American Economic Review 2009, 99:1, 243–264
http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.1.243

Rare Disasters, Asset Prices, and Welfare Costs

By Robert J. Barro*

A representative-consumer model with Epstein-Zin-Weil preferences and i.i.d. shocks, including rare disasters, accords with observed equity premia and risk-free rates if the coefficient of relative risk aversion equals 3–4. If the intertemporal elasticity of substitution exceeds one, an increase in uncertainty lowers the price-dividend ratio for equity, and a rise in the expected growth rate raises this ratio. Calibrations indicate that society would willingly reduce GDP by around 20 percent each year to eliminate rare disasters. The welfare cost from usual economic fluctuations is much smaller, though still important, corresponding to lowering GDP by about 1.5 percent each year. (*JEL* E13, E21, E22, E32)

In a previous study, Barro (2006), I used the Thomas A. Rietz (1988) idea of rare economic disasters to explain the equity premium and related asset-pricing puzzles. My quantitative examination of large macroeconomic contractions in 35 countries during the twentieth century suggested a disaster probability of roughly 2 percent per year. The size distribution of GDP contractions during these events ranged between 15 percent (the arbitrary lower bound) and over 60 percent. A simple representative-agent economy, calibrated to accord with this disaster experience, can explain an equity premium of around 4–6 percent and a risk-free real interest rate of 1–2 percent. With power-utility preferences, these results require a coefficient of relative risk aversion of 3–4. The analysis applies in a Lucas-tree economy with i.i.d. production shocks or to an "AK model" with endogenous saving and stochastic depreciation.

The present analysis extends the framework to consider additional aspects of asset pricing and to assess the welfare cost of consumption uncertainty. As observed by Ravi Bansal and Amir Yaron (2004), power-utility preferences with a coefficient of relative risk aversion above one generate two implausible predictions. First, an increase in uncertainty *raises* the price-dividend ratio for equities and, second, a rise in the mean growth rate *lowers* the price-dividend ratio. More reasonable predictions require an intertemporal elasticity of substitution (IES) above one. In the power-utility framework, however, this property conflicts with a coefficient of relative risk aversion greater than one—a condition needed to match observed equity premia. Therefore, to fit a broad set of asset-pricing "facts," it is essential to use a preference specification, such as that of Larry G. Epstein and Stanley E. Zin (1989) and Philippe Weil (1990), that delinks the IES from the coefficient of relative risk aversion. Power utility, although attractive for its simplicity, cannot work.

The framework is still a representative-consumer model with i.i.d. shocks to production. In this setting, the key asset-pricing conditions under Epstein-Zin-Weil (henceforth, EZW) preferences resemble those with power utility. However, two key differences emerge. First, under EZW preferences, consumption enters into asset-pricing formulas with an exponent that involves the

* Economics Department, Harvard University, Cambridge, MA 02138 (e-mail: rbarro@harvard.edu). This research is supported by the National Science Foundation (grant SES 0617253). I appreciate comments from Fernando Alvarez, Marios Angeletos, John Campbell, Raj Chetty, John Cochrane, Xavier Gabaix, Francois Gourio, Kai Guo, Lars Hansen, David Laibson, Jong-Wha Lee, Greg Mankiw, Ian Martin, Casey Mulligan, Marty Weitzman, and Iván Werning.

coefficient of relative risk aversion, not the IES. Second, the formulas involve an effective rate of time preference, denoted ρ^*, that deviates from the usual rate of time preference, ρ, when the coefficient of relative risk aversion is unequal to the reciprocal of the IES. The value of ρ^* depends on ρ, the IES, the coefficient of relative risk aversion, and the other parameters of the model—including parameters that describe expected growth and uncertainty.

With i.i.d. shocks, the EZW framework ends up as simple as the power-utility setting, and it accords with a broader set of asset-pricing facts. First, when calibrated to the observed frequency and size distribution of macroeconomic disasters, the model can explain the equity premium and risk-free rate, still with a coefficient of relative risk aversion of 3–4. Second, with an IES above one, the model predicts that an increase in uncertainty lowers the dividend-price ratio, whereas a rise in the expected growth rate raises this ratio. Third, in an *AK* model that allows for endogenous saving, the IES above one implies that more uncertainty *lowers* the saving ratio (because substitution effects dominate when the IES exceeds one).

Robert E. Lucas, Jr. (1987, ch. 3; 2003, sect. II) argued that the welfare gain from eliminating aggregate consumption uncertainty is trivial. One problem with his calculation, apparent from Rajnish Mehra and Edward C. Prescott (1985), is that calibrations of Lucas's model do not get into the right ballpark for explaining the high equity premium. This asset-pricing failure suggests, as observed by Andrew Atkeson and Christopher Phelan (1994), that the model misses important aspects of consumption uncertainty. Hence, the model's estimates of welfare effects from aggregate consumption uncertainty are unlikely to be accurate.

A reasonable principle is that analyses of the impacts of consumption uncertainty should be carried out within models that at least roughly replicate the way that asset markets price this uncertainty. This Atkeson-Phelan principle was followed by Fernando Alvarez and Urbann J. Jermann (2004) and is also adopted in the present paper. In my case, the prospects of rare economic disasters, as in Rietz (1988), are critical for matching asset-pricing facts. Within this setting, changes in consumption uncertainty that reflect shifts in the probability of disaster have major implications for welfare. Individuals would willingly relinquish as much as 20 percent of GDP each year in exchange for eliminating all chances for macroeconomic disaster. The welfare cost from usual economic fluctuations is much smaller, though still important—corresponding to lowering GDP by around 1.5 percent each year.

Section I works out the Lucas-tree model with rare disasters. The key asset-pricing formulas under EZW preferences are derived here. Section II computes welfare costs within this model, first for marginal changes in uncertainty and then for large changes. Section III discusses the sensitivity of the welfare-cost calculations to the two key preference parameters: the coefficient of relative risk aversion and the intertemporal elasticity of substitution. Section IV allows for endogenous labor supply. A key conclusion is that any wage elasticity of labor supply is compatible with a given coefficient of relative risk aversion. Section V includes endogenous saving and investment and shows how adjustments of saving affect welfare costs. Section VI concludes by emphasizing the effects of policies and institutions on disaster probabilities and sizes.

I. A Lucas Fruit-Tree Model

The initial model is a version of Lucas's (1978) representative-agent, fruit-tree economy with exogenous, stochastic production. Output of fruit in period t equals real GDP, Y_t. Population is constant. The number of trees is fixed; that is, there is neither investment nor depreciation. (The model in Section IV allows for investment.) Government purchases are nil. Since the economy is closed and all output is consumed, consumption, C_t, equals Y_t.

The log of output evolves as a random walk with drift:

$$\log(Y_{t+1}) = \log(Y_t) + g + u_{t+1} + v_{t+1}. \tag{1}$$

The random term u_{t+1} is i.i.d. normal with mean 0 and variance σ^2. This term reflects "normal" economic fluctuations. The parameter $g \geq 0$ is a constant that reflects exogenous productivity growth.

The random term v_{t+1} in equation (1) picks up low-probability disasters, as in Rietz (1988) and Barro (2006). In these rare events, output and consumption jump down sharply. The probability of a disaster is the constant $p \geq 0$ per unit of time. The probability of more than one disaster in a period is assumed to be small enough to neglect; later, the arbitrary period length shrinks to zero. In a disaster, output contracts by the fraction b, where $0 < b < 1$. The distribution of v_{t+1} is given by

$$\text{probability } 1 - p\text{: } v_{t+1} = 0,$$

$$\text{probability } p\text{: } v_{t+1} = \log(1 - b).$$

The disaster size, b, follows some probability distribution (gauged subsequently by the empirical distribution of these sizes).

Unlike Lucas (1987, ch. 3), but in line with Maurice Obstfeld (1994), the shocks u_{t+1} and v_{t+1} in equation (1) represent permanent effects on the level of output, rather than transitory disturbances to the level. That is, the economy has no tendency to revert to a deterministic trend line.

John H. Cochrane (1988, table 1) used variance-ratio statistics for k-year differences to assess the extent of reversion to a deterministic trend in the log of US real per capita GNP for 1869–1986. He found evidence for reversion in that the ratio of the k-year variance (divided by k) to the one-year variance was between 0.30 and 0.36 for k between 20 and 30 years. Therefore, at large k, the empirical variance ratio was much less than the value 1.0 predicted by equation (1). However, Timothy Cogley (1990, table 2) showed that the Cochrane finding was particular to the United States. For 9 OECD countries, including the United States, from 1871 to 1985, the mean of the variance ratio at 20 years was 1.1, hence, close to the value 1.0 predicted by equation (1).

Cogley's results hold up for a broader sample comprising 19 OECD countries. The data on per capita GDP are for 1870–2005 from Angus Maddison (2003), updated from the World Bank *World Development Indicators* (and using US data from Nathan S. Balke and Robert J. Gordon (1989) before 1929). For $k = 20$, the mean of the variance ratios for the 19 countries is 1.22 and the median is 1.00, while for $k = 30$, the corresponding values are 1.30 and 0.96. These values accord with equation (1). The United States—with variance ratios of 0.42 when $k = 20$ and 0.38 when $k = 30$—has the lowest ratios at these values of k among the 19 countries.[1] The critical factor for the United States is that the turbulence of the Great Depression and World War II happened to be followed by the log of per capita GDP reverting roughly to the pre-1930 and pre-1914 trend lines. Most other countries do not look like this.

