On the Welfare Costs of Consumption Uncertainty*

Robert J. Barro

Harvard University

September 2006

Abstract

Satisfactory calculations of the welfare cost of aggregate consumption uncertainty require a framework that replicates major features of asset prices and returns, such as the high equity premium and low risk-free rate. A Lucas-tree model with rare but large disasters is such a framework. In a baseline simulation, the welfare cost of disaster risk is large—society would be willing to lower real GDP by about 20% each year to eliminate all disaster risk, including wars. In contrast, the welfare cost from usual economic fluctuations is much smaller, though still important—corresponding to lowering GDP by around 2% each year.

*This research is supported by the National Science Foundation. I am grateful to Jong-Wha Lee for motivating me to think about this topic. I appreciate comments from John Cochrane, Xavier Gabaix, and Ivan Werning.

Lucas (1987, Ch. 3; 2003, section II) argued that the welfare gain from eliminating uncertainty in aggregate consumption is trivial. He got this answer by using parameters for the time series of real per capita consumer expenditure from U.S. post-World War II macroeconomic data, along with plausible values for the coefficient of relative risk aversion.

One problem with this calculation, apparent from Mehra and Prescott (1985), is that simulations with the same model and parameters do not get into the right ballpark for explaining well-known asset-pricing puzzles, such as the high equity premium and low risk-free rate. These failures with respect to asset returns suggest, as observed by Atkeson and Phelan (1994), that the model misses important aspects of consumption uncertainty. Hence, the model's estimates of welfare effects from consumption uncertainty are likely to be inaccurate. A possibly satisfactory framework, used here, is one with rare economic disasters, as in Rietz (1988). Barro (2006) shows that this model can replicate some prominent features of asset returns. In this model, changes in consumption uncertainty that reflect shifts in the probability or size of disasters have major implications for welfare.

I. A Lucas Fruit-Tree Model

The 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. 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:

(1) $$\log(Y_{t+1}) = \log(Y_t) + \gamma + u_{t+1} + v_{t+1},$$

where $\gamma \geq 0$ is a constant that reflects exogenous, long-term productivity growth. 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 random term v_{t+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. If a disaster occurs, output contracts proportionately by the fraction b. The idea is that the probability of disaster in a period is small but b is large. The distribution of v_{t+1} is given by

probability e^{-p}: $v_{t+1} = 0$,

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

Unlike Lucas (1987, Ch. 3), but in line with Obstfeld (1994), the shocks u_{t+1} and v_{t+1} in Eq. (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.

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 U.S. 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 1-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 Eq. (1). However, 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 Eq. (1).

Cogley's results hold up for a broader sample comprising 19 OECD countries. The data on per capita GDP are for 1870-2004 from Maddison (2003), updated from Economist Intelligence Unit, *Country Data* (and using U.S. data from Balke and 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 Eq. (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 with models that predict reversion to a fixed, deterministic trend. In contrast, the results on variance ratios seem consistent with the simple stochastic-trend specification in Eq. (1). Therefore, I use this specification to represent the evolution of GDP and consumption. Richer models—such as those with trend-breaks analyzed in Banerjee, Lumsdaine, and Stock (1992)—may fit better than Eq. (1). However, the results for the welfare costs of consumption uncertainty may be similar because the trend-

[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 U.S. ratios at these values of k are, respectively, 1.30, 1.34, and 0.94.

break and stochastic-trend models share the property of great uncertainty about GDP and consumption in the distant future. In contrast, the important, but apparently counter-factual, characteristic of models with reversion to a fixed trend is the comparatively low uncertainty about the distant future.

Previous research (Barro [2006, Table 1 and Figure 1]) gauged the probability and size of disaster events from time series on per capita GDP for 35 countries for the full 20th century from Maddison (2003).[2] For contractions of 15% or more over consecutive years (such as 1939-44 for France and 1929-33 for the United States), 60 events were found.[3] 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% per year.

The contraction proportion b for the observed 20th century disasters ranged from 15% to 64%, with a mean of 29%.[4] However, with substantial risk aversion (for example, with a coefficient of relative risk aversion of 3 or 4), the effective average value of b is much higher than the mean. It turns out that a constant b of around 40% generates about the same equity premium and welfare effects as the empirically observed frequency

[2] In the fruit-tree model, GDP and consumption coincide. More generally, consumption would be more appropriate than GDP for an analysis of welfare costs. However, long-term data on real consumer expenditure are not reported by Maddison (2003) and are not readily available for many countries. In ongoing research, Jose Ursua and I are assembling a data set on long-term real consumer expenditure for as many countries as possible.

[3] This analysis excludes five post-war GDP contractions that did not involve large declines in real consumer expenditure. The lower limit of 15% is arbitrary. Extending to 10% brings in another 21 contractions for the 35 countries. However, the inclusion of these smaller contractions has a minor effect on the results.

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

distribution of b. My previous calibration exercise allowed b to be generated from the empirical distribution. However, for present purposes, it is satisfactory to simplify by treating b as a constant, 40%.[5]

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 2004, 14 exhibit negative skewness, and the only substantially positive values are for France, the Netherlands, and Switzerland.

The expected growth rate of the economy depends not only on the growth-rate parameter, γ, but also on σ, p, and b. As the length of the period approaches zero, the expected growth rate of GDP and consumption, defined to be γ*, is given by

$$\gamma^* = \gamma + (1/2)\cdot\sigma^2 - pb. \quad (2)$$

The representative consumer maximizes a familiar time-additive utility function with iso-elastic preferences:

$$U_t = E_t \sum_{i=0}^{\infty} e^{-\rho i} \cdot [(C_{t+i})^{1-\theta} - 1]/(1-\theta), \quad (3)$$

where $\rho \geq 0$ is the rate of time preference, and $\theta > 0$ is the coefficient of relative risk aversion and the reciprocal of the intertemporal elasticity of substitution for consumption. The simplicity of the underlying structure (i.i.d. shocks, representative consumer with time-additive and iso-elastic preferences, closed economy with no investment) allows for a closed-form solution for expected utility as a function of the underlying parameters of

[5] I neglect two other features of my earlier analysis that matter for the equity premium but not for the welfare calculations—leverage in the ownership structure for trees and default possibilities on bonds.

preferences and of the stochastic process for output. Obstfeld (1994) derived analogous closed forms in a model without disaster risk.

A key variable turns out to be the market value, V, of a tree that initially produces one unit of fruit. This value can be calculated by summing the prices of equity claims on future "dividends," $C_{t+i} = Y_{t+i}$. (In order to correspond to the summation in Eq. [3], it is convenient to treat C_t, rather than C_{t+1}, as the first payout on tree equity.) These prices follow from the usual first-order conditions for maximizing U_t. As the arbitrary period length approaches zero, the reciprocal of V turns out to be[6]

$$(4) \qquad 1/V = \rho + (\theta-1)\cdot\gamma - (1/2)\cdot(\theta-1)^2\sigma^2 - p\cdot[(1-b)^{1-\theta} - 1].$$

We can think of V as the price-earnings ratio for an unlevered equity claim on a tree.

The right-hand side of Eq. (4) equals the difference between the expected rate of return on unlevered equity, given by[7]

$$(5) \qquad r^e = \rho + \theta\gamma - (1/2)\cdot\theta^2\sigma^2 + \theta\sigma^2 - p\cdot[(1-b)^{1-\theta}-1+b],$$

and the expected growth rate, γ^*, given in Eq. (2). The transversality condition, which guarantees that the market value of a tree is positive and finite, is that the right-hand side of Eq. (4) be positive—that is, $r^e > \gamma^*$.

Expected utility as of period t, given in Eq. (3), can be determined to be[8]

$$(6) \qquad U_t = V\cdot(Y_t)^{1-\theta}/(1-\theta),$$

[6] See Barro (2006, Eq. [17]).
[7] See Barro (2006, Eq. [9]).
[8] For $\theta=1$, expected utility is $U_t = (1/\rho)\cdot\log(Y_t)$. The derivation of Eq. (6) depends on the condition $C_t=Y_t$ and on time-additive, iso-elastic preferences but not on the particular stochastic process for output in Eq. (1). However, the constancy of the price-earnings 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-earnings ratios for stock-market claims. However, the model can match this volatility if the disaster probability, p_t, moves around. Gabaix (2006) shows that the main implications of the model for asset pricing go through if p_t evolves exogenously in random-walk-like fashion.

up to an inconsequential additive constant. Equation (4) determines V. However, for subsequent purposes, we want a formula for V in terms of the expected growth rate, γ^*, rather than the growth-rate parameter, γ. The result follows from Eq. (2):

$$(7) \qquad 1/V = \rho + (\theta-1)\cdot\gamma^* - (1/2)\cdot\theta\cdot(\theta-1)\cdot\sigma^2 - p\cdot[(1-b)^{1-\theta} - 1 - b\cdot(\theta-1)].$$

II. Baseline Calculation of Welfare Effects

Equations (6) and (7) determine the effects on expected utility from changes in the expected growth rate, γ^*, and the parameters that govern consumption risk: σ, p, and b. These effects can be compared with that from proportionate shifts in the initial level of GDP and consumption, Y_t.

The marginal effect on utility from a proportionate change in Y_t is given from Eq. (6) by

$$(8) \qquad \frac{\partial U_t}{\partial Y_t}\cdot Y_t = (Y_t)^{1-\theta}\cdot V\,.$$

The marginal effect from a change in γ^* follows from Eqs. (6) and (7) as

$$(9) \qquad \frac{\partial U_t}{\partial \gamma^*} = (Y_t)^{1-\theta}\cdot V^2\,.$$

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

$$(10) \qquad \frac{-\partial U_t/\partial \gamma^*}{(\partial U_t/\partial Y_t)\cdot Y_t} = -V\,.$$

This result gives the proportionate decrease in Y_t that compensates, at the margin, for an increase in γ^*—in the sense of preserving expected utility. Equation (10) shows that this compensating output change depends only on the combination of parameters that enter into the price-earnings ratio, V, as determined in Eq. (4) or Eq. (7).

To pin down a reasonable magnitude for V, start with the already mentioned specification p=0.017 per year and b=0.4. The other parameters are assumed to equal the values used (and defended) in the main calibration exercise carried out in Barro (2006, Table 5, col. 2). I assume a rate of time preference of $\rho=0.03$ per year, a coefficient of relative risk aversion of $\theta=4$, a growth parameter of $\gamma=0.025$ per year, and a standard deviation for the u_t shocks of $\sigma=0.02$ per year.[9] The assumed parameter values yield V=24.1 in Eq. (3).

A key parameter in the determination of V and, therefore, for calculations of welfare costs is the coefficient of relative risk aversion, taken here to be $\theta=4$. Barro (2006, Table 5) shows that this coefficient (along with the other assumed parameters and with additional assumptions about leverage and default probability on bonds) generates empirically reasonable equity premia. As discussed later, much lower (or higher) values of θ—for example, $\theta<3$ with the rest of the specification unchanged—generate empirically implausible asset returns. Therefore, from the standpoint of assessing the welfare implications of consumption uncertainty, it would be unsatisfactory to use the present framework along with values of $\theta<3$. Since the model would then generate inaccurate patterns of asset returns, there would be no reason to trust the model with respect to its implications for welfare costs. However, since the model with $\theta\approx4$ generates realistic asset returns, the specification is, on this ground, a plausible candidate for assessing welfare costs.

[9] For 21 OECD countries over 1954-2004—a tranquil period with no disaster events for these countries—the median of the growth rates of real per capita consumer expenditure was 0.026 per year. The median standard deviation of the growth rates was 0.024. The U.S. values were 0.024 and 0.018, respectively.

With V=24.1, Eq. (10) implies that a small rise in the expected growth rate, γ*—for example, by 0.1% per year—has to be compensated by a fall in the initial level of GDP, Y_t, by 2.4%. Despite differences in specification, this result accords with the one found by Lucas (1987, Ch. 3, p. 24). 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 expected utility, U_t, from a change in σ is given from Eqs. (6) and (7) by

$$\text{(11)} \qquad \frac{\partial U_t}{\partial \sigma} = -(Y_t)^{1-\theta} \cdot \theta \sigma V^2 .$$

Note that this calculation holds fixed γ*, as given in Eq. (2), not γ. Hence, the loss in utility shown in Eq. (11) reflects only consumption uncertainty, not an effect on the expected growth rate.

The utility rate of transformation between proportionate changes in Y_t and σ is given by

$$\text{(12)} \qquad \frac{-(\partial U_t / \partial \sigma) \cdot \sigma}{(\partial U_t / \partial Y_t) \cdot Y_t} = \theta \sigma^2 V .$$

This expression gives the proportionate increase in initial GDP required to compensate, at the margin, for a proportionate rise in σ. The parameters specified before imply $\theta\sigma^2 V = 0.039$. Therefore, to maintain expected utility, a small increase in σ—for example, by 10% (from 0.020 to 0.022)—requires a rise in the initial level of GDP by 0.39%. Since we are holding fixed the expected growth rate, γ*, this proportionate rise in

GDP level should be viewed as applying each year.[10] We could also modify the calculations to allow for a growth effect from a change in σ; that is, for given γ, γ* rises with σ in Eq. (2).

Consider now the welfare consequences from a change in the disaster probability, p, for given disaster size, b. Equations (6) and (7) imply

$$(13) \qquad \frac{\partial U_t}{\partial p} = -(Y_t)^{1-\theta} \cdot V^2 \cdot [(1-b)^{1-\theta} - 1 - b \cdot (\theta - 1)]/(\theta - 1) .$$

Again, this formula applies while holding fixed γ* in Eq. (2). The utility rate of transformation between proportionate changes in Y_t and p is given by

$$(14) \qquad \frac{-(\partial U_t / \partial p) \cdot p}{(\partial U_t / \partial Y_t) \cdot Y_t} = pV \cdot [(1-b)^{1-\theta} - 1 - b \cdot (\theta - 1)]/(\theta - 1) .$$

With the parameter values used before, the right-hand side of Eq. (14) equals 0.33. Hence, an increase in p by 10% (from 0.0170 to 0.0187) matches up with a proportionate rise in initial GDP by 3.3%. Again, this change in GDP level applies each year. We could also modify the calculations to allow for a growth effect from a change in p; that is, for given γ, γ* falls with p in Eq. (2).[11]

We can similarly assess the welfare consequences from changes in disaster size, b, while holding fixed disaster probability, p. The formula, analogous to Eq. (14), is

$$(15) \qquad \frac{-(\partial U_t / \partial b) \cdot b}{(\partial U_t / \partial Y_t) \cdot Y_t} = bpV \cdot [(1-b)^{-\theta} - 1] .$$

---

[10] 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 purely transitory (one-year) disturbance to the level of output.

[11] Barro (2006, section 8) allows for endogenous investment in an "AK" version of the fruit-tree model. In this model, the saving rate (and, hence, ratio of investment to GDP) rises with p if θ>1. The full effect of p on the expected growth rate of GDP is then ambiguous, because the saving effect is offset by the direct negative impact of p on the expected growth rate. The same considerations apply to the growth effects from changes in σ. See Barlevy (2004) for a discussion of models in which uncertainty affects the long-run growth rate.

With the usual parameter values, the right-hand side equals 1.10. Hence, an increase in b by 10% (from 0.40 to 0.44) matches up with a proportionate rise in initial GDP by 11%. Once again, we could modify the calculations to allow for a growth effect from a change in b; that is, for given γ, γ^* falls with b in Eq. (2).[12]

The formulas in Eqs. (10), (12), (14), and (15) apply locally; that is, to small changes in Y_t, γ^*, σ, p, and b. We can instead use Eqs. (6) and (7) to assess the effects on expected utility from large changes. Let V and Y_t be the values that apply for the baseline specification of parameters. Let V* and $(Y_t)^*$ be values that apply in some alternative situation that delivers the same expected utility, U_t. Then the formula for U_t in Eq. (6) implies

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

As an aside, Alvarez and Jermann (2004) try to go as far as possible to gauge the welfare costs of consumption uncertainty just by observing or estimating various asset prices. The general idea is to obtain welfare-cost estimates that are not sensitive to specifications of preferences and production structure. Equation (16) 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-earnings ratio for unlevered equity claims on trees. However, the price V* is unlikely to be observable. V* is the price-earnings ratio for tree equity in a hypothetical economy—for example, one in which all uncertainty has been reduced to zero. Moreover, Eq. (16) shows that the welfare cost, measured by the

[12] In the model described in n.11, an increase in b raises the saving rate if $\theta>1$. The full effect of b on the expected growth rate of GDP is then ambiguous.

compensating output change $(Y_t)^*/Y_t$, depends on the coefficient of relative risk aversion, θ, for given V and V*.

Lucas focused on the consequences of eliminating all consumption uncertainty associated with usual business fluctuations—in the present context, this exercise corresponds to setting σ=0. The formula for V in Eq. (7) implies for this case

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

Substituting into Eq. (16) yields[13]

$$(17) \qquad (Y_t)^*/ Y_t = [1 + (1/2)\cdot\theta\cdot(\theta-1)\cdot\sigma^2 V]^{1/(1-\theta)}.$$

With the parameter values assumed before, $(Y_t)^*$ is 1.9% below Y_t. That is, society would be willing to give up about 2% of output each year to eliminate all of the customary economic fluctuations represented by σ. As noted before (n. 10), 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, based on Eqs. (7) and (16), is[14]

$$(18) \qquad (Y_t)^*/Y_t = \{1 + pV\cdot[(1-b)^{1-\theta} - 1 - b\cdot(\theta-1)]\}^{1/(1-\theta)}.$$

With the same parameter values as before, $(Y_t)^*$ is 20.6% below Y_t. Hence, when gauged by the compensating proportionate change in output, eliminating disaster risk is worth about 10 times as much as eliminating normal economic fluctuations.

[13] If the magnitude of $(1/2)\cdot\theta\cdot(\theta-1)\cdot\sigma^2V$ is much less than one, the result simplifies to $\log[(Y_t)^*/Y_t] \approx -(1/2)\cdot\theta\sigma^2V$; that is, one-half the size of the local effect shown in Eq. (12). This approximation is satisfactory at the baseline specification, where $(1/2)\cdot\theta\cdot(\theta-1)\cdot\sigma^2V = 0.058$.

[14] If θ is close to 1, the result simplifies to $\log[(Y_t)^*/Y_t] \approx -(1/2)\cdot pb^2V$. However, this approximation is unsatisfactory at the baseline specification, where θ=4.

We can also consider the elimination of all consumption uncertainty by setting σ=0 and p=0 (or b=0) simultaneously. In this case, $(Y_t)^*$ is 21.3% below Y_t. Not surprisingly, the main effect comes from setting p=0.

III. Sensitivity of the Welfare-Cost Estimates

The welfare-cost estimates, including the effects from eliminating all disaster risk, depend particularly on the coefficient of relative risk aversion, θ=4, and the probability and size of disasters, p=0.017 per year and b=0.4. The welfare-cost estimates decline if θ, p, and b are smaller. However, these changes also affect the model's predictions for asset returns, including the equity premium and risk-free rate. Specifications that deviate sharply from observed asset returns probably should not be taken seriously with regard to their implications for welfare costs.

To illustrate, in the baseline specification with θ=4, the elimination of all disaster risk (setting p or b to zero) balances, in the sense of maintaining expected utility, with a proportionate cut in the initial output level by 20.6%. If we reduce θ but keep the rest of the specification unchanged, the compensating proportionate decline in initial GDP turns out to be 14.6% when θ=3, 9.4% when θ=2, and 6.1% when θ=1. Thus, the welfare gain from eliminating disaster risk remains substantial for θ as low as 1, but the quantitative effect in terms of the compensating change in output is reduced by a factor of more than 3 compared to the situation when θ=4.

The problem with these results is that substantial reductions in the risk-aversion coefficient, θ, for given values of the other parameters, cause the model to fail in its predictions on asset returns. For example, at θ=4, the model predicts an equity premium (excess of the expected rate of return on levered equity over that on short-term bonds) of 4.4% per year.[15] This result compares with empirically observed equity premia for OECD countries over long periods of around 6% (see Table 1). Thus, the prediction is at least in the right ballpark. However, the predicted equity premium falls to 2.4% at θ=3, 1.2% at θ=2, and 0.5% at θ=1. Therefore, if θ is well below 4, some other change in specification is necessary for the model to match observed asset returns.

As an example, it is possible to restore reasonable predictions for the equity premium at low θ if the disaster probability, p, is raised above 1.7% per year. At θ=3, p has to be 3.2% per year to generate the same levered equity premium, 4.4% per year, as originally. But then the elimination of all disaster risk (setting p or b to zero) turns out to balance against a proportionate decline in the initial output level by 27.7%, even more than the 20.6% calculated originally. At θ=2, p has to be 6.6% per year, and the elimination of disaster risk corresponds to a proportionate decline in the initial output level by 36.7%. The point is not that these unrealistically high disaster probabilities—p rising above 6% per year—should be taken seriously. Rather, the inference is that disciplining the model to be consistent with observed asset returns, notably the high equity premium, tends to generate large estimates for the welfare benefits from eliminating all disaster risk.

[15] These results on the levered equity premium, calculated from Barro (2006, Eq.[23]), assume a debt-equity ratio of 0.5. Short-term bonds, such as government bills, are assumed to have a chance for partial default during major crises. In the calculations, the conditional default probability is 0.4, and the extent of default (when it occurs) equals b, which is 0.4 in the present context.

IV. Alvarez-Jermann and Risk-Free Interest Rates

As mentioned before, Alvarez and Jermann (2004), henceforth AJ, gauge the welfare cost of aggregate consumption uncertainty from an approach that relies on asset prices and, thereby, avoids much dependence on forms of consumer preferences or the production structure. In a general way, my analysis follows their idea that welfare-cost estimates should be disciplined to be consistent with observed patterns in asset prices and returns. An earlier expression of this idea is in Atkeson and Phelan (1994).

AJ observe (p. 1225) that the "marginal cost of consumption fluctuations [equals] the ratio of the values of two securities: a claim to the consumption trend … and a claim to aggregate consumption …" In the Lucas-tree economy, the claim to aggregate consumption corresponds to an equity share on a tree and has the price V determined in Eq. (4). This price is positive and finite if the expected rate of return on equity, r^e in Eq. (5), exceeds the expected growth rate of GDP and consumption, γ^* in Eq. (2). This inequality is the transversality condition for the model.

A problem arises with the price of a claim to the consumption trend—the payments on this claim would start at the initial level of per capita GDP and consumption and then grow at the constant rate γ^*. The rate of return on this security would be the risk-free rate, denoted r^f. The price of such a claim is finite only if $r^f > \gamma^*$. However, in the model described in Barro (2006), there is no reason for this inequality to hold—in fact, the risk-free rate can be negative. For the baseline calibration of the present model, the risk-free rate turns out to be positive—1.3% per year—but less than the expected growth rate of per capita GDP and consumption—1.8%.[16]

Table 1 shows growth rates and real rates of return for 1880-2004 for 11 OECD countries with available data.[17] The table reports four concepts of growth rates: for per capita and levels of real consumer expenditure and GDP. In the Lucas-tree model with constant population, the four concepts coincide. When consumption and GDP diverge, consumption—proxied by real consumer expenditure in the data—would be the relevant variable for asset pricing. However, this distinction is empirically unimportant in the present context, because the long-term growth rates of real consumer expenditure and GDP are similar.

When population is growing, assets would still be priced in relation to variations over time in per capita consumption. However, if new people enter the economy as children tied altruistically to parents, the AJ claim to the consumption trend would relate to the trend of the level of consumption, not consumption per capita. (A family dynasty would want to provide in a risk-free way for new members as well as for rising consumption per person.) Thus, in this context, the relevant growth rates in Table 1 are those for the levels variables. These growth rates include population growth and are, not surprisingly, substantially higher than the per capita growth rates.

---

[16]The formula for the risk-free rate is $r^f = \rho + \theta\gamma - (1/2)\cdot\theta^2\sigma^2 - p\cdot[(1-b)^{-\theta} - 1]$—see Barro (2006, Eq. [12]). The expected growth rate is higher, 2.0%, using the empirical distribution of disaster sizes, b, rather than the constant b=0.4.

[17]I excluded 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. See the notes to Table 1 for a listing of years of missing data for the 11 included countries. From a selection standpoint, the most serious concerns are the missing data on bond returns in Germany for 1880-1923 and Sweden for 1880-1921. These omissions cover World War I and the German hyperinflation and, therefore, miss times of low real rates of return on bonds. The mean values shown in Table 1 for bond returns in Germany and Sweden for 1880-2004 use bill returns for the periods of missing bond data.

Table 1 shows three real rates of return—on stocks, short-term bills (analogous to U.S. Treasury bills[18]), and long-term bonds (typically 10-year government bonds). The first observation is that the average real rate of return on stocks—7.4% per year—clearly exceeded the average growth rates. A second observation is that the average real rate of return on bills—1.0%—fell short of the average growth rates. Finally, the average real rate of return on bonds—2.2%—was close to the average per capita growth rates—1.9% for consumer expenditure and 2.1% for GDP—but less than the average growth rates of levels—2.8% and 3.0%, respectively.

In the model (with i.i.d. shocks to GDP), the term structure of risk-free rates is flat. Therefore, in this setting, one could equally well examine short-term or long-term risk-free rates. More generally, a long-term risk-free rate would be relevant for pricing a claim on the consumption trend. A problem in the data is the absence of risk-free assets—indexed U.S. or U.K. government bonds come close, but these securities exist only recently. The observed real returns on conventional short-term bills or long-term bonds are not risk-free returns, given especially the uncertainty about inflation. Unexpected inflation amounts to a form of partial default on nominal debt and tends to be high in bad economic times, especially wartime economic crises. This pattern means that hypothetical risk-free rates tend to be lower than observed averages of real rates of return on bills and bonds. Thus, my inference from the data in Table 1 is that, for the OECD countries in the long run, risk-free rates were likely lower than expected growth rates, particularly growth rates of levels of consumption and GDP.[19]

[18] In some cases, such as in the United States before 1922, the data are for commercial paper.
[19] Another consideration that reinforces this conclusion is that after-tax real returns tend to be lower than the measured gross returns.

When the risk-free rate is below the expected growth rate, the AJ measure of the "marginal cost of all consumption uncertainty" is infinity. AJ demonstrate (section V) that their marginal cost is an upper bound for the total cost of all consumption uncertainty in the sense of Lucas (1987) and in my estimates (with p and σ both set to zero). Obviously, this result is consistent with the finding that the AJ measure of the marginal cost of all consumption uncertainty is likely to be infinite.[20]

This discussion does not invalidate the AJ calculations for the marginal cost of business-cycle fluctuations (their Eq. [4] on p. 1229). This cost refers to the elimination of the part of consumption fluctuations that can be represented as variation around a "trend," defined to be a finite moving-average of past consumption. The AJ expression for this cost is finite even if the risk-free rate is less than the expected growth rate.

V. Concluding Observations

The baseline parameter value σ=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 when considered in comparison to the first half of the 20th century, a turbulent time that featured World Wars I and II and the Great Depression. Hence, a reduction in σ amounts to making

[20] AJ recognize this possibility (p. 1228): "as the [risk-free] yield … gets close to the growth rate g, this value tends to infinity." However, in section V, they argue that the Lucas concept of the total cost of consumption uncertainty (which has to be finite) can be approximated as one-half of their marginal cost (which is likely to be infinite). The apparent contradiction arises because their derivation depends on the assumption that "consumption fluctuations are small," a condition that ensures that the risk-free interest rate exceeds the expected growth rate. Note that the AJ concept of "marginal" that underlies their calculations differs from the one in my derivation of utility rates of transformation for p in Eq. (14) and σ in Eq. (12). The AJ concept is based on what an individual has to pay, given asset prices, to shift the portfolio away from 100% equity to a little less equity and a positive quantity of risk-free claims on the consumption trend. The utility rates of transformation reveal how much the representative household would willingly give up in terms of output to obtain a small reduction in uncertainty (p or σ). These rates of transformation are finite as long as V is finite; that is, if the transversality condition holds.

milder the business fluctuations that were already strikingly tame. Not surprisingly, the benefit from this change—corresponding to around 2% of GDP each year—is only moderate (though still important).

In contrast, the probability parameter p and size parameter b refer to major economic disasters, 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 2001-02. 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 20^{th} century. Included here would be nuclear conflicts, large-scale natural disasters (tsunamis, hurricanes, earthquakes, asteroid collisions), and epidemics of disease (Black Death, avian flu).

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 growth rate, γ^*. 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 γ^*). However, commentaries on monetary policy frequently also stress the roles of central banks in exacerbating or moderating major economic crises. For example, Friedman and 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 1937-38. 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. These policy actions have more to do with lowering p and b than decreasing σ. A key, unresolved issue is whether and how a monetary authority can act—for example, by reacting to potential financial crises—in ways that reduce the probability, p, and size, b, of economic collapses.

Other governmental institutions and policies 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 (γ^*), sometimes focusing on the role of international trade in goods and assets. However, from a political perspective, the main force behind the adoption of these institutions was likely the desire to avoid a repeat 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—for example, 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, F. and U.J. Jermann (2004). "Using Asset Prices to Measure the Cost of Business Cycles," *Journal of Political Economy*, 112, December, 1223-1256.

Atkeson, A. and C. Phelan (1994). "Reconsidering the Cost of Business Cycles with Incomplete Markets," *NBER Macroeconomics Annual*.

Balke, N.S. and R.J. Gordon (1989). "The Estimation of Prewar Gross National Product: Methodology and New Evidence," *Journal of Political Economy*, 97, February, 38-92.

Banerjee, A., R.L. Lumsdaine, and J.H. Stock (1992). "Recursive and Sequential Tests of the Unit-Root and Trend-Break Hypotheses: Theory and International Evidence," *Journal of Business & Economic Statistics*, 10, July, 271-287.

Barlevy, G. (2004). "The Cost of Business Cycles under Endogenous Growth," *American Economic Review*, 94, September, 964-990.

Barro, R.J. (2006). "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, August, 823-866.

Cochrane, J.H. (1988). "How Big Is the Random Walk in GNP?" *Journal of Political Economy*, 96, October, 893-920.

Cogley, T. (1990). "International Evidence on the Size of the Random Walk in Output," *Journal of Political Economy*, 98, June, 501-518.

Friedman, M. and A.J. Schwartz (1963). *A Monetary History of the United States, 1867-1960*, Princeton, Princeton University Press.

Gabaix, X. (2006). "An Equilibrium Asset-Pricing Model with Time-Varying Importance of Rare Disasters," MIT, unpublished.

Lucas, R.E. (1978). "Asset Prices in an Exchange Economy," *Econometrica*, 46, November, 1429-1445.

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

Lucas, R.E. (2003). "Macroeconomic Priorities," *American Economic Review*, 93, March, 1-14.

Maddison, A. (2003). *The World Economy: Historical Statistics*, Paris, OECD.

Mehra, R. and E.C. Prescott (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, March, 145-161.

Obstfeld, O. (1994). "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, 38, August, 1471-1486.

Rietz, T.A. (1988). "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, July, 117-131.

Taylor, B. (2005). "GFD Guide to Total Returns on Stocks, Bonds and Bills," available on the Internet from Global Financial Data at www.globalfindata.com.

Table 1
Growth Rates and Rates of Return for OECD Countries, 1880-2004

	Growth rates				Real rates of return		
Country	**Δc/c**	**Δy/y**	**ΔC/C**	**ΔY/Y**	**stocks**	**bills**	**bonds**
Australia	0.015	0.015	0.032	0.033	0.102	0.012	0.035
Canada	0.021	0.022	0.038	0.039	0.076	0.016	0.035
Denmark	0.018	0.020	0.026	0.028	0.070	0.031	0.032
France	0.016	0.021	0.020	0.025	0.055	-0.011	0.002
Germany	0.023	0.022	0.030	0.027	0.072	-0.019	-0.002*
Italy	0.020	0.022	0.025	0.028	0.048	0.002	0.010
Japan	0.024	0.029	0.034	0.040	0.089	0.004	0.020
Norway	0.019	0.023	0.027	0.030	0.066	0.018	0.026
Sweden	0.020	0.020	0.025	0.026	0.088	0.023	0.031*
U.K.	0.015	0.015	0.019	0.020	0.063	0.017	0.028
U.S.	0.020	0.020	0.034	0.034	0.081	0.015	0.023
Means	**0.019**	**0.021**	**0.028**	**0.030**	**0.074**	**0.010**	**0.022**

Note: c is real per capita consumer expenditure, y is real per capita GDP, C is real consumer expenditure, and Y is real GDP. Growth rates are for annual data for year t compared to year t-1. Real rates of return are calculated from arithmetic annual returns during each year, based on nominal total return indexes and consumer price indexes. (For some country-years, stock returns are based on stock-price indexes and estimates of dividend yields.) Periods for growth rates are 1881-2004 and for returns are 1880-2004, except for the following missing data. Australia is missing growth rates of c and C for 1880-1901. Canada is missing stock returns for 1880-1915 and bill returns for 1880-1899. Denmark is missing growth rates of c and C for 1915-21 and stock returns for 1880-1914. France is missing growth rates of c and C for 1947-49 and stock returns for 1940-41. Germany is missing growth rates of c and C for 1914-25, 1939-40, and 1945-48 and bond returns for 1880-1923. Italy is missing growth rates of c and C for 1881-85 and stock returns for 1880-1905. Japan is missing growth rates of c and C for 1881-85, stock returns for 1880-1914, and bill returns for 1880-82. Norway is missing stock returns for 1880-1914. Sweden is missing stock returns for 1880-1901 and bond returns for 1880-1921. Data on real GDP and population are from Maddison (2003), updated from Economist Intelligence Unit, *Country Data*. U.S. GDP data before 1929 are from Balke and Gordon (1989). Real consumer expenditure is from various sources, based on ongoing research. Data on asset returns and consumer price indexes are from Global Financial Data, discussed in Taylor (2005).

*German bond returns are available only for 1924-2004. The mean value shown is calculated from bill returns for 1880-1923 and bond returns for 1924-2004. Swedish bond returns are available only for 1922-2004. The mean value shown is calculated from bill returns for 1880-1921 and bond returns for 1922-2004.