = 0.02$, $c/w = 0.06$ is that consumption is deterministic.

B. *A Reduction in Variance*

I now investigate an alternative counterfactual in which the variance of log consumption growth is reduced by α^2. (It is possible to consider such an adjustment in variance alone—leaving higher cumulants unchanged—because the Brownian component of log consumption growth affects only the second cumulant. Conversely, it is not clear how to adjust, say, kurtosis without changing other cumulants.) The new CGF is then

$$\tilde{\boldsymbol{c}}(\theta) = \boldsymbol{c}(\theta) + \alpha^2\theta/2 - \alpha^2\theta^2/2. \tag{11}$$

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 $\boldsymbol{c}_\alpha(1) = \boldsymbol{c}(1)$.

Substituting (11) into (9) and replacing $\rho - \boldsymbol{c}(1-\gamma)$ with the observable c/w, we find

$$\phi_\alpha = \left\{1 + \frac{W_0}{C_0}\left[1 - e^{-\frac{1}{2}\alpha^2\gamma(\gamma-1)}\right]\right\}^{1/(\gamma-1)} - 1. \tag{12}$$

With $\gamma = 4$, and setting $c/w = 0.06$ as usual, it follows from (12) that a reduction in variance of 0.0003 —a decline in the standard deviation of log consumption growth from 2 percent to 1 percent, for example—is equivalent in welfare terms to an increase in current consumption (or equivalently wealth) of 1.0 percent. Most of the cost of uncertainty can be attributed to higher-order cumulants.

IV. Conclusion

Cumulant-generating functions render the general power utility–i.i.d. model tractable. The mere fact that they simplify notation makes them useful modelling tools. In more complicated settings (Martin 2007), it may even be easier to work with a CGF than to consider a special case such as lognormality, simply because the CGF's progress can be easily tracked through the algebra.

The other theme of this paper is that it is desirable, when thinking about disasters, to try to make statements that are not sensitively dependent on the assumed pattern of higher cumulants. Section III showed that it is possible to use the observed consumption-wealth ratio to estimate the welfare cost of uncertainty without specifying a consumption process, and argued also that the cost is high.

REFERENCES

Barro, Robert J. 2006a. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–66.

Barro, Robert J. 2006b. *On the Welfare Costs of Consumption Uncertainty.* National Bureau of Economic Research Working Paper 12763.

Cochrane, John H. 2005. *Asset Pricing.* 2nd ed. Princeton, NJ: Princeton University Press.

Cogley, Timothy. 1990. "International Evidence on the Size of the Random Walk in Output." *Journal of Political Economy*, 98(3): 501–18.

Lucas, Robert E., Jr. 2003. "Macroeconomic Priorities." *American Economic Review*, 93(1): 1–14.

Martin, Ian W. R. 2007. "The Lucas Orchard." Unpublished.

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics*, 22(1): 117–31.