My inference from the long-term GDP data for the OECD countries is that the evidence conflicts strongly with reversion to a fixed, deterministic trend. The key, counterfactual prediction from this model is the comparatively low uncertainty about the distant future. In contrast, the variance-ratio results are consistent with the stochastic-trend specification in equation (1). Therefore, I use this model for the present analysis. Richer models of GDP and consumption that I am

[1] The next smallest values for $k = 20$ are 0.55 for New Zealand, 0.68 for Germany, and 0.77 for Switzerland. At $k = 30$, the next smallest values are 0.40 for New Zealand, 0.53 for Germany, and 0.54 for Canada. For smaller values of k, the mean and median of the variance ratios are, respectively, 1.16 and 1.18 at $k = 2$, 1.23 and 1.31 at $k = 5$, and 1.13 and 1.06 at $k = 10$. The US ratios at these values of k are, respectively, 1.30, 1.34, and 0.94.

currently studying (in joint work with Emi Nakamura, Jon Steinsson, and Jose Ursua) allow for trend breaks (analyzed starting from Anindya Banerjee, Robin Lumsdaine, and James H. Stock 1992) and for gradual reversion to past levels after major disasters, such as destructive wars and financial crises.

Previous research (Barro 2006, table 1 and figure 1) gauged the probability and size distribution of disaster events from time series on real per capita GDP for 35 countries for the full twentieth century.[2] That study defined a macroeconomic disaster as a decline in real per capita GDP by at least 15 percent over consecutive years (such as 1939–1944 for France and 1929–1933 for the United States). These kinds of events are rare—only 60 cases were found in the long-term experiences of the 35 countries; that is, less than 2 per country.[3] Therefore, to use history to gauge the probability and size distribution of macroeconomic disasters, it is hopeless to rely on the experience of a single country, such as the United States, even if we are willing to assume that the US economic structure remained fixed for 100 years or more.[4] In contrast, in long time series for a broad international sample, enough disaster realizations are available to allow for reasonably accurate inferences about disaster probabilities and size distributions. Underlying this calculation, of course, is the assumption that the underlying probability distributions are reasonably similar across countries, as well as roughly stable over time.

Table 1—Rates of Return for OECD Countries, 1880–2005

	Mean real rates of return	
Country	Stocks	Bills
Australia	0.103	0.012
Canada	0.077	0.016
Denmark	0.074	0.030
France	0.056	−0.011
Germany	0.073	−0.018
Italy	0.049	0.002
Japan	0.093	0.004
Norway	0.069	0.018
Sweden	0.091	0.023
UK	0.064	0.017
US	0.080	0.014
Means	**0.075**	**0.010**

Notes: Data on asset returns and consumer price indexes are from Global Financial Data, discussed in Bryan Taylor (2005). Real rates of return are calculated from arithmetic annual returns during each year, based on nominal total return indexes and consumer price indexes. In some cases, such as the United States before 1922, the data on bill returns are for commercial paper. For some country-years, stock returns are based on stock price indexes and estimates of dividend yields. Periods for returns are 1880–2005, except for the following missing data. Canada is missing stock returns for 1880–1915 and bill returns for 1880–1899. Denmark is missing stock returns for 1880–1914. France is missing stock returns for 1940–1941. Italy is missing stock returns for 1880–1905. Japan is missing stock returns for 1880–1914 and bill returns for 1880–1882. Norway is missing stock returns for 1880–1914. Sweden is missing stock returns for 1880–1901. The table excludes countries that were missing data on asset returns during major crises—Austria, Belgium, and the Netherlands around World Wars I and II; Finland, New Zealand, Portugal, and Switzerland around World War I; and Spain during the Spanish Civil War.

[2] The GDP data were from Maddison (2003). In the fruit-tree model, GDP and consumption coincide. More generally, consumption would be more appropriate than GDP for analyses of asset pricing and welfare costs. However, long-term data on real consumer expenditure are not reported by Maddison and are not readily available for many countries. An ongoing research project, described in Barro and Jose F. Ursua (2008a, b), involves the assembly of a dataset on long-term real personal consumer expenditure for as many countries as possible.

[3] The 60 cases exclude 5 postwar GDP contractions that did not seem to involve large declines in real personal consumer expenditure. The lower limit of 15 percent is arbitrary. Extending to 10 percent brings in another 21 contractions for the 35 countries. However, the inclusion of these smaller contractions has a minor effect on the results.

[4] Satyajit Chatterjee and Dean Corbae (2007, 1534) use the US history of the unemployment rate to note that there is "only one depression episode in the sample." From these data—and an assumption of unchanged economic structure since 1900—they infer a probability of moving from normalcy to depression of once every 83 years. This probability and the size distribution of depressions cannot be gauged accurately from this one time series. Moreover, they assume, without discussion, that real GDP always reverts to a deterministic trend line, although, as already noted, Cogley's (1990) and other analyses indicate that the data for most countries strongly reject this hypothesis. Kevin D. Salyer's (2007) analysis is similar in spirit to that of Chatterjee and Corbae.

For the 35 countries, the main global disasters were World War II (18 countries with large GDP contractions), the Great Depression (16 countries), World War I (13 countries), and post–World War II depressions in Latin America and Asia (11 country-events). The empirical frequency—60 events for 35 countries over 100 years—corresponds to a disaster probability, p, of 1.7 percent per year. (The disasters need not be independent across countries; in fact, they tend to congregate into events such as world wars, the Great Depression, the Asian financial crisis, and the Latin American debt crisis.)

The contraction proportion b for the observed twentieth century disasters ranged from 15 percent to 64 percent, with a mean of 29 percent.[5] However, with substantial risk aversion, the effective average value of b is substantially above the mean. For example, with a coefficient of relative risk aversion of four, a constant b of around 40 percent generates about the same equity premium and welfare effects as the empirically observed frequency distribution of b.

The formulation neglects rare bonanzas. With substantial risk aversion, bonanzas do not count nearly as much as disasters for the pricing of assets and for welfare effects. Moreover, long-term data on annual growth rates of per capita GDP tend to exhibit negative skewness. For 19 OECD countries from 1880 to 2005, 14 exhibit negative skewness, and the only substantially positive values are for France, the Netherlands, and Switzerland.

The expected growth rate of real GDP depends not only on the growth-rate parameter, g, but also on the uncertainty parameters. As the length of the period approaches zero, the specification in equation (1) implies that the expected growth rate of GDP and consumption, denoted by g^*, is given by

$$g^* = g + (½)\sigma^2 - p \times Eb, \tag{2}$$

where Eb is the expected value of b—0.29 in the sample of 60 observed crises. In practice, the term $(½)\sigma^2$ tends to be negligible—0.0002 in the calibrations considered later, for which $\sigma = 0.02$. However, the term $p \cdot Eb$ is nontrivial—0.0049 when $p = 0.017$ and $Eb = 0.29$. In this case, $g = 0.025$ corresponds to $g^* = 0.020$, the value used in the main calibrations.

I start with the familiar formulation where the representative consumer maximizes a time-additive utility function with iso-elastic preferences:

$$U_t = E_t \sum_{i=0}^{\infty} \frac{1}{(1+\rho)^i} [(C_{t+i})^{1-\gamma} - 1]/(1-\gamma), \tag{3}$$

where $\rho \geq 0$ is the rate of time preference. As is well known, this power-utility specification implies that the key parameter $\gamma > 0$ represents both the coefficient of relative risk aversion and the reciprocal of the intertemporal elasticity of substitution, henceforth denoted IES. This restriction matters for welfare-cost calculations, as observed by Obstfeld (1994), and also generates predictions about asset prices that are probably counterfactual, as argued by Bansal and Yaron (2004). Therefore, I soon generalize to a preference formulation—due to Epstein and Zin (1989) and Weil (1990)—that de-links the coefficient of relative risk aversion from the IES.

[5] The 29 percent figure refers to raw levels of per capita GDP. With an adjustment for trend growth, the mean contraction size was 35 percent.

Asset prices and rates of return can be determined in the usual way from the first-order conditions for consumption over time. With the power-utility formulation from equation (3), the familiar first-order conditions are

$$C_t^{-\gamma} = \left(\frac{1}{1+\rho}\right) E_t(R_t\, C_{t+1}^{-\gamma}), \tag{4}$$

where R_t is the gross return on any asset from time t to $t + 1$.

A key variable is the market value, V, of a tree that initially produces one unit of fruit. One way to calculate this value is to sum the prices of equity claims on future "dividends," $C_{t+i} = Y_{t+i}$. (In order to correspond to the summation in equation (3), it is convenient to treat C_t, rather than C_{t+1}, as the first payout on tree equity.) These prices can be determined readily from equation (4). As the arbitrary period length approaches zero, the reciprocal of V turns out to be

$$1/V = \rho + (\gamma - 1)\, g^* - (½)\, \gamma\, (\gamma - 1)\, \sigma^2 - p\left[E(1-b)^{1-\gamma} - 1 - (\gamma - 1)\, Eb\right], \tag{5}$$

where g^* is the expected growth rate (of dividends) from equation (2), $E(1-b)^{1-\gamma}$ is the expectation of $(1-b)^{1-\gamma}$, and Eb is the expectation of b. The variable V corresponds to the price-dividend ratio for an unlevered equity claim on a tree.

Given the pricing formula in equation (5), the expected rate of return on unlevered equity can be determined (when the period length approaches zero) to be

$$r^e = \rho + \gamma g^* - (½)\, \gamma\, (\gamma - 1)\, \sigma^2 - p\left[E(1-b)^{1-\gamma} - 1 - (\gamma - 1)\, Eb\right]. \tag{6}$$

Therefore, the right-hand side of equation (5) is the difference between r^e and g^*. The transversality condition, which guarantees that the market value of a tree is positive and finite, is that this right-hand side be positive, that is, $r^e > g^*$.

The risk-free rate, r^f, can also be determined from equation (4). The result (when the period length approaches zero) is

$$r^f = \rho + \gamma g^* - (½)\, \gamma\, (\gamma + 1)\, \sigma^2 - p\left[E(1-b)^{-\gamma} - 1 - \gamma\, Eb\right]. \tag{7}$$

(Since the model has i.i.d. shocks, the term structure of risk-free rates is flat; that is, r^f is the short-term and long-term risk-free rate.) Depending on the uncertainty parameters—particularly p and the distribution of b—r^f can be very low. In fact, r^f can be less than g^* and even less than zero. The equity premium is

$$r^e - r^f = \gamma\sigma^2 + pE\{b\left[(1-b)^{-\gamma} - 1\right]\}. \tag{8}$$

The first term, $\gamma\sigma^2$, tends to be very small and corresponds to the no-disaster equity premium of Mehra and Prescott (1985). The second term brings in disasters and is proportional to the disaster probability, p. The disaster size, b, enters as the expectation of the product of b and the proportionate excess of the marginal utility of consumption in a disaster state, $\left[(1-b)^{-\gamma} - 1\right]$.

Table 1 shows average real rates of return on stocks and government bills from 1880 to 2005 for 11 OECD countries that have the necessary long-term data. The equity premium, in the sense of the difference between the two average rates of return, is 0.065 per year. Since the stock returns reflect leverage, the premium for unlevered equity would be smaller. For example, with a debt-equity ratio of 0.5 (corresponding to recent US values for nonfinancial corporations), the predicted premium would be (0.065/1.5) = 0.043.

For the model to get into the right ballpark for explaining the equity premium, the coefficient of relative risk aversion, γ, has to be well above one. Barro (2006) showed, for plausible values of the uncertainty parameters, especially p and the distribution of b, that $\gamma = 4$ was satisfactory.[6] In any event, γ could not be less than about three.

One difficulty is that, if $\gamma > 1$, equation (5) delivers the likely counterfactual prediction that an increase in uncertainty (higher σ or p or a shift in the distribution of b toward larger values), for given g^*, implies a *higher* price-dividend ratio, V. Bansal and Yaron (2004, 1487) make an analogous observation about the connection between the volatility of consumption growth and the price-dividend ratio in their model. The prediction for a positive relationship between the extent of uncertainty and the price-dividend ratio conflicts with the usual view that an increase in aggregate uncertainty tends to depress stock prices. The reason that the model makes this counterintuitive prediction is that, with power utility, the IES is constrained to equal the reciprocal of the coefficient of relative risk aversion. Therefore, I now relax this restriction (as do Bansal and Yaron 2004) by adopting the preference specification of Epstein and Zin (1989) and Weil (1990).

Using a minor modification of the Weil (1990) formulation, the extended utility formula is

$$U_t = \frac{\{C_t^{1-\theta} + (1/(1+\rho))[(1-\gamma)E_tU_{t+1}]^{(1-\theta)/(1-\gamma)}\}^{(1-\gamma)/(1-\theta)}}{(1-\gamma)}, \tag{9}$$

where $1/\theta > 0$ is the IES and $\gamma > 0$ is the coefficient of relative risk aversion. Equation (3) is the special case of equation (9) when $\theta = \gamma$.

In general, EZW preferences do not allow for simple formulas for pricing assets. However, when the underlying shocks are i.i.d., as already assumed, the analysis simplifies dramatically. A key property of the solution under i.i.d. shocks is that attained utility, U_t, ends up as a simple function of contemporaneous consumption, C_t:

$$U_t = \Phi C_t^{1-\gamma}, \tag{10}$$

where the constant Φ depends on the parameters of the model.[7]

The application of a standard perturbation argument to equation (10) leads to the first-order conditions for utility maximization:

$$C_t^{-\gamma} = \left(\frac{1}{1+\rho^*}\right) E_t(R_t\, C_{t+1}^{-\gamma}), \tag{11}$$

where R_t is the gross, one-period return on any asset. As usual, these first-order conditions will be the basis for asset pricing. Thus, an important result is that, with i.i.d. shocks, the conditions for

[6] That analysis also took account of partial default on bills, typically due to high wartime inflation.

[7] Alberto Giovannini and Philippe Weil (1989, appendix) show that, with the utility function in equation (9), attained utility, U_t, is proportional to wealth raised to the power $1-\gamma$. The form in equation (10) follows because C_t is optimally chosen as a constant ratio to wealth in the i.i.d. case. The formula for Φ is, if $\gamma \neq 1$ and $\theta \neq 1$,

$$\Phi = \left(\frac{1}{1-\gamma}\right)\left\{\rho + (\theta-1)g^* - (½)\gamma(\theta-1)\sigma^2 - \left(\frac{\theta-1}{\gamma-1}\right)p\,[E(1-b)^{1-\gamma} - 1 - (\gamma-1)Eb]\right\}^{(\gamma-1)/(1-\theta)}.$$

asset pricing under EZW preferences look similar to those in the power-utility model, described by equation (4). However, two key features of the EZW results are worth stressing. First, the exponents on C_t and C_{t+1} in equation (11) involve γ, the coefficient of relative risk aversion, not θ, the reciprocal of the IES. Second, the effective rate of time preference, ρ^*, differs from ρ when γ and θ diverge. The formula for ρ^* is, if $\gamma \neq 1$,

$$(12) \qquad \rho^* = \rho - (\gamma - \theta)\left\{g^* - (½)\,\gamma\sigma^2 - \left(\frac{p}{\gamma - 1}\right)[E(1-b)^{1-\gamma} - 1 - (\gamma - 1)\,Eb]\right\}.$$

In this and subsequent cases, results when γ (or, subsequently, θ) approach one can be derived from standard limit calculations. Note from equation (12) that ρ^* depends not only on preference parameters—ρ, γ, and θ—but also on parameters that describe expected growth and uncertainty—g^*, σ, p, and the distribution of b.

The results imply that, in the i.i.d. case, asset-pricing formulas derived under EZW preferences coincide with formulas under power utility if ρ^* replaces ρ. In particular, the formulas for V, r^e, and r^f in equations (5)–(7) remain valid with the substitution of ρ^* for ρ. Therefore, in the EZW case, the IES, $1/\theta$, affects the price-dividend ratio (equation (5)) and levels of rates of return (equations (6) and (7))—through influences on ρ^*—but not the equity premium (equation (8)). The equity premium depends on the coefficient of relative risk aversion, γ, exactly as in the power-utility case. Since the power-utility model accorded reasonably well with observed equity premia when $\gamma = 4$, it follows that the EZW specification fits the equity premium just as well when $\gamma = 4$.

With EZW preferences, the formula for the price-dividend ratio, V, in equation (5) becomes, after replacement of ρ by ρ^* from equation (12),

$$(13) \qquad 1/V = \rho + (\theta - 1)g^* - (½)\gamma(\theta - 1)\sigma^2 - p\left(\frac{\theta - 1}{\gamma - 1}\right)[E(1-b)^{1-\gamma} - 1 - (\gamma - 1)\;Eb],$$

if $\gamma \neq 1$. For any $\gamma > 0$, the condition $\theta < 1$ implies that, with g^* held fixed, V is lower if uncertainty is greater (higher σ or p or a shift of the b-distribution toward higher values). That is, a once-and-for-all increase in an uncertainty parameter reduces stock prices (as seems plausible) if and only if $\theta < 1$, so that the IES > 1.[8]

Equation (13) also implies, if $\theta < 1$, that V is higher if the mean growth rate, g^*, is higher, for given uncertainty parameters. This condition is important in Bansal and Yaron (2004), who propose to explain the equity premium not by disaster risk but rather by shocks to their counterpart of g^*. They also allow for a time-varying variance of these shocks. One limitation of their approach is that quantitative success depends on high risk aversion. The coefficient γ has to be around ten to account for observed equity premia in their model. Thus, my inference is that

[8] Although equation (13) applies to unlevered equity, the effects from once-and-for-all variations in uncertainty parameters tend also to hold for levered equity. If firms issue a mixture of equity and risk-free debt, V applies to the total market value of the firm. Then, if the debt is short term, a one-time change in an uncertainty parameter (σ, p, or the distribution of b) has a negligible effect on the market value of the outstanding debt. In this case, changes in the value of levered equity absorb the full change in the firm's overall market value. Thus, we still have that more uncertainty lowers the price of levered equity if and only if $\theta < 1$. For longer-term debt, a once-and-for-all shift in an uncertainty parameter also generates changes in the market value of the firm's bonds. If the debt is risk-free and the probability of future changes in uncertainty parameters is nil, then greater uncertainty would raise the market value of the outstanding bonds. In this case, the value of levered equity falls by more than the decline in the overall market value of the firm. Hence, the price of levered equity would decrease even if θ is somewhat greater than one. In more general situations—for example, when the firm's debt has default risk—this result can change because greater uncertainty might lower the market value of the firm's existing bonds.

fluctuating long-run growth rates, g^*, may usefully supplement analyses that include disaster risk but probably cannot be the main basis for explaining the equity premium.

Given the GDP process in equation (1), data on rates of return, such as r^e and r^f, and the price-dividend ratio, V, pin down γ and the effective rate of time preference, ρ^*. Since ρ^* depends on a combination of ρ and θ (in equation (12)), the data would not allow separate identification of ρ and θ, a finding that relates to the observational-equivalence point of Narayana R. Kocherlakota (1990). However, the parameters ρ and θ could be separately identified from other information; for example, if we know how V responds to one-time changes in the uncertainty parameters—σ, p, and the distribution of b—or the expected growth rate, g^*, in equation (13). Alternatively, in the model with endogenous saving considered in Section IV, identification would follow from information about how the saving ratio reacts to one-time changes in σ, p, and the distribution of b.

To summarize, the model with Epstein-Zin-Weil preferences, disaster risk, and i.i.d. shocks can accord with some central asset-pricing "facts." First, if the coefficient of relative risk aversion, γ, is around four, the equity premium and risk-free rate can be roughly correct. Second, if $\theta < 1$, so that the IES is greater than one, the price-dividend ratio, V, relates to aggregate uncertainty and expected growth in the "right" directions—more uncertainty goes along with lower V and higher expected growth goes along with higher V.

The simplicity of the underlying structure (i.i.d. shocks, representative consumer, closed economy with no investment) allows for a closed-form solution for attained utility, U_t, as a function of the underlying parameters of preferences and the output process. Obstfeld (1994) derived analogous closed forms in a model without disaster risk. A convenient representation uses equations (5) and (10) to express U_t as a function of the price-dividend ratio, V. The formula, when $\gamma \neq 1$ and $\theta \neq 1$, is, up to an inconsequential additive constant:[9]

$$U_t = \left(\frac{1}{1-\gamma}\right) V^{(1-\gamma)/(1-\theta)}\, Y_t^{1-\gamma}. \tag{14}$$

Equation (14), in conjunction with equation (13), allows for assessments of the welfare effects of uncertainty.

II. Calculation of Welfare Effects

Equations (13) and (14) determine the dependence of attained utility, U_t, on the expected growth rate, g^*, and the parameters that govern consumption risk: σ, p, and the distribution of b. These effects can be compared with those from proportionate shifts in the initial level of GDP and consumption, Y_t.

A. *Local Effects on Welfare*

The marginal effect on utility from a proportionate change in Y_t is given from equation (14) by

$$\frac{\partial U_t}{\partial Y_t} Y_t = V^{(1-\gamma)/(1-\theta)} (Y_t)^{1-\gamma}. \tag{15}$$

[9] The form of equation (14) does not depend on the particular stochastic process for output in equation (1). However, the constancy of the price-dividend ratio, $V_t = V$, depends on the i.i.d. form of the shocks, u_t and v_t. A constant V conflicts with the observed volatility of price-dividend ratios for stock-market claims. The model can match this volatility if the parameters of uncertainty, such as the disaster probability, p_t, move around. Xavier Gabaix (2008) shows that the main asset-pricing results go through if p_t evolves exogenously in nearly random-walk-like fashion.

The marginal effect from a change in g^* follows from equations (13) and (14) as

$$\frac{\partial U_t}{\partial g^*} = V^{[1+(1-\gamma)/(1-\theta)]}(Y_t)^{1-\gamma}. \tag{16}$$

Therefore, the utility rate of transformation between proportionate changes in Y_t and changes in g^* is given by

$$\frac{-\partial U_t/\partial g^*}{(\partial U_t/\partial Y_t)\, Y_t} = -V. \tag{17}$$

This result gives the proportionate decrease in Y_t that compensates, at the margin, for an increase in g^*—in the sense of preserving attained utility. Equation (17) shows that this compensating output change depends only on the combination of parameters that enter into the price-dividend ratio, V, determined in equation (13).

To pin down a reasonable magnitude for V, start with the already mentioned specification $p = 0.017$ per year. This and subsequent calibration parameters are collected in Table 2. The probability distribution for b is the historical one mentioned before, for which $Eb = 0.29$. Some other baseline parameters are the same as those used in the main calibration exercise in Barro (2006, table 5). The coefficient of relative risk aversion is $\gamma = 4$, the standard deviation of the u_t shocks is $\sigma = 0.020$ per year, the growth-rate parameter is $g = 0.025$ per year, and the expected growth rate is $g^* = 0.020$ per year (from equation (2)).[10] Since this earlier exercise assumed power-utility preferences, where $\theta = \gamma$, the IES, $1/\theta$, was constrained to be 0.25. As already mentioned, an IES this low produces implausible results concerning effects of uncertainty and growth parameters on the price-dividend ratio. The EZW case now being considered requires a separate calibration for the IES.

Macroeconomic estimates of the IES, $1/\theta$, represented by Hall (1988), come from regressions of consumption growth rates on real rates of return, for example, on short-term real interest rates. The resulting estimates of $1/\theta$ cover a broad range and tend to be well below one. However, as observed by Bansal and Yaron (2004, 1501) and Barro (2005, sect. VIII), these coefficient estimates tend to be biased sharply toward zero because sample fluctuations in real interest rates likely reflect, to a considerable extent, variations in uncertainty parameters (and the usual instrumentation using lagged variables does not help in this respect). This type of regression approach with macroeconomic data yields satisfactory estimates of $1/\theta$ only if the fluctuations in real interest rates stem mainly from movements in the expected growth rate, g^*, for given uncertainty parameters.

Because of the shortcomings of macroeconomic estimates of the IES, it is worthwhile to consider microeconomic evidence. The Jonathan Gruber (2006) analysis is particularly attractive because it uses cross-individual differences in after-tax real interest rates that derive from arguably exogenous differences in tax rates on capital income. For present purposes, the key point is that the Gruber estimate of the IES is around 2.0. Thus, the baseline calibration assumes $\theta = 0.5$.

[10] The values for g and σ come from data on real personal consumer expenditure for 21 OECD countries for 1954–2005, a tranquil period with no disaster events for these countries. The largest contraction was 14 percent for per capita real consumer expenditure (12 percent for per capita GDP) for Finland for 1989–1993. For 1954–2005, the median of the growth rates of real per capita personal consumer expenditure for the 21 countries was 0.026 per year, and the median standard deviation of the growth rates was 0.024. The US values were 0.024 and 0.018, respectively. With $\gamma = 4$, the expectations associated with the historical distribution of disaster sizes, b, are $Eb = 0.29$, $E(1-b)^{-\gamma} = 7.69$, and $E(1-b)^{1-\gamma} = 4.05$.

TABLE 2—PARAMETERS IN BASELINE CALIBRATION EXERCISE

Parameter	Value
p: disaster probability (per year)	0.017
b: disaster size—uses historical frequency distribution	mean = 0.29
g: growth-rate parameter (per year)	0.025
g^*: expected growth rate (per year), from equation (2)	0.020
σ: s.d. of normal growth-rate fluctuations (per year)	0.020
γ: coefficient of relative risk aversion	4
θ: reciprocal of intertemporal elasticity of substitution (IES)	0.5
ρ: rate of time preference (per year)	0.052
ρ^*: effective rate of time preference (per year), from equation (12)	0.027

The final parameter needed is the rate of time preference, ρ. The main calibrations in Barro (2006) used $\rho = 0.030$ per year. However, the pure rate of time preference is not directly observable. Typically, a reasonable value for ρ is inferred from its connection to levels of rates of return, including the risk-free rate. Thus, a first point is that, in the EZW context, the link to rates of return involves the effective rate of time preference, ρ^*, given in equation (12), not ρ, per se. Hence, I proceed by assuming that ρ takes on a value that, given the other baseline parameters, generates a ρ^* that yields a plausible risk-free rate. This procedure means that fitting the risk-free rate is not a test of the model.

Table 1 shows that the real rate of return on government bills (or analogous short-term claims) for 11 OECD countries from 1880 to 2005 averaged 0.010 per year. These bill returns are not risk-free and include some low realizations associated with war-related inflations (such as in Germany around World War I). Therefore, risk-free rates (not directly observed) would likely be somewhat lower than the average real rate of return on bills. However, I take 0.010 as an approximation to the risk-free rate. Given the other baseline parameters, it turns out that $\rho^* = 0.027$ is required to generate $r^f = 0.010$ in the model (from equation (7) with ρ^* substituted for ρ). Equation (12) then implies that the required value of ρ is 0.052.

The full set of baseline parameters, shown in Table 2, generates a price-dividend ratio, V, in equation (13) of 20.7. This value for V implies that a small rise in the expected growth rate, g^*—for example, by 0.1 percent per year—has to be compensated by a fall in the initial level of GDP, Y_t, by 2.1 percent. Despite differences in specification, this result accords with the one reported by Lucas (1987, ch. 3). An economy should be willing to give up a lot in its initial level of GDP to obtain a small increase in its long-term growth rate.

The Lucas calculations about consumption uncertainty relate in the present model to the parameter σ. The marginal effect on attained utility, U_t, from a change in σ is given from equations (13) and (14) by

$$\frac{\partial U_t}{\partial \sigma} = -(Y_t)^{1-\gamma}\, V^{(\theta-\gamma)/(1-\theta)}\, \gamma\sigma V^2. \tag{18}$$

Therefore, equation (15) implies that the utility rate of transformation between proportionate changes in Y_t and changes in σ is given by

$$\frac{-(\partial U_t/\partial\sigma)}{(\partial U_t/\partial Y_t)\, Y_t} = \gamma\sigma V. \tag{19}$$

This expression gives the proportionate increase in initial GDP required to compensate for a small rise in σ.[11] The parameters specified before imply $\gamma\sigma V = 1.66$. Therefore, to maintain attained utility, an increase in σ by, say, 10 percent (from 0.020 to 0.022) requires a rise in the initial level of GDP by approximately 0.33 percent. Since the expected growth rate, g^*, is held fixed, this proportionate rise in GDP level should be viewed as applying each year.[12]

These calculations apply for small changes in σ. Large changes, considered in the next section, recognize that the utility rate of transformation tends to rise with σ on the right-hand side of equation (19). This consideration means that the welfare gain from reducing σ from 0.020 to zero is smaller in magnitude than the amount—3.3 percent—that would be calculated from equation (19) if the utility rate of transformation were constant.

Consider now the welfare consequences from a change in the disaster probability, p, for a given distribution of disaster sizes, b. Equations (13) and (14) imply

$$(20) \qquad \frac{\partial U_t}{\partial p} = -(Y_t)^{1-\gamma}\, V^{(\theta-\gamma)/(1-\theta)}\, V^2\left[E(1-b)^{1-\gamma} - 1 - Eb\,(\gamma-1)\right]/(\gamma-1).$$

This formula applies while holding fixed the expected growth rate, g^*; that is, it does not allow for the negative effect of p on g^*, for given g, in equation (2). The utility rate of transformation between proportionate changes in Y_t and changes in p is given by

$$(21) \qquad \frac{-(\partial U_t/\partial p)}{(\partial U_t/\partial Y_t)\, Y_t} = V\left[E(1-b)^{1-\gamma} - 1 - Eb\,(\gamma-1)\right]/(\gamma-1).$$

With the parameter values used before, the right-hand side equals 15.1. As before, the result applies to small changes. An increase in p by 10 percent (from 0.0170 to 0.0187) matches up approximately with a proportionate rise in initial GDP by 2.6 percent. Again, this change in GDP level applies each year.

We can modify the calculations to allow for a growth effect from a change in p; that is, for given g, g^* falls with p in equation (2).[13] The result modifies equation (21) to

$$(22) \qquad \frac{-(\partial U_t/\partial p)}{(\partial U_t/\partial Y_t)\, Y_t}\,(\textit{incl. growth effect}) = V\left[E(1-b)^{1-\gamma} - 1\right]/(\gamma-1).$$

With the same parameter values as before, the right-hand side equals 21.0. Therefore, a rise in p by 10 percent now matches up with a proportionate increase in GDP by 3.6 percent— larger than before because of the decline in g^*.

B. *Welfare Effects from Large Changes*

Equations (17), (19), (21), and (22) assess welfare effects from small changes in Y_t, g^*, σ, and p. We can instead use equations (13) and (14) to assess the effects on attained utility from large

[11] We can instead compute the increase in the expected growth rate, g^*, needed to compensate for a small rise in σ. Equations (17) and (19) imply that the answer is $\gamma\sigma$, which does not involve V.

[12] Obstfeld (1994) observes that Lucas (1987, ch. 3) gets far smaller estimates for the welfare cost of consumption uncertainty because he treats the shock, analogous to u_t in the present model, as a transitory disturbance to the level of output.

[13] See Gadi Barlevy (2004) for a discussion of models in which uncertainty affects the expected growth rate of GDP.

changes. Let V and Y_t be values that apply for the baseline specification of parameters in Table 2. Let V^* and $(Y_t)^*$ be values that apply in an alternative situation that delivers the same attained utility, U_t. Then, the formula for U_t in equation (14) implies[14]

$$(Y_t)^*/Y_t = (V/V^*)^{1/(1-\theta)}. \tag{23}$$

The result in equation (23) relates to Alvarez and Jermann (2004), who try to go as far as possible to gauge the welfare costs of consumption uncertainty by observing or estimating various asset prices.[15] Equation (23) provides insight for the present model on the extent to which welfare costs can be assessed from observations of asset prices related to equity shares. The price V may be observable—in the Lucas-tree economy, V is the price-dividend ratio for an unlevered equity claim on a tree. However, the price V^* is unlikely to be observable: V^* is the price-dividend ratio for unlevered tree equity in a hypothetical economy, such as one with zero uncertainty.

If the hypothetical price-dividend ratio, V^*, could be observed or estimated, equation (23) shows that the welfare gain, measured by the compensating output change $(Y_t)^*/Y_t$, depends on the parameter θ, for given V and V^*. The baseline specification in Table 2 uses the value $\theta = 0.5$, corresponding to an IES of two. For this case, equation (23) implies an inverse-square law for the relation between the welfare effect and the ratio of equity prices:

$$(Y_t)^*/Y_t = (V/V^*)^2. \tag{24}$$

Suppose, for example, that a reduction in uncertainty (a decrease in σ or p or a shift in the b-distribution toward smaller values) results in an increase in the price-dividend ratio by 1 percent—that is, V^* is 1 percent above V. In this case, the compensating output change is about 2 percent—that is, $(Y_t)^*$ is roughly 2 percent below Y_t.

Lucas (1987, ch. 3; 2003, sect. II) focused on the consequences of eliminating all consumption uncertainty associated with usual business fluctuations—in the present context, this exercise corresponds to setting $\sigma = 0$. The formula for V in equation (13) implies for this case

$$1/V^* = 1/V + (½)\,\gamma\,(\theta - 1)\,\sigma^2.$$

Substitution into equation (23) yields

$$(Y_t)^*/\,Y_t = [1 + (½)\,\gamma\,(\theta - 1)\,\sigma^2\,V]^{1/(1-\theta)}. \tag{25}$$

Suppose that the magnitude of $(½)\,\gamma\,(\theta - 1)\,\sigma^2\,V$ is much less than one—a condition likely to hold because $(½)\,\gamma\,(\theta - 1)\,\sigma^2\,V = -0.0083$ in the baseline specification. In this case, the result in equation (25) simplifies to

$$\log[(Y_t)^*/Y_t] \approx -(½)\,\gamma\,\sigma^2 V. \tag{26}$$

[14] Equation (23) determines the compensating income change in the sense of John R. Hicks (1946, 330–31) for a shift in a parameter, such as g^*, σ, or p.

[15] Part of the Alvarez-Jermann analysis depends on the pricing of a claim to a "consumption trend." The price of such a claim is finite only if the risk-free rate, r^f, exceeds the expected growth rate, g^*. This condition need not hold in my model; that is, $r^f < g^*$ can apply in equation (7). Moreover, the data in Table 1 indicate that the average real rate of return on government bills, 0.010, was below the long-term average growth rate for OECD countries. These growth rates averaged around 0.020 for per capita GDP and consumption and 0.030 for levels of GDP and consumption. Under these circumstances, the price of an Alvarez-Jermann consumption-trend claim is infinity.

That is, the welfare benefit from reducing σ to zero is approximately one-half the effect that would be calculated from the local impact of a change in σ given by equation (19).

With the parameter values assumed before, equation (25) implies that $(Y_t)^*$ is 1.6 percent below Y_t. That is, society would be willing to give up 1.6 percent of output each year to eliminate all of the customary economic fluctuations represented by σ. As noted before (fn. 12), this effect is much larger than that found by Lucas (1987), mainly because the impact of a shock, u_t, on the GDP level is permanent in the present model.

Setting the disaster probability, p, to zero (or, equivalently, the disaster size, b, to zero) has much greater consequences for welfare. The formula, derived from equations (13) and (23), is

$$(27) \qquad \frac{Y_t^*}{Y_t} = \left\{1 + pV\left(\frac{\theta - 1}{\gamma - 1}\right)\left[E(1-b)^{1-\gamma} - 1 - (\gamma - 1)\,Eb\right]\right\}^{1/(1-\theta)}.$$

Note that this formula holds fixed the expected growth rate, g^*; that is, it does not allow for the inverse relation between p and g^*, for given g, in equation (2). With the same parameter values as before, $(Y_t)^*$ is 24.0 percent below Y_t. Hence, when gauged by the compensating proportionate change in output, eliminating disaster risk is worth 15 times as much as eliminating normal economic fluctuations.

These large welfare costs of disasters arise even though the present analysis considers only the utility lost from reduced consumption. For wars, natural disasters, and epidemics, an allowance for the direct utility losses from death, injury, and disease would raise the welfare effects. See Gregory D. Hess (2003) for a discussion in the context of conflicts.

We can again modify the calculations to allow for a growth effect from a change in p; that is, for given g, g^* falls with p in equation (2). The revised formula for the welfare gain is

$$(28) \qquad \frac{Y_t^*}{Y_t} = \left\{1 + pV\left(\frac{\theta - 1}{\gamma - 1}\right)\left[E(1-b)^{1-\gamma} - 1\right]\right\}^{1/(1-\theta)}.$$

With the usual parameter values, $(Y_t)^*$ is 32.5 percent below Y_t. This result is larger than before because the reduction in p raises g^*.

We can also consider the elimination of all consumption uncertainty by setting $\sigma = 0$ and $p = 0$ (or $b = 0$) simultaneously. If g^* is held fixed, $(Y_t)^*$ is 25.4 percent below Y_t. Allowing for the inverse relation between p and g^*, the result is 33.9 percent. These results correspond, as a good approximation, to the sum of the effects from setting $\sigma = 0$ and $p = 0$ separately. Thus, the main effects in each case come from setting $p = 0$.

III. Sensitivity of the Welfare-Cost Estimates

The welfare estimates, including the effects from eliminating all disaster risk, depend on the coefficient of relative risk aversion, γ, and the IES, $1/\theta$. Table 3 shows how the computed welfare effects depend on these preference parameters. The line shown in bold, where $\gamma = 4$ and $\theta = 0.5$, is the baseline specification already discussed.

The first four lines of Table 3 show the impact of raising θ, while maintaining $\gamma = 4$. One complication is that, for given ρ, changes in θ influence the effective rate of time preference, ρ^*, given in equation (12). The spirit of the calibration exercise was to choose ρ to generate a ρ^* that produced reasonable levels of rates of return, including the risk-free rate. To accord with this perspective, ρ is varied in the table each time θ or γ changes to keep ρ^* at its baseline value, 0.027. For example, for $\gamma = 4$, $\rho = 0.054$ when $\theta = 0.25$, 0.052 when $\theta = 0.50$, 0.048 when

$\theta = 1$, and 0.027 when $\theta = 4$. Since γ and ρ^* are held fixed, the rates of return, r^e and r^f, and the price-dividend ratio, V, do not change as θ varies. For example, the equity premium remains fixed at 0.059 in these cases.

The general pattern in Table 3 is that an increase in θ—implying a decrease in the IES—lowers the welfare benefits from eliminating uncertainty. However, for a given γ, since ρ^* is held constant, an increase in θ—say from 0.25 to 4—has only a minor effect on the welfare gain from setting σ to zero. This result is apparent from equation (26) because, as an approximation, the benefit does not depend on θ, for given γ and V. For example, when $\gamma = 4$, the welfare gain from setting $\sigma = 0$ declines only slightly from 1.65 percent of output at $\theta = 0.25$ to 1.60 percent at $\theta = 4$. The negative effect from raising θ on welfare is more pronounced for setting $p = 0$. For example, when $\gamma = 4$, the benefit decreases from 24.7 percent of output at $\theta = 0.25$ to 22.6 percent at $\theta = 1$ and 17.3 percent at $\theta = 4$. If we restrict attention to the range where $\theta < 1$, so that the IES > 1, the changes in θ have relatively small consequences for the welfare effects.

Table 3 shows, not surprisingly, that decreases in the coefficient of relative risk aversion, γ, reduce the welfare benefit from eliminating uncertainty. These effects are more important than those from changing θ (given that ρ^* is maintained at 0.027 in all cases). For example, if θ is fixed at 0.50, the welfare benefit from setting $\sigma = 0$ declines from 1.65 percent of output when $\gamma = 4$ to 1.30 percent at $\gamma = 3.5$, 1.12 percent at $\gamma = 3$, and 0.74 percent at $\gamma = 1$. The corresponding gain from setting $p = 0$ falls from 24.0 percent to 16.1 percent, 11.8 percent, and 4.6 percent. Thus, the large estimated welfare gains from eliminating disaster risk depend on agents having a substantial degree of risk aversion.

A problem with the calculations for low values of γ is that the predicted equity premium deviates from observed values, which suggest an unlevered premium of 4–5 percent. Table 3 shows that the model's predicted premium is 5.9 percent at $\gamma = 4$, 3.9 percent at $\gamma = 3.5$, 2.6 percent at $\gamma = 3$, and only 0.3 percent at $\gamma = 1$. Hence, even with the presence of disaster risk, the predictions deviate from observed equity premia unless γ is at least 3.5. The model's implications for welfare costs of uncertainty likely should not be taken seriously in the range of values for γ where the model fails to get into the right ballpark for explaining the equity premium. Thus, it seems best to focus on welfare effects corresponding to a value for γ of at least 3.5. For this case, when $\theta = 0.5$, the welfare gain from setting $p = 0$ is 16.1 percent of output.

It is possible to restore reasonable predictions for the equity premium at lower values of γ if the disaster probability, p, is raised substantially above 1.7 percent per year. For example, at $\gamma = 3$, p has to be 4.1 percent to generate the same equity premium, 5.9 percent, as in the baseline case. With this unrealistically high p, the elimination of all disaster risk (setting p or b to zero) turns out to balance against a proportionate decline in output by 60 percent, well above the 24 percent calculated originally.

IV. Endogenous Labor Supply

The model can be extended to encompass a simple model of productive labor and labor-leisure choice. Suppose that the output of each tree is given by

$$Y_t = A_t L_t^{\alpha}, \tag{29}$$

where A_t is exogenous productivity, L_t is the quantity of labor employed, and $0 < \alpha < 1$. The log of productivity is generated in the same way as output in the baseline Lucas-tree model; that is, $\log(A_{t+1})$ follows the stochastic process given by the form of equation (1). Thus, the underlying uncertainty in this model is the same as in the original setting. All labor is equally productive

TABLE 3—EFFECTS OF PREFERENCE PARAMETERS ON RATES OF RETURN AND WELFARE COSTS

							Welfare effects (percent)	
γ	θ	ρ	ρ^*	r^e	r^f	V	$\sigma = 0$	$p = 0$
4	0.25	0.054	0.027	0.069	0.010	20.7	1.65	24.7
4	**0.50**	**0.052**	**0.027**	**0.069**	**0.010**	**20.7**	**1.65**	**24.0**
4	1	0.048	0.027	0.069	0.010	20.7	1.64	22.6
4	4	0.027	0.027	0.069	0.010	20.7	1.60	17.3
3.5	0.25	0.062	0.027	0.074	0.035	18.7	1.31	16.5
3.5	0.50	0.059	0.027	0.074	0.035	18.7	1.30	16.1
3.5	1	0.054	0.027	0.074	0.035	18.7	1.30	15.5
3.5	4	0.022	0.027	0.074	0.035	18.7	1.27	12.7
3	0.25	0.063	0.027	0.074	0.048	18.7	1.12	12.0
3	0.50	0.060	0.027	0.074	0.048	18.7	1.12	11.8
3	1	0.053	0.027	0.074	0.048	18.7	1.12	11.5
3	4	0.014	0.027	0.074	0.048	18.7	1.10	9.9
1	0.25	0.041	0.027	0.047	0.044	37.1	0.74	4.7
1	0.50	0.036	0.027	0.047	0.044	37.1	0.74	4.6
1	1	0.027	0.027	0.047	0.044	37.1	0.74	4.6
1	4	−0.030	0.027	0.047	0.044	37.1	0.73	4.3

Notes: The baseline results are in bold, γ is the coefficient of relative risk aversion, θ is the reciprocal of the IES in the formula for utility in equation (9), ρ is the rate of time preference, and ρ^* is the effective rate of time preference, given in equation (12); ($\rho = \rho^*$ holds when $\gamma = \theta$). The formulas for the expected rate of return on equity, r^e, the risk-free rate, r^f, and the price-dividend ratio, V, are given in equations (6), (7), and (5), respectively, after replacing ρ by ρ^*. The value of ρ^* is set at 0.027 to generate $r^f = 0.010$ with the baseline parameters. The value for ρ (0.052 in the baseline specification) is then varied in each case to maintain $\rho^* = 0.027$ (in equation (12)). Since ρ^* is held constant, the values for r^e, r^f, and V depend on γ but not on θ. Each welfare effect gives the percentage reduction in initial output, $1 - (Y_t)^*/Y_t$, that maintains attained utility while setting to zero either the standard deviation, σ, of normal economic fluctuations or the disaster probability, p. The effects are for a given expected growth rate, g^*, given in equation (2). The values for $1 - (Y_t)^*/Y_t$ come from equation (23).

and earns the common real wage rate, w_t. Since the labor market is competitive, w_t equals the marginal product of labor, determined from equation (29).

Each person is endowed with one unit of time, which can be allocated between leisure and market work. Utility now depends on each period's consumption, C_t, and leisure, $1 - L_t$. One straightforward way to model preferences is to use the Epstein-Zin-Weil formulation of utility from equation (9), but replace $C_t^{1-\theta}$ by $[C_t(1 - L_t)^\lambda]^{1-\theta}$.[16] The new parameter $\lambda > 0$ is the constant elasticity of substitution between consumption and leisure at a point in time. This form is consistent with the prescription of Robert G. King, Charles I. Plosser, and Sergio Rebelo (1988) that preferences accord with the property that work effort, L_t, be constant in the long run, that is, when w_t and C_t advance at the same rate due to steady productivity growth. In the present setting, which lacks capital accumulation, this property also holds in the short run, so that L_t ends up constant in equilibrium.

The new set of first-order conditions involves substitution between leisure and consumption at each point in time:

$$\left[\frac{\partial u/\partial(1 - L)}{\partial u/\partial C}\right]_t = w_t. \tag{30}$$

[16] The basic results go through under the more general specification $[C_t \cdot \omega(L_t)]^{1-\theta}$, where the function ω satisfies $\omega(L_t) > 0$ and $\omega'(L_t) < 0$.

Given the assumed form of the utility function, these conditions imply

$$L_t = 1 - \lambda\,(C_t/w_t), \tag{31}$$

which can be viewed as a labor-supply function.

The production function in equation (29) and the condition $C_t = Y_t$ imply

$$C_t/w_t = (1/\alpha)\,L_t. \tag{32}$$

This result, in conjunction with equation (31), implies

$$L_t = \alpha/(\alpha + \lambda). \tag{33}$$

Hence, the fraction of time worked is constant—invariant with shocks to productivity, A_t. This result applies because substitution effects (from changing w_t) exactly offset income effects (associated with changing C_t).

Since L_t is constant, output and consumption, $Y_t = C_t$, and dividends paid on equity claims are all proportional to A_t. Therefore, the pricing of equity claims (and other claims) is the same as in the initial model.

Raj Chetty (2006) shows, within an expected-utility setting, that labor-supply elasticities and the extent of leisure-consumption complementarity imply restrictions on the admissible range for the coefficient of relative risk aversion—the parameter γ in equation (9). In particular, he argues that empirical estimates of income-compensated labor-supply elasticities suggest $\gamma < 2$. Thus, he suggests that the expected-utility framework has to be abandoned to accommodate the higher values of γ needed to accord with observed behavior in asset and insurance markets.

The last finding fits with my results in the EZW framework. The (consumption-compensated) wage elasticity can be computed from equation (31) by taking a derivative with respect to w_t, while holding fixed C_t, to get

$$\text{compensated wage elasticity of labor supply} = \lambda/\alpha. \tag{34}$$

Given the production-function parameter, α, the compensated wage elasticity can be anything, depending on λ, the elasticity of substitution between consumption and leisure. Thus, in the EZW framework, labor-supply elasticities place no restriction on the permissible range for γ. The EZW model, extended to incorporate labor-leisure choice, has three independent parameters: one governing risk aversion (γ), another for the IES (θ), and a third for consumption-leisure substitution (λ).

V. Endogenous Saving and Investment

In an endowment economy, agents do not react to changes in uncertainty by altering saving and investment. Generally, the potential for such adjustments affects welfare costs—not at the margin (by the envelope theorem) but for large changes in parameters. This section illustrates this process by using a version of the tractable *AK* model of endogenous saving and investment developed in Barro (2006, sect. VIII).

The quantity of trees is now variable and corresponds to the capital stock, K_t. Production of fruit is given by an *AK* production function:

$$Y_t = AK_t. \tag{35}$$

Unlike the original model, the productivity level, $A > 0$, is now constant. Output can be consumed as fruit, C_t, or invested as seed, I_t, so that

$$(36) \qquad C_t = Y_t - I_t = AK_t - I_t.$$

The creation of new trees through planting seeds (that is, investment) is assumed to be rapid enough so that, as in the conventional one-sector production framework, the price of trees (capital) in units of fruit is pegged at a price normalized to one. This setting corresponds to "Tobin's q" always equaling one—unlike in the previous model, where the market price of trees was variable.

The capital stock evolves because of gross investment and depreciation, $\delta_{t+1}K_t$:

$$(37) \qquad K_{t+1} = K_t + I_t - \delta_{t+1}K_t.$$

The depreciation rate is stochastic and equal to

$$(38) \qquad \delta_{t+1} = \delta + u_{t+1} + v_{t+1},$$

where $0 < \delta < 1$. The u_{t+1} shock, normally distributed with mean 0 and variance σ^2, represents normal fluctuations, as in the previous setting. The v_{t+1} shock represents rare disasters, again as in the earlier model. With probability $1 - p$, $v_{t+1} = 0$, and with probability p, $v_{t+1} = -b$, that is, the fraction b $(0 < b < 1)$ of the trees is destroyed. The analysis requires $0 < \delta_{t+1} < 1$. This restriction holds with probability one as the length of the period approaches zero, assuming that σ^2 and p are proportional to the length of the period.

Since the market price of trees is pegged at one, the expected rate of return on equity shares is given immediately by

$$(39) \qquad r^e = A - \delta - p\,Eb.$$

Because the shocks, u_{t+1} and v_{t+1}, are i.i.d. (permanent to the levels of capital stock and GDP), the ratio of gross investment (and gross saving) to the capital stock will be optimally chosen as a constant, denoted by ν. One way to determine ν is to use the usual consumption-based asset-pricing formula for equity shares, combined with the condition that the price of these shares equals unity.

The saving ratio, ν, can be determined, as the length of the period approaches zero, to be

$$(40) \qquad \nu = \delta + (1/\theta)\left\{A - \delta - \rho + (½)\,\gamma\,(\theta - 1)\,\sigma^2 + \left(\frac{\theta - 1}{\gamma - 1}\right) p\,[E(1 - b)^{1-\gamma} - 1]\right\},$$

assuming $\gamma \neq 1$. One point from equation (40) is that, if $\gamma > 0$, the sign of the effect of uncertainty $(\sigma, p$, or the distribution of $b)$ on the saving ratio, ν, depends on the IES, $1/\theta$, not the degree of risk aversion, γ. Moreover, if $\theta < 1$—the case that we emphasized previously—so that the IES exceeds one, the "substitution effect" dominates, and more uncertainty (higher σ or p or a shift of the b-distribution toward larger values) leads to a lower saving ratio, ν.

The expected growth rate of GDP and the capital stock is

$$(41) \qquad E_t(K_{t+1}/K_t - 1) = \nu - \delta - p\,Eb.$$

Therefore, a higher saving ratio, ν, in equation (40) implies a higher expected growth rate in equation (41).

The calibration of the endogenous-saving model can be matched to the endowment economy. Thus, I use the baseline parameter values assumed before, including $\gamma = 4$ and $\theta = 0.5$. To get a full correspondence, the expected growth rate, given in equation (41), has to equal the value $g^* = 0.020$ used before (see equation (2)). Equations (40) and (41) imply that the expected growth rate determines the parameter combination A-δ-ρ, which turns out to equal 0.024. The formula for the risk-free rate in the endogenous-saving model[17] pins down the value of ρ needed to generate $r^f = 0.010$, as before. The result is $\rho = 0.052$ (as in the baseline case for the endowment economy) and, hence, $A - \delta = 0.074$. The expected rate of return on equity, r^e, given by equation (39), then equals 0.069, the same as in the endowment economy. Therefore, the equity premium, $r^e - r^f = 0.059$, is also the same as before.

Substitution of the various parameter values into equation (40) yields a gross saving ratio, ν, of $0.025 + \delta$. (The parameters A and δ cannot be separated, but this limitation does not affect the welfare analysis.) As an example, if $\delta = 0.05$, then $\nu = 0.075$. That is, annual gross saving and investment equal 7.5 percent of the capital stock.

The new results on welfare costs apply to large changes, for example, setting $\sigma = 0$ or $p = 0$. We can, as before, express the results in terms of proportionate declines in initial GDP (and, in the present context, also the capital stock) that would be willingly exchanged for each kind of reduction in uncertainty. *These welfare effects in the endogenous-saving model coincide with those for the endowment economy if the gross saving ratio, ν, is constrained to remain fixed at its initial value* (*0.075*). Specifically, the offsetting proportionate reductions in GDP would be 1.65 percent for setting $\sigma = 0$ and 32.5 percent for setting $p = 0$. (This effect for setting $p = 0$ coincides with that for the endowment economy if the latter calculation includes the growth effect from reducing p in the formula for g^* in equation (2)—see equation (28).) In effect, with ν held fixed, the endogenous-saving model operates like an endowment economy.

The results are different if the saving ratio, ν, is free to adjust to the changes in σ and p, in accordance with the optimal response given by equation (40). Since the optimal saving response cannot make the situation worse, the compensating output variations for eliminating uncertainty must be at least as large as those in the endowment economy.[18] For the calibration parameters already mentioned, the results are[19]

- Setting $\sigma = 0$: saving ratio, ν, rises from 0.0751 to 0.0759, welfare effect = 1.65 percent;
- Setting $p = 0$: ν rises from 0.0751 to 0.0924, welfare effect = 35.7 percent.

For the case where σ is set to zero, the impact of allowing for the small increase in the saving ratio (from 0.0751 to 0.0759) is trivial. Hence, the welfare effect, 1.65 percent, is essentially the same as that for the endowment economy.[20] However, when considering $p = 0$, the significant rise in the saving ratio (from 0.0751 to 0.0924) generates a detectable increase in the welfare effect: the output that would be relinquished to eliminate disaster risk rises from 32.5 percent to 35.7 percent. Note that we could instead start from $\sigma = 0$ or $p = 0$ and compute the proportionate

[17] The formula is $r^f = A - \delta - \gamma\sigma^2 - p[E(1-b)^{-\gamma} - E(1-b)^{1-\gamma}]$.

[18] If $\theta = 1$, ν does not depend on σ or p, and the results are the same as those as in the case where ν is constrained not to vary.

[19] The welfare calculations can be made from a formula that modifies equation (23): $(Y_t)^*/Y_t = (V/V^*)^{1/(1-\theta)}(A - \nu)/(A - \nu^*)$. The variables V and V^* correspond to the price-dividend ratios determined in equation (13), except that g^* is replaced by the expected growth rate given in equation (41). The variables ν and ν^* are investment ratios in the initial and hypothetical situations, as determined by equation (40).

[20] Anne Epaulard and Aude Pommeret (2003) also modified the calculations of welfare costs of uncertainty to allow for adjustments of saving in an AK model. However, because they excluded disaster risks, the quantitative significance of the saving adjustment was minor—as in the present case for setting $\sigma = 0$.

increase in GDP required to compensate for an increase in σ or p. In this case, the optimal adjustment of the saving ratio (downward) *reduces* the welfare effect in the sense of the compensating, proportionate increase in GDP.

VI. Concluding Observations

The baseline parameter value $\sigma = 0.02$ per year represents the extent of business fluctuations during the tranquil post–World War II years in the United States and other OECD countries. This period was calm for the OECD countries (at least until 2008!) when considered in comparison to the first half of the twentieth century, a turbulent time that featured World Wars I and II and the Great Depression. Hence, a reduction in σ amounts to making milder the business fluctuations that were already strikingly tame. Not surprisingly, the benefit from this change—corresponding to around 1.5 percent of GDP each year—is only moderate, though still important.

In contrast, the disaster probability, p, and size, b, refer to major economic crises, such as those that occurred in many countries during World Wars I and II and the Great Depression. Outside of the OECD, we can also think of p and b as relating to events such as the Asian financial crisis of the late 1990s, the Latin-American debt crisis of the early 1980s, and the Argentine exchange-rate crisis of 2001–2002. A reduction in p amounts to lowering the chance of repeating these kinds of extreme events, and a fall in b amounts to decreasing the likely size of these events. To go further, decreases in p or b constitute reductions in the probability or size of disasters not yet seen or, at least, not seen in the twentieth century. Included here would be nuclear conflicts, large-scale natural disasters (tsunamis, hurricanes, earthquakes, asteroid collisions), and epidemics of disease (Black Death, avian flu). My estimates indicate that the welfare consequences from eliminating all uncertainty of this kind are large—roughly 15 times the effects for normal economic fluctuations. Moreover, these large welfare effects from disasters arise even though the analysis considers only the utility losses from reduced consumption. A broader analysis would include direct utility effects from changes in mortality and health.

Macroeconomic stabilization policies, including monetary policy, relate to both types of uncertainty—σ on the one hand and p or b on the other hand. The policies may also affect the long-term expected growth rate, g^*. Well known is the success of OECD countries in achieving low and stable inflation since the mid-1980s. This success is sometimes argued to have contributed to milder business fluctuations (lower σ) and perhaps to stronger average economic growth (higher g^*). However, commentaries on monetary policy frequently also stress the roles of central banks in exacerbating or moderating major economic crises. For example, Milton Friedman and Anna J. Schwartz (1963) blame the Federal Reserve for the severity of the Great Depression in the United States, as well as for the sharp recession of 1937–1938. Observers of Alan Greenspan's tenure as Fed chair often focus on his role in apparently moderating the consequences of the global stock-market crash of 1987 and the Long-Term Capital Management/Russian crisis of 1998. Difficulties in 2007–2008 related to mortgage and other financing may be another such case—one that makes or breaks the reputations of Ben Bernanke and other central bankers. The policy actions during these kinds of crises—if actually effective—have more to do with lowering p and b than decreasing σ. A key, unresolved issue is whether and how a monetary authority can reduce the probability, p, and size, b, of economic collapses.

Another possibility is that moderating ordinary business fluctuations, represented by σ, has an indirect payoff in terms of reducing the probability, p, and size, b, of major crises. That is, preventing or lessening mild recessions may lower the chances of experiencing downturns that magnify into depressions. This perspective would greatly amplify the rewards from ordinary stabilization policies, such as those practiced regularly by central bankers. However, at this point, this idea is just a conjecture, well worth further investigation.

Governmental institutions and policies that are not directly related to macroeconomic stabilization can also affect disaster probabilities and sizes. For example, the formation of the European Union and the adoption of the euro have often been analyzed as influences on the extent of business fluctuations (σ) and the average rate of economic growth (g^*), sometimes focusing on the role of international trade in goods and assets. From a political perspective, however, the main force behind the adoption of these institutions was likely the desire to avoid a repetition of World War II, that is, to reduce the disaster probability, p, applicable to war. This perceived impact on disaster probability related to war is likely to be a key element in explaining why these institutions exist in Western Europe. Of course, this perception may be inaccurate—forcing Germany and France to share monetary, fiscal, and other policies may ultimately create more conflict than it eliminates. Thus, an important research topic is the actual influence of various policies and institutions on the probability and size of disasters, including wars.

REFERENCES

Alvarez, Fernando, and Urban J. Jermann. 2004. "Using Asset Prices to Measure the Cost of Business Cycles." *Journal of Political Economy*, 112(6): 1223–56.

Atkeson, Andrew, and Christopher Phelan. 1994. "Reconsidering the Costs of Business Cycles with Incomplete Markets." In *NBER Macroeconomics Annual 1994*, ed. S. Fischer and J. J. Rotemberg, 187–207. Cambridge: MIT Press.

Balke, Nathan S., and Robert J. Gordon. 1989. "The Estimation of Prewar Gross National Product: Methodology and New Evidence." *Journal of Political Economy*, 97(1): 38–92.

Banerjee, Anindya, Robin L. Lumsdaine, and James H. Stock. 1992. "Recursive and Sequential Tests of the Unit-Root and Trend-Break Hypotheses: Theory and International Evidence." *Journal of Business and Economic Statistics*, 10(3): 271–87.

Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, 59(4): 1481–1509.

Barlevy, Gadi. 2004. "The Cost of Business Cycles under Endogenous Growth." *American Economic Review*, 94(4): 964–90.

Barro, Robert J. 2005. "Rare Events and the Equity Premium." National Bureau of Economic Research Working Paper 11310.

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

Barro, Robert J., and Jose F. Ursua. 2008a. "Consumption Disasters since 1870." *Brookings Papers on Economic Activity*, (1): 225–350.

Barro, Robert J., and Jose F. Ursua. 2008b. "Consumption Disasters in the Twentieth Century." *American Economic Review*, 98(2): 58–63.

Chatterjee, Satyajit, and Dean Corbae. 2007. "On the Aggregate Welfare Cost of Great Depression Unemployment." *Journal of Monetary Economics*, 54(6): 1529–44.

Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *American Economic Review*, 96(5): 1821–34.

Cochrane, John H. 1988. "How Big Is the Random Walk in GNP?" *Journal of Political Economy*, 96(5): 893–920.

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

Epaulard, Anne, and Aude Pommeret. 2003. "Recursive Utility, Endogenous Growth, and the Welfare Cost of Volatility." *Review of Economic Dynamics*, 6(3): 672–84.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4): 937–69.

Friedman, Milton, and Anna J. Schwartz. 1963. *A Monetary History of the United States, 1867–1960.* Princeton: Princeton University Press.

Gabaix, Xavier. 2008. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." National Bureau of Economic Research Working Paper 13724.

Giovannini, Alberto, and Philippe Weil. 1989. "Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model." National Bureau of Economic Research Working Paper 2824.

Gruber, Jonathan. 2006. "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution." National Bureau of Economic Research Working Paper 11945.

Hall, Robert E. 1988. "Intertemporal Substitution in Consumption." *Journal of Political Economy*, 96(2): 339–57.

Hess, Gregory D. 2003. "The Economic Welfare Cost of Conflict: An Empirical Assessment." Unpublished.

Hicks, John R. 1946. *Value and Capital*. 2nd ed. Oxford: Oxford University Press.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. "Production, Growth and Business Cycles: I. The Basic Neoclassical Model." *Journal of Monetary Economics*, 21(2/3): 195–232.

Kocherlakota, Narayana R. 1990. "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result." *Journal of Finance*, 45(1): 175–90.

Lucas, Robert E., Jr. 1978. "Asset Prices in an Exchange Economy." *Econometrica*, 46(6): 1429–45.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.

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

Maddison, Angus. 2003. *The World Economy: Historical Statistics*. Paris: OECD.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, 15(2): 145–61.

Obstfeld, Maurice. 1994. "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability." *European Economic Review*, 38(7): 1471–86.

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

Salyer, Kevin D. 2007. "Macroeconomic Priorities and Crash States." *Economics Letters*, 94(1): 64–70.

Taylor, Bryan. 2005. "GFD Guide to Total Returns on Stocks, Bonds and Bills." http://www.globalfindata.com.

Weil, Philippe. 1990. "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics*, 105(1): 29–42.