The Role of Uncertainty in the Economics of Catastrophic Climate Change

Martin L. Weitzman*

June 9, 2007. Slightly preliminary. Comments appreciated

Abstract

Using climate change as a prototype motivating example, this paper analyzes the implications of structural uncertainty for the economics of catastrophes. The paper shows that having an uncertain multiplicative parameter, which amplifies or scales exogenous impulses and is updated by Bayesian learning, induces a critical "tail fattening" of posterior-predictive distributions. Such fattened tails have strong implications for situations (like climate change) where a catastrophe is theoretically possible because prior knowledge cannot place sufficiently narrow bounds on overall damages. Under strict relative risk aversion, the impact of an uncertain scale of damages on cost-benefit analysis outweighs the impact of discounting.

1 Introduction

A long-standing theme with important ramifications concerns a basic distinction between risk and uncertainty. "Risk" is intended to signify a random situation where probabilities are known measurable frequencies but realized future outcomes are not yet known. "Uncertainty" reflects a lack of knowledge at a deeper level, where the probabilities driving future outcomes are themselves unknown because information is limited, vague, or ambiguous.[1] Loosely speaking, in the distinction I am trying to make throughout this paper I identify

*Department of Economics, Harvard University, Cambridge, MA 02138 (e-mail: mweitzman@harvard.edu). For helpful detailed comments on earlier drafts of this paper, but without implicating them for its remaining defects, I am grateful to Frank Ackerman, Richard Carson, Stephen DeCanio, Michael Mastrandrea, Cedric Philibert, Richard Posner, Richard Tol, Gary Yohe, and Richard Zeckhauser.

[1] Warning: what economists name "risk" and associate narrowly with known probabilities, scientists call "uncertainty" – while what economists name "uncertainty" and associate narrowly with unknown probabilities, scientists label as "deep uncertainty" or "structural uncertainty."

"risk" with having *known PDFs* (probability density functions) while "uncertainty" is about *PDFs of PDFs.* I will argue in this paper that a useful way to think about the difference between risk and uncertainty identifies risk with the stochastic outcome of a data-generating process whose structural parameters are *known*, while uncertainty refers to the part of a data-generating process concerned with structural parameters that are *unknown* and must be estimated. This distinction will be shown to have surprisingly strong consequences for applications of expected utility theory under strict relative risk aversion.

The key unknown structural parameter in this paper is a crucial multiplier that amplifies (or scales) uncertain exogenous impulses, which are subsequently propagated throughout the system. Very crudely – at an extremely high level of abstraction and without taking the analogy too literally – the role of this uncertain multiplicative amplifier or scale parameter can be illustrated by the role of uncertain "climate sensitivity" in discussions of global warming.

Let $\Delta \ln CO_2$ be sustained relative change in atmospheric carbon dioxide while ΔT is equilibrium temperature response. Climate sensitivity S is an amplifying or scaling multiplier for converting $\Delta \ln CO_2$ into ΔT by the (reasonably accurate) linear approximation $\Delta T \approx (S/\ln 2) \times \Delta \ln CO_2$. From the IPCC4 (2007) Executive Summary: "The equilibrium climate sensitivity is a measure of the climate system response to sustained radiative forcing. It is not a projection but is defined as the global average surface warming following a doubling of carbon dioxide concentrations. It is *likely* to be in the range 2 to 4.5°C with a best estimate of 3°C, and is *very unlikely* to be less than 1.5°C. Values substantially higher than 4.5°C cannot be excluded, but agreement of models with observations is not as good for those values." (For IPCC4, "likely" is $P > 66\%$ while "very unlikely" is $P < 10\%$.)

In this paper I am mostly concerned with the 17% of those S "values substantially higher than 4.5°C" which "cannot be excluded." Eighteen recent studies of climate sensitivity with 18 PDFs of S lie behind the above-quoted IPCC4 summary statement. From Figure 1 in Box 10.2 of IPCC4 (2007), it is apparent that the right tails of these 18 PDFs tend to be long (and thick). For graphical neatness, Box 10.2 arbitrarily truncates all PDFs at 10°C, but they are actually more open-ended than this in most of the underlying studies. The upper 5% confidence level averaged over all 18 studies is $\overline{S}$=6.2°C, which I take as approximately meaning that $P[S > 6.2°C] \approx 5\%$. For the purposes of this paper I assume $P[S > 8°C] \approx 2\%$, which is very roughly consistent with the averaged right tails in Box 10.2. Such numbers are highly-speculative ballpark estimates, but the subject matter of this paper concerns just such kind of highly-speculative open-ended scale parameters and my example here does not depend at all on precise numbers.

Societies and ecosystems whose average temperatures have changed in the course of one or two centuries by $\Delta T > 8°C$ (for U.S. readers: $8°C = 14.4°F$) are located in *terra incognita*, since such high temperatures have not existed for at least tens of millions of years. IPCC4 (2007) gives several truly frightening examples of possible consequences of such kind of temperature changes: sudden disintegration of the Greenland or West Antarctic ice sheets with dramatic raising of sea level, shutdowns or even reversals of the warming component of large-scale oceanic circulation systems like the Gulf Stream, major disruptions of large-scale weather patterns like monsoons, runaway greenhouse warming due to endogenous heat-induced rapid releases of the immense amounts of GHGs (greenhouse gases) currently sequestered in arctic permafrost or offshore methane hydrates, and so forth. It is difficult to imagine what $\Delta T > 8°C$ might mean for life on earth, but such temperature changes (which would be geologically instantaneous) are larger than what separates us now from past ice ages. Furthermore, we are far from knowing that anthropogenically-injected GHG stocks will be stabilized at anything like twice pre-industrial-revolution levels. Given current trends in GHG emissions, we will attain such a doubling within about 30-40 years and will then go well beyond that amount unless relatively drastic measures are taken starting soon. A rough guesstimate of the equilibrium temperature response to a *tripling* of GHG concentrations (which, projecting current trends in GHG emissions, would be attained relative to pre-industrial-revolution levels within about a century) might in this case be $P[\Delta T > 8°C] \approx$ 3-4%.

The numbers for probabilities of extreme disasters being cited above are all such crude ballpark estimates that there is a tendency in the literature to dismiss them on the grounds that they are much too highly speculative to be taken seriously. By contrast with the conventional wisdom of *not* taking seriously extreme-temperature-change probabilities *because* such crude estimates are highly speculative, the purpose of this paper is to prove the exact opposite logic by giving a rigorous sense in which (other things being equal) the more speculative and fuzzy are the tail probabilities of extreme events, the more serious is the situation for an agent whose welfare is measured by present discounted expected utility. The core difficulty comes from combining strict relative risk aversion with a highly-uncertain, effectively-open-ended impulse-amplifying multiplicative factor.

How warm the climate ultimately gets is approximately a product of two factors – the amount of GHG concentrations and a critical climate-sensitivity-like scaling multiplier. Both of these factors are uncertain, but the scaling multiplier is much more open-ended on the high side with a much longer right tail. The main point I am trying to make with these highly-speculative extreme climate-change examples is to suggest that the great degree of *scientific* indeterminacy enters cost-benefit analysis in the reduced form of some highly-uncertain open-ended climate-sensitivity-like amplifying or scaling factor. This critical scale parameter is then used as a multiplier for converting aggregated GHG emissions – which are largely *economic* variables or unknowns – into eventual temperature changes. The paper will show that a generalization of this same *modus operandi* packaged in the form of a macroeconomic model with essentially the same reduced form (structural uncertainty about some unknown open-ended scale-multiplier parameter) can have very strong consequences for cost-benefit analysis, because it can drive applications of expected utility theory under strict relative risk aversion. It will turn out in this paper that the fact that empirical climate sensitivity estimates have fat-tailed PDFs is preordained in a certain sense (from being constructed out of inductive experience), rather than being merely coincidental.

When it is fed back into an economic analysis, the great open-ended uncertainty about eventual temperature changes becomes yet-greater yet-more-open-ended uncertainty about eventual changes in utility. Not only is it very difficult to estimate tail probabilities of high-ΔT outcomes, but translating this via ambiguous weather consequences and unknown adaptation capabilities into utility-equivalent units for people living a century or two from now introduces enormous further fuzziness, especially when such utility includes the existence values which people at that future time may place on wild species, *in situ* conservation, and natural or historical habitats. Even if climate sensitivity were bounded above by some very high number, the value of what might be called "welfare sensitivity" is effectively unbounded.

Without further belaboring the point, the overall utility effects of global warming that might accompany a 2% chance of $S > 8°\text{C}$ are sufficiently open-ended and fuzzy that it seems fair to say offhand there might be a non-negligible probability of a catastrophe. In his book *Catastrophe: Risk and Response*,[2] Richard A. Posner defines the word catastrophe "... to designate an event that is believed to have a very low probability of materializing but that if it does materialize will produce a harm so great and sudden as to seem discontinuous with the flow of events that preceded it." Posner adds: "The low probability of such disasters – frequently the *unknown* probability, as in the case of bioterrorism and abrupt global warming – is among the things that baffle efforts at responding rationally to them." In this paper I address what rational economic analysis in the form of expected present discounted utility theory might offer in the way of guidance for thinking coherently about the economics of uncertain catastrophes (using Posner's approximate sense of the term).

The prime general example of an unknown structural parameter that best illustrates the critical distinction I am trying to make between risk and uncertainty is a standard-deviation-like scale parameter whose way of interacting parallels the interaction of a climate-sensitivity-like amplifying multiplier. Suppose the true value of this generic scale parameter is unknown because of limited past experience, which situation can be modeled *as if* inferences must

[2] Posner (2004). See also the insightful review by Parson (2007).

be made inductively from a finite number of data observations. At a very high level of abstraction, each data point might be interpreted as representing an outcome from a particular study. This paper shows that having an uncertain scale parameter in such a setup invariably adds a significant tail-fattening effect to posterior-predictive expectations, even when Bayesian learning takes place with arbitrarily large (but finite) amounts of data. The driving mechanism is that, loosely speaking, the operation of taking "expectations of expectations" or "probability distributions of probability distributions" over an uncertain scale parameter invariably spreads apart and fattens the tails of the reduced-form compounded posterior-predictive probability distribution. Rare disasters located in the stretched-out fattened tails of such posterior-predictive distributions must have a large component of uncertainty because it is inherently impossