THE ECONOMIC AND POLICY CONSEQUENCES OF CATASTROPHES*

by

Robert S. Pindyck
Massachusetts Institute of Technology

Neng Wang
Columbia University

This draft: February 26, 2009

Abstract: What is the likelihood that the U.S. will experience a devastating catastrophic event over the next few decades – something that would substantially reduce the capital stock, GDP and wealth? What does the possibility of such an event imply for the behavior of economic variables such as investment, interest rates, and equity prices? And how much should society be willing to pay to reduce the probability or likely impact of such an event? We address these questions using a general equilibrium model that describes production, capital accumulation, and household preferences, and includes as an integral part the possible arrival of catastrophic shocks. Calibrating the model to average values of economic and financial variables yields estimates of the implied expected mean arrival rate and impact distribution of catastrophic shocks. We also use the model to calculate "willingness to pay," i.e., the amount of capital or consumption society would sacrifice to reduce the probability or impact of a shock.

JEL Classification Numbers:

Keywords:

*Comments from ..., research assistance of Ben Lockwood, ... TO BE ADDED.

1 Introduction.

What is the likelihood that the U.S. will experience a devastating catastrophic event over the next few decades? Even if the probability is small, what does the possibility of such an event imply for the behavior of economic variables such as capital investment, interest rates, and equity prices? And how much should society be willing to pay to reduce the probability or the likely impact of such an event?

By "catastrophic event," we mean something national or global in scale that would substantially reduce the capital stock and/or the productive efficiency of capital, thereby substantially reducing GDP, consumption, and wealth. Examples that we have in mind (you can come up with your own) include a nuclear or biological terrorist attack (far worse than even 9/11), a highly contagious "mega-virus" that spreads uncontrollably, a global environmental disaster, or a financial and economic crisis on the order of the Great Depression. Unlike more locally contained events such as Hurricane Katrina or the recent Asian tsunami, as terrible as they were, the events of concern to us would destroy part of the country's (or the world's) productive capital and raise future costs of operating the remaining capital.[1]

The questions we raised above have been the focus of an emerging literature, the roots of which go back to the observation by Rietz (1988) that low-probability catastrophes could, in theory at least, explain the equity premium puzzle, i.e., could help reconcile a relatively large equity premium (4 to 7%) and low real risk-free rate of interest (0 to 2%) with a moderate degree of risk aversion on the part of households. Rietz's article received little or no attention for nearly two decades, until the recent work of Barro (2006, 2009) and Weitzman (2007). Barro (2006) assembled data on "consumption disasters," defined as reductions in real GDP of 15% or greater, for a panel of NNN countries over the past century. Barro estimated the Poisson arrival rate of such events (just under 2% per year), the mean reduction in GDP (MMM%, which he equated to a reduction in the capital stock), and the distribution of the

[1] Those readers who are incurable optimists and/or have limited imaginations should read Posner (2004), who provides additional examples and argues that society fails to take these risks sufficiently seriously, and also Sunstein (2007). For a sobering discussion of the likelihood and possible impact of nuclear terrorism, see Allison (2004). In an excellent review article of Posner's book, Parson (2007) points out the need for a general cost-benefit framework to address these risks in a consistent way.

drop in GDP. Using a pure exchange (fruit tree) model of the economy similar to that of Rietz, Barro showed that these numbers are roughly consistent with the observed equity premium and real risk-free rate in the U.S.[2]

Barro (2009) extended his earlier work by generalizing the model to include an AK production technology and Epstein-Weil-Zin recursive preferences, thereby endogenizing savings and investment, and disentangling the index of relative risk aversion from the elasticity of intertemporal substitution. Using his earlier estimates of the mean arrival rate and impact distribution, he found that the model could again match the observed equity premium and real risk-free rate. He then used the model to estimate the welfare costs (in terms of lost utility from reduced consumption) of rare disasters and found those costs to be very large: depending on the parameters, equivalent to a 15 to 30% reduction in the initial capital stock (and thus initial GDP).[3]

Barro's (2009) AK production model can explain the equity premium and real risk-free rate, but its calibration is inconsistent with other basic economic variables. For example, the model predicts a consumption-investment ratio of about 1:3, instead of matching the roughly 3:1 ratio in the data. Also, because consumption and investment goods are freely interchangable in his model, Tobin's q (marginal and average) always equals one. This means that physical capital does not earn any rents (i.e., there is no difference between physical and financial capital), which is clearly inconsistent with empirical evidence. Finally, like other studies in this emerging literature, Barro uses past data to estimate the mean arrival rate and intensity of the events, and then uses those estimates as exogenous inputs to the model. As discussed below, we take a different approach to determining these event characteristics.

Several authors have extended Barro's work. Gourio, for example, used an exchange economy model with E-W-Z preferences but allowed consumption disasters to have finite

[2]Concurrently, Weitzman (2007) showed that the equity premium and real risk-free rate puzzles could be explained by "structural uncertainty" in which one or more key parameters, such as the true variance of equity returns, and is estimated through Bayesian updating.

[3]We use the word "model," but Barro's paper actually has three different models: a pure exchange economy with E-W-Z preferences, a model with endogenous labor supply, and the AK model that we discuss here, and which Barro focused on for his calibration exercise and welfare calculations.

duration. He found that the effect of recoveries on the equity premium could be positive or negative, depending on the intertemporal elasticity of substitution. Wachter introduced a time-varying Poisson disaster arrival rate, and showed that this could explain the high volatility of the overall stock market (in addition to the equity premium and real risk-free rate).[4]

Other studies have tried to improve upon Barro's (2006) estimates of the event arrival rate and impact. For example, Barro and Ursúa (2008) use an extended dataset ...., consumption instead of GDP, ... [FILL IN]. [Also, Barro, Nakamura, ... (2008). ADD OTHERS?] While these studies provide a better understanding of the characteristics of historical "consumption disasters," they are limited in two fundamental respects. First, many of the events that are treated as independent data points are manifestations of three major global events — the two World Wars and the Great Depression. [ADD SOME NUMBERS ON THIS POINT.] Second, the possible catastrophic events that we think are of greatest interest today have little or no historical precedent — there is simply no historical data, for example, on the frequency or impact of nuclear or biological terrorist attacks.

Consider the forty-year period beginning around 1950 and ending with the breakup of the Soviet Union. During that period there was one potential catastrophic event that dominated all others: the possibility of an all-out nuclear war between the U.S. and the Soviet Union. The fear of such of an event was based partly on the possibility of a mistake: One side might see something threatening on a radar screen, and, unable to get sufficient reassurance from a phone call, launch its own missiles. What was the mean arrival rate for such an event, and what was the probability distribution for its impact? Although numerous studies were done by the Department of Defense, the RAND Corporation, and other organizations to address these questions, there was no historical precedent on which to base estimates.

We take a different approach to analyzing the economics of catastrophes. Unlike the existing literature, we do not try to estimate the mean arrival rate and impact distribution of catastrophic events from historical data, nor do we use the estimates of others. Instead,

[4] Other related studies include Bansal and Yaron (2004), ... ADD OTHERS, DISCUSS BRIEFLY.]

we estimate these characteristics as a calibration *output* of our analysis. In effect, we are assuming that the calibrated characteristics of catastrophes are those perceived by firms and households, in that they are consistent with behavior, and thus with the data for key economic variables.

Behavioral reactions to possible catastrophic events depend in part on preferences. Like some other recent studies, we assume that households have recursive preferences, which involve three behavioral parameters: the rate of pure time preference, the index of relative risk aversion, and the elasticity of intertemporal substitution. There is little agreement among economists regarding the "correct" values for these parameters, but our calibration exercise provides insight into their plausible ranges and relative magnitudes. In particular, we provide evidence that the elasticity of intertemporal substitution is below 0.5, and that expected CRRA utility is a good approximation for modeling preferences.

We specify an AK model of production more general than that used by Barro (2007) in that it includes adjustment costs, so that consumption and investment goods are not freely interchangeable and q need not equal one. Our calibration ties down the level and slope of the adjustment cost function, and thus is consistent with any well-behaved function. However, a specific adjustment cost function is needed for welfare calculations; we specify a quadratic function as a second-order approximation to a more general one.

We model catastrophes as Poisson events with some mean arrival rate, and with an impact that can be completely characterized by one-parameter power probability distribution (which can be transformed into an equivalent exponential or Pareto distribution). Thus the characteristics of catastrophes are completely captured by two parameters.[5] Leaving these two parameters unconstrained, we calibrate our production model so that it fits the basic data for the consumption-investment ratio, the risk-free interest rate, the equity premium, Tobin's q, and the average real growth rate. We thereby calculate the implied characteristics of catastrophes, and also determine how those characteristics vary over the range of reasonable

[5] We are assuming that the impact of a catastrophe is permanent, but note that we can relate this permanent impact to an equivalent (in terms of effects on observed economic variables) larger but temporary impact. The use of a Pareto distribution was suggested by Barro at a December 2008 seminar at Columbia University.

values for the preference parameters.[6] Also, conditional on a value for the index of risk aversion, the calibration yields values for the elasticity of intertemporal substitution, and thus allows us to impose and test the constraint of expected (CRRA) utility.

We use our calibrated model to address the third question that we raised at the outset: How much should society be willing to pay to reduce the probability or the likely impact of a catastrophic event? We calculate a tax-based measure of "willingness to pay" (WTP). In our model a permanent tax on consumption is non-distortionary, equivalent to a lump-sum tax, and equivalent to a reduction in the current capital stock by an amount equal to the tax rate. Thus our WTP is the permanent percentage tax rate that society should be willing to accept to reduce the mean arrival rate of a catastrophic event from its calibrated value to some fraction of that value. This approach allows us to avoid estimating the cost of reducing that mean arrival rate, which is presumably a convex function of the size of the reduction.

In the next section, we lay out a general equilibrium model and show how it incorporates catastrophic shocks, discuss the solution of the model, and explain how its calibration yields information about the nature of those shocks. Section 3 shows the calibration results, and discusses the implications for the nature of catastrophic shocks, and for household preferences. Section 4 discusses our application of the model to policy analysis, and in particular, the calculation of WTP. Section 5 concludes.

2 Framework.

Our analysis requires a tractable general equilibrium model that describes production, capital accumulation, and household preferences, and that includes as an integral part the possible arrival of catastrophic events. Furthermore, the model, when calibrated, should match basic economic variables such as the consumption-investment ratio, the real risk-free rate, the equity premium, Tobin's q, and the real growth rate.

[6] A more general version of the model would allow the impact of a catastrophe to be temporary by assuming that following an initial drop, productivity mean reverts to its original level. This would introduce a third parameter to the characterization of catastrophes — the rate of mean reversion. The three catastrophe parameters could still be estimated as an output of the model's calibration. We leave that to future work.

We construct a general equilibrium continuous-time model in which: (i) a representative consumer has recursive (Duffie-Epstein-Zin-Weil) preferences; (ii) output is given by an AK technology; (iii) investment involves a generalized adjustment cost that reflects the cost and time required to install capital; so that while there are no irreversibilities, we still have $q \neq 1$; (iv) catastrophic shocks occur as Poisson arrivals resulting in the loss of a random fraction of the capital stock. Despite its generality, the model remains tractable, and yields closed-form solutions for equilibrium allocation and pricing.

2.1 Building Blocks.

Preferences. We use the Duffie and Epstein (1992) continuous-time version of Epstein-Zin-Weil preferences, and assume that a representative consumer has homothetic recursive preferences given by:

$$V_t = \mathcal{E}_t \left[\int_t^{\infty} f(C_s, V_s) ds \right] , \tag{1}$$

where

$$f(C, V) = \frac{\rho}{1-\psi^{-1}} \frac{C^{1-\psi^{-1}} - ((1-\gamma)V)^{\omega}}{((1-\gamma)V)^{\omega-1}} . \tag{2}$$

Here $\rho > 0$ is the rate of time preference, ψ is the elasticity of intertemporal substitution, and γ is the coefficient of relative risk aversion. Unlike standard time-additive separable isoelastic utility specifications, this recursive preference allows us to separate risk aversion from the intertemporal elasticity of substitution. To capture how this differs from the standard expected utility, we define the parameter ω as follows:

$$\omega = \frac{1-\psi^{-1}}{1-\gamma} . \tag{3}$$

If $\gamma = \psi^{-1}$ so that $\omega = 1$, we have the standard time-additive separable constant-relative-risk-aversion (CRRA) expected utility, represented by additively separable aggregator:

$$f(C, V) = \frac{\rho C^{1-\gamma}}{1-\gamma} - \rho V. \tag{4}$$

One of the questions we address when we calibrate the model is whether ω is in fact close to 1, so that the simple CRRA utility function is a reasonable approximation to the more general preferences given by eqn. (1). In particular, we analyze and test how equilibrium allocation and pricing constrains the model's parameters, including the elasticity of intertemporal substitution and index of risk aversion.

Production. Aggregate output has the following constant-returns-to-scale AK production technology:

$$Y = AK \ , \tag{5}$$

where A is a constant. It is common to interpret K in this case as a composite capital stock that includes all forms of capital — plant and equipment as well as human and infrastructure capital — but doing so complicates the interpretation of investment and Tobin's q. An alternative interpretation of eqn. (5), which we prefer, is that K is physical capital as traditionally measured, and is used in fixed proportions with labor.[7] Thus if we were calibrating the model to U.S. data, we would expect A to be in the range of 0.2 to 0.3, corresponding to rough measures of the output-capital ratio.

Catastrophic Shocks. We model catastrophic shocks as Poisson arrivals with mean arrival rate λ. Each shock destroys a stochastic fraction $1 - Z$ of the capital stock K, where the remaining capital stock, Z, is a random variable with $0 \leq Z \leq 1$. We assume that Z follows a power distribution over (0,1) with parameter $\alpha > 0$:

$$f_Z(z) = \alpha z^{\alpha - 1} \ \ ; \ \ 0 \leq z \leq 1 \ , \tag{6}$$

so that $\Pr(Z \geq z) = 1 - z^{\alpha}$, and $\mathcal{E}(Z) = \alpha/(\alpha + 1)$. Thus if $\alpha >> 1$, when a catastrophic event occurs, there is a relatively small probability that the resulting Z will be small and the loss $(1 - Z)$ will be large. This would capture the notion that catastrophic events (large loss) occur at low frequencies.

Note that the distribution given by eqn. (6) is quite general. If $\alpha = 1$, Z follows a uniform distribution. More generally, eqn. (6) also implies that $-\ln Z$ is exponentially distributed with mean $\mathcal{E}(-\ln Z) = 1/\alpha$. In addition, the inverse of the remaining fraction of the capital

[7] For example, if $Y = A'K^a L^{1-a}$ and $L = nK$, we have $Y = AK$ with $A = A'n^{1-a}$.

stock, $m = 1/Z$, follows a Pareto distribution given by

$$f(m) = \alpha m^{-\alpha-1} \quad , \quad m > 1 \quad .$$

The Pareto distribution is often used to model random events with fat-tailed distributions.

Investment and Capital Accumulation. Letting I denote aggregate investment, the capital stock K evolves as:

$$dK(t) = \Phi(I(t), K(t))dt + \sigma K(t)dW(t) - (1 - Z)K(t)dJ(t) \ . \tag{7}$$

Here the volatility parameter σ captures "normal" volatility, and $J(t)$ is a pure jump process with mean arrival rate λ that captures catastrophic events; if a jump occurs, K falls by the random amount $(1 - Z)$. The function $\Phi(I, K)$ is a generalized cost of adjustment function that captures the effects of depreciation and the costs of installing capital and making it productive. We assume that $\Phi(I, K)$ is homogeneous of degree one in I and K and thus can be written as:

$$\Phi(I, K) = \phi(i)K \ , \tag{8}$$

where $i = I/K$ and $\phi(i)$ is increasing and concave. The solution (and thus calibration) of the model uses only the level of $\phi(i^*)$ and it slope $\phi'(i^*)$ at the equilibrium investment-capital ratio, i^*, and thus does not require any particular specification for this function. However, WTP calculations use off-equilibrium values for i, making it necessary to specify a particular function. We use a quadratic specification for the adjustment cost function, which can be viewed as a second-order approximation to a more general one: $\phi(i) = i + \frac{1}{2}\theta i^2 - \delta$.

Competitive Equilibrium. We add two more conditions: (1) Investment always equals savings, i.e., $I(t) = Y(t) - C(t)$ at all $t \geq 0$. (2) Financial markets clear, i.e., there is a zero net supply of bonds, and total wealth is all marketable equity and thus equals the market value of the capital stock $Q(t)$. With these conditions, we can solve the social planner's problem to obtain the competitive equilibrium.

The Bellman equation for the social planner's problem is:

$$0 = \max_C \left\{ f(C, V) + \Phi(I, K)V'(K) + \frac{1}{2}\sigma^2 K^2 V''(K) + \lambda \mathcal{E}\left[V(ZK) - V(K)\right] \right\} \ . \tag{9}$$

Using the identity $C + I = Y$, we have the following first-order condition for I:

$$f_C(C,V) = \phi'(i)V'(K) \ . \tag{10}$$

The left-hand side of eqn. (10) is the marginal benefit of consumption and the right-hand side is the marginal cost of consumption, which is equal to the product of the marginal value of capital $V'(K)$ and the marginal efficiency of converting a unit of the consumption good into a unit of capital, $\phi'(i)$.

2.2 Model Solution.

Exploiting the homogeneity property of the value function, we conjecture and later confirm that value function takes the following form:

$$V(K) = \frac{1}{1-\gamma}(bK)^{1-\gamma} \, , \tag{11}$$

where b is a coefficient determined as part of the solution. Let $c = C/K = A - i$. (Throughout this paper, we use lower-case letters to express quantities relative to the capital stock K.) In the Appendix, we show that b is related to the equilibrium levels of consumption and investment by:

$$b = (A-i)^{1/(1-\psi)}\left(\frac{\rho}{\phi'(i)}\right)^{-\psi/(1-\psi)} \, , \tag{12}$$

where the equilibrium investment-capital ratio i solves the following implicit equation:

$$A - i = \frac{1}{\phi'(i)}\left[\rho + (\psi^{-1}-1)\left(\phi(i) - \frac{\gamma\sigma^2}{2} - \frac{\lambda}{1-\gamma}\mathcal{E}\left(1 - Z^{1-\gamma}\right)\right)\right] . \tag{13}$$

Equilibrium capital accumulation is then given by

$$dK(t)/K(t) = gdt + \sigma dW(t) - (1-Z)dJ(t) \ , \tag{14}$$

where $g = \phi(i^*)$ is the expected growth rate (conditional on no disaster) and i^* is the solution of eqn. (13).

Using the solution sketched out above, we show in the Appendix that the model can be summed up by five key equations, which can be interpreted as the decentralized competitive-market implementation of the social planner's solution:

$$i = A - c \tag{15}$$

$$q = \frac{1}{\phi'(i)} \tag{16}$$

$$c = \left[\rho + (\psi^{-1} - 1)\left(\phi(i) - \frac{\gamma\sigma^2}{2} - \frac{\lambda}{1-\gamma}\mathcal{E}\left(1 - Z^{1-\gamma}\right)\right)\right] q \tag{17}$$

$$r = \rho + \psi^{-1}\phi(i) - \frac{\gamma(\psi^{-1}+1)\sigma^2}{2} - \lambda\mathcal{E}\left[\left(\psi^{-1} - \gamma\right)\left(\frac{1 - Z^{1-\gamma}}{1-\gamma}\right) + \left(Z^{-\gamma} - 1\right)\right] \tag{18}$$

$$r^e - r = \gamma\sigma^2 + \lambda\mathcal{E}\left[(1-Z)\left(Z^{-\gamma} - 1\right)\right] \tag{19}$$

Eqn. (15) is simply an accounting identity that equates saving and investment. Eqn. (16) is the first-order condition for producers. Re-writing it as $\phi'(i)q = 1$, it equates the marginal benefit of an extra unit of investment (which at the margin yields $\phi'(i)$ units of capital, each of which is worth q) with its marginal opportunity cost (1 unit of the consumption good).[8]

Eqn. (17) is the first-order condition for consumers. It equates consumption (normalized by the capital stock) to the marginal propensity to consume (MPC) out of wealth (everything in the square brackets) times q, which is the value of a marginal unit of wealth. Note that from the consumer's perspective, the entire capital stock is marketable, and its value is qK. What drives the marginal propensity to consume? Looking inside the square brakets, ρ is the rate of time preference; if $\psi = 1$, wealth and substitution effects just offset each other, and the MPC $= C/W = c/q$ is pinned down by ρ. More generally, if $\psi^{-1} < (>)\ 1$, the MPC $< (>)\ \rho$ by an amount proportional to the quantity in the large parentheses, which is the risk-adjusted growth rate: With no stochastic changes in K, the growth rate g is $\phi(i)$; that rate is reduced by "normal" (diffusive) fluctuations in K (the second term) and by the prospect of downward jumps in K (the last term). (Note that the percentage change of value function $(1 - Z^{1-\gamma})/(1-\gamma)$ is always positive conditioning on a downward jump.)

Eqn. (18) for the equilibrium interest rate r may be interpreted as a generalized Ramsey rule. If $\psi^{-1} = \gamma$ so that preferences simplify to CRRA expected utility, and if there were no stochastic changes in K, the deterministic Ramsey rule $r = \rho + \gamma g$ would hold. More generally, the term $\gamma\sigma^2/2$ is the standard adjustment for continuous stochastic fluctuations in K, and the last term adjusts for downward jumps in K.

[8] The AK model of Barro (2009) has no adjustment costs, so $\phi'(i) = q = 1$. In his model, capital is perfectly liquid and commands no rents.

Finally, eqn. (19) describes the equity risk premium. The first term on the RHS is the usual risk premium in diffusion models (see, e.g., Breeden (1979) and Lucas (1978)), and the second term is the increase in the premium that is due to jumps in K. When a jump occurs, $(1 - Z)$ is the fraction of loss, and $(Z^{-\gamma} - 1)$ is the percentage increase in marginal utility from that loss, i.e., the price of risk. The incremental expected excess return is given by the product of these two terms. Note that the fraction of losses and the increase of marginal utility are positively correlated. This positive correlation substantially contributes to the risk premium. For example, consider the limiting case where the loss is close to 100%, then the marginal utility increase is effectively infinite.

2.3 Calibration Procedure.

We calibrate the model using "average" or "consensus" values for the following variables or parameters, based on U.S. data: the consumption-investment ratio c/i, marginal (and average) q, the real risk-free interest rate r, the equity risk premium rp, the average normal growth rate g, and the normal volatility σ^2. In addition, we specify values for the index of risk aversion γ and the rate of time preference ρ.

This leaves six unknowns: the economic variables c, i and A; the elasticity of intertemporal substitution ψ, and the two parameters describing the characteristics of catastrophic shocks, λ and α. We have the five equations (15) to (19), so the set of unknowns is under-identified. We therefore calibrate the model for different values for A, and consider solutions over the range of admissible values for A, i.e., values for which the resulting $\lambda > 0$ and $\psi > 0$. In practice, this limits A to a fairly narrow range (and thus roughly identifies A). We also calibrate the model imposing the constraint of expected CRRA utility, i.e., $\psi = 1/\gamma$, in which case the five remaining unknowns are exactly identified.

3 Results.

Our objective is to estimate λ and α, as well as the elasticity of intertemporal substitution ψ as outputs of the model's calibration. We will examine how these endogenous parameters vary as we change key unknown inputs, such as the index of risk aversion γ and the productivity parameter A. We proceed in steps, and begin by first reviewing some results of Barro (2009), in which λ and the distribution for Z were exogenous inputs. A comparison of our results with those of Barro is useful because it illustrates the importance and implications of adjustment costs, and the implications of certain parameter choices. We then discuss some of the key parameter choices for our model, present a set of calibration results, and discuss the implications of those results for preference parameters, and for the values of λ and α.

3.1 Initial Calibration: Exogenous Catastrophe Characteristics.

Based on historic "consumption disasters" for a large panel of countries, Barro (2009) estimated the mean arrival rate λ to be .017. He set $\gamma = 4$, and based on an empirical distribution for the remaining post-disaster capital stock Z, he estimated the three moments $\mathcal{E}(Z)$, $\mathcal{E}(Z^{1-\gamma})$, and $\mathcal{E}(Z^{-\gamma})$. He also chose the following values for the remaining parameters: $\psi = 2$, $\rho = .052$, $\sigma = .02$, and $A = .174$. As we noted earlier, economists differ in their views about ψ, but a value of 2 is certainly at the high end of the range of numbers that have appeared in the literature. Likewise, a more typical value for ρ, the rate of pure time preference, would be closer to .02.

In Barro's model, consumption and capital goods are freely interchangeable, i.e., there are no adjustment costs, so $q = 1$. Table 1 shows this calibration of Barro's model; it also shows how the results change as the adjustment cost parameter θ is increased so that q increases. Note that with no adjustment costs ($\theta = 0$), the model gives a sensible estimate of the risk-free rate r, but yields an investment-to-capital ratio close to 3, whereas the actual ratio is closer to 1/4 to 1/3. This is a fundamental inconsistency with the basic economic facts.

If we increase the adjustment cost parameter θ, i falls and c increases. When θ is in the range of 6 to 8, i/c roughly matches the actual data, and Tobin's q is around 1.4, which is also a rough match to the data. However, the real risk-free rate falls below -3%, so that once again there is a miss-match. The basic problem is that the preference parameters (particularly ψ and ρ) and the exogenous inputs for λ and the moments of Z simply cannot

Table 1: Calibration to Barro Parameters

Note: $\gamma = 4$, $\psi = 2$, $\rho = .052$, $\sigma = .02$, $A = .174$, $\lambda = .017$, $\mathcal{E}(Z) = .71$, $\mathcal{E}(Z^{1-\gamma}) = 4.05$, $\mathcal{E}(Z^{-\gamma}) = 7.69$				
θ	i	c	r	q
0	0.126	0.048	0.011	1.000
2	0.088	0.086	-0.012	1.213
4	0.062	0.112	-0.025	1.333
6	0.047	0.127	-0.032	1.396
8	0.038	0.136	-0.036	1.434
10	0.031	0.143	-0.039	1.458
20	0.017	0.157	-0.045	1.510

match all of the basic economic fact, even when we allow for adjustment costs. The model is too constrained.

3.2 Endogenous Catastrophe Characteristics.

We turn now to calibrations of the model for which the catastrophe characteristics are endogenous. In particular, the mean arrival rate λ and the impact distribution parameter α are unconstrained, as are the productivity parameter A and the elasticity of intertemporal substitution, ψ. Constrained variables for which we impose numerical values are as follows: the consumption-investment ratio $c/i = 3$, marginal (and average) $q = 1.5$, the real risk-free interest rate $r = .02$, the equity risk premium $rp = .06$, the average "normal" (i.e., absent a catastrophe) growth rate $g = .025$, and the normal volatility $\sigma = .02$. In addition, for the rate of pure time preference, we use $\rho = .02$, and we use two alternative values for the index of risk aversion γ, 2 and 4.

The calibration results are summarized in Table 2. For each value of γ, the table shows the admissible range of A, i.e., the values of A for which the resulting values of ψ and λ are positive. This admissible range for A turns out to be about .11 to .12 or .13, numbers that are somewhat low compared to historical data for the output-to-capital ratio. (If we set q to be 2 instead of 1.5, this range is instead about .15 to .17.) For each A, the table shows as calibration outputs the resulting values of ψ, λ, α and $\mathcal{E}(1 - Z)$. (Note that α and $\mathcal{E}(1 - Z)$

Table 2: Calibration 2: $q = 1.5$

Note: $r = .02$, $c/i = 3$, $rp = .06$, $\rho = .02$, $\sigma = .02$, $g = .025$				
$\gamma = 2$ and $q = 1.5$				
A	ψ	λ	α	$\mathcal{E}(1-Z)$
0.112	0.3769	0.0031	2.0983	0.3228
0.114	0.3421	0.0064	2.1918	0.3133
0.116	0.3081	0.0098	2.2818	0.3047
0.118	0.2745	0.0135	2.369	0.2968
0.120	0.2413	0.0173	2.454	0.2895
0.122	0.2084	0.0212	2.5373	0.2827
0.124	0.1757	0.0253	2.619	0.2763
0.126	0.1432	0.0296	2.6996	0.2703
0.128	0.1108	0.034	2.779	0.2646
$\gamma = 4$ and $q = 1.5$				
0.112	0.3584	0.0053	4.2956	0.1888
0.114	0.3149	0.0111	4.5407	0.1805
0.116	0.2746	0.0173	4.761	0.1736
0.1173	**0.25**	**0.0214**	**4.8929**	**0.1697**
0.118	0.2362	0.0239	4.966	0.1676
0.12	0.1991	0.0308	5.1605	0.1623
0.122	0.1628	0.0381	5.3472	0.1576
0.124	0.1272	0.0457	5.5279	0.1532

directly related.)

There are several important conclusions that can be drawn from the results in this table. First, these calibrations yield values values of ψ that are low — well below 0.5. Estimates of ψ in the literature vary considerably, and include low numbers in the range of those in Table 2 to values close to 2.[9] But as these calibrations (as well as others that follow) show, basic macro data are inconsistent with values of ψ above 0.5.

Second, these results are consistent with the view that restricting preferences to expected CRRA utility is not a bad approximation for modeling purposes. When $\gamma = 2$, expected utility implies that $\psi = 0.5$, and although our estimates of ψ are below 0.5, they are not far

[9] ADD FOOTNOTE WITH REFERENCES: GRUBER, VISSING-JORGENSON, OTHERS.

below. When $\gamma = 4$, expected utility implies $\psi = 0.25$, which is in the middle of the range of our estimates for ψ, and is shown as the row of boldface numbers in Table 2.

Third, we are unable to estimate the mean arrival rate λ with any kind of precision. For $\gamma = 2$ ($\gamma = 4$), the estimates of λ over the admissible range of A range from roughly .003 to .03 (.005 to .05), i.e., they vary by an order of magnitude. On the other hand, the estimates of the expected loss $\mathcal{E}(1-Z)$ depend on γ, but given γ, are reasonably well identified. If, for example, one had a a strong prior that $\gamma = 2$, then the expected loss should a catastrophic shock occur is in the range of 26 to 32 percent. But if $\gamma = 4$, that range becomes 15 to 19 percent. Thus given a value for γ, the variation in our estimates of λ is much greater than the variation in the estimated expected loss.

As mentioned before, the values of A that result from these calibrations are lower than what the data would suggest. Higher values of q, however, result in larger values of A. Table 3 shows calibration results for $q = 2$, with the other parameters and variable the same, and again γ is 2 and 4.

Observe that now A is in the range of .15 to .17. The resulting values of ψ are about the same as in Table 2, as are the ranges for λ and $\mathcal{E}(1-Z)$. (Once again, the numbers that correspond to expected utility are shown in bold.) Higher values of q help to match the basic macroeconomic facts (and certainly $q > 1$ is needed), but do not change the estimates of λ or $\mathcal{E}(1-Z)$ in any significant way.

In summary, the model calibrates well against the data for basic macro variables (particularly if one is willing to accept that q exceeds 1.5). Given a prior on the index of risk aversion, the model provides a good fix on the expected loss, $\mathcal{E}(1-Z)$. However, the model does not allow us to pin down a value for the mean arrival rate λ. Thus when we use the model for policy analysis (see below), we must do so for different values of A and hence λ.

4 Policy Implications.

We now turn to the last question raised in the first paragraph of this paper: how much should society be willing to pay to reduce the probability or the likely impact of such a catastrophic

Table 3: Calibration 3: $q = 2$

Note: $r = .02$, $c/i = 3$, $rp = .06$, $\rho = .02$, $\sigma = .02$, $g = .025$				
$\gamma = 2$ and $q = 2$				
A	ψ	λ	α	$\mathcal{E}(1-Z)$
0.148	0.3947	0.0015	2.0499	0.3279
0.15	0.3682	0.0039	2.1221	0.3203
0.152	0.3421	0.0064	2.1918	0.3133
0.154	0.3165	0.009	2.2596	0.3068
0.156	0.2912	0.0116	2.3257	0.3007
0.158	0.2662	0.0144	2.3904	0.2949
0.16	0.2413	0.0173	2.454	0.2895
0.162	0.2166	0.0202	2.5166	0.2844
0.164	0.192	0.0233	2.5783	0.2795
0.166	0.1676	0.0264	2.6393	0.2748
0.168	0.1432	0.0296	2.6996	0.2703
0.17	0.1189	0.0329	2.7592	0.266
$\gamma = 4$ and $q = 2$				
0.148	0.3823	0.0026	4.1573	0.1939
0.150	0.3471	0.0067	4.3601	0.1866
0.152	0.3149	0.0111	4.5407	0.1805
0.154	0.2845	0.0157	4.7076	0.1752
0.156	0.2552	0.0205	4.8651	0.1705
0.1564	**0.2500**	**0.0214**	**4.8929**	**0.1697**
0.158	0.2268	0.0256	5.0155	0.1662
0.160	0.1991	0.0308	5.1605	0.1623
0.162	0.1718	0.0362	5.3011	0.1587
0.164	0.1450	0.0418	5.4382	0.1553
0.166	0.1184	0.0476	5.5723	0.1522

event? To address this question, we consider a consumption tax. Our measure of willingness to pay is the maximum *permanent* consumption tax rate τ that society would be willing to accept if the resulting stream of government revenue could finance whatever activities would be necessary to permanently reduce the mean arrival rate of a catastrophe from λ to λ'. (Later we also consider a tax that would reduce the average impact of a catastrophe.)

What is the effect of a permanent consumption tax? Given investment $I(t)$ and output $Y(t)$, households pay a stochastic stream of tax $\tau(Y - I)$ to the government and consume the remaining amount:

$$C(t) = (1 - \tau)[Y(t) - I(t)] \ . \tag{20}$$

Therefore, the first-order condition of eqn. (10) is now

$$(1 - \tau)f_C(C, V) = \phi'(i)V'(K) \ . \tag{21}$$

A consumption tax has two opposing effects. First, it lowers the marginal benefit of current consumption, making investment more attractive, all else equal. Second, because the tax is permanent, the marginal benefit of investing is lower because any future output generated from the additional capital will also be taxed at the same rate. As shown below and in the Appendix, in equilibrium, these two forces exactly offset each other, leaving the investment decision unaffected. That is, the optimal investment-capital ratio i still solves eqn. (13). Further more, the permanent consumption tax is equivalent in terms of welfare to an intitial one-time reduction of the capital stock by the amount of the tax rate, τ.

The intuition behind this result is that because a permanent proportional tax lowers consumption by the same fraction in each period, it leaves households' intertemporal marginal rate of substitution unchanged. Therefore, holding the likelihood of a catastrophic shock, λ, fixed, the equilibrium investment and growth paths are the same with or without the tax. A permanent consumption tax in our model is therefore non-distortionary; households have no incentive to change their intertemporal consumption/investment decisions in response to the tax.

4.1 Willingness to Pay.

How large a tax rate would society accept to reduce the likelihood of a catastrophe from λ to λ'? Holding λ fixed, let $V(K;\lambda,\tau)$ and $b(\lambda;\tau)$ denote the value function and the corresponding value function coefficent when there is a permanent consumption tax rate at rate τ, and let $V^*(K;\lambda)$ and b^* denote the same quantities in the absence of a tax, i.e., $V^*(K;\lambda) = V(K;\lambda,0)$ and $b^*(\lambda) = b(\lambda;0)$. We can conjecture and then verify (see the Appendix) that

$$V(K;\lambda,\tau) = \frac{1}{1-\gamma}\left(b(\lambda;\tau)K\right)^{1-\gamma}, \tag{22}$$

where

$$b(\lambda;\tau) = (1-\tau)b^*(\lambda). \tag{23}$$

Recall that $b^*(\lambda)$ is the value function coefficient given in eqn. (12), which is evaluated at the equilibrium i^* (without a tax).

Therefore, households are indifferent between (a) no tax and a likelihood of catastrophe λ and (b) paying a permanent tax τ to reduce the likelihood to λ' if the following condition holds:

$$V(K;\lambda',\tau) = V(K;\lambda,0). \tag{24}$$

Using equations (22 and (23), we obtain the following result:

$$b(\lambda';\tau) = (1-\tau)b^*(\lambda') = b^*(\lambda). \tag{25}$$

To reduce the likelihood of a catastrophe from λ to λ', households would be willing to pay a consumption tax at the following constant tax rate:

$$\tau(\lambda,\lambda') = 1 - \frac{b^*(\lambda)}{b^*(\lambda')}. \tag{26}$$

As can be seen from eqns. (22), (23), and (26), a permanent tax at rate τ is equivalent to giving up a fraction τ of the capital stock.

4.2 Calibration Results.

Table 4 shows the maximum permanent tax rate that society would accept to reduce λ by 20 percent (i.e., $\lambda'/\lambda = .8$), 50 percent, and 100 percent for various values of A corresponding

Table 4: WTP Calculations

Note: $r = .02$, $c/i = 3$, $rp = .06$, $\rho = .02$, $\sigma = .02$, $g = .025$ $\gamma = 2$ and $q = 1.5$					
			$\lambda'/\lambda = .8$	$\lambda'/\lambda = .5$	$\lambda'/\lambda = 0$
A	λ	$\mathcal{E}(1-Z)$		τ^*	
0.112	.0031	0.3228	.03	.06	.12
0.122	.0212	0.2827	.18	.36	.54
0.128	.0340	0.2646	.36	.59	.75

to the calibrations in Table 2. For each value of A, we also show the resulting values of λ and $\mathcal{E}(1-Z)$. Whatever the percentage reduction in λ, these tax rates vary substantially depending on the base value of λ.

These results are not surprising. As we saw in Tables 2 and 3, and as can be seen in Table 4, λ varies greatly over the range of admissible values of A, but $\mathcal{E}(1-Z)$ varies much less. Thus we would expect that the WTP to reduce λ from .0340 to zero when the expected loss is 26 percent would be much larger than the WTP to reduce λ from .0031 to zero when the expected loss is 32 percent.

ADD TABLE WITH FIXED λ?

Given that we cannot pin down λ, we are also unable to pin down the WTP for, say, the elimination of catastrophic risk. We can only calculate that WTP conditional on some prior regarding λ. Estimating λ remains a key problem.

5 Conclusions.

TO BE ADDED.

Appendix

A. Solution of Model.

TO BE ADDED.

B. Proof that Permanent Consumption Tax is Non-Distortionary.

TO BE ADDED.

References

Allison, Graham, *Nuclear Terrorism: The Ultimate Preventable Catastrophe*, Henry Holt & Company, 2004.

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

Barro, Robert J., "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, forthcoming, March 2009.

Barro, Robert J., and José Ursúa, "Macroeconomic Contractions since 1870," *Brookings Papers on Economic Activity*, Spring 2008. 255–350.

Breeden, Douglas T., "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 1979, **7**, 265–296.

Gourio, François, "Disasters and Recoveries," *American Economic Review*, May 2008, **98**, 68–73.

Lucas, Robert E, Jr, "Asset Prices in an Exchange Economy," *Econometrica*, **46**, 1429–45

Parson, Edward A., "The Big One: A Review of Richard Posner's *Catastrophe: Risk and Response*, *Journal of Economic Literature*, March 2007, **45**, 147–164.

Posner, Richard A., *Catastrophe: Risk and Response*, Oxford University Press, 2004.

Sunstein, Cass R. *Worst-Case Scenarios*, Harvard University Press, 2007.

Wachter, Jessica, "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" NBER Working Paper No. 14386, October 2008.

Weitzman, Martin L., "Prior-Sensitive Expectations and Asset-Return Puzzles," *American Economic Review*, 2007(a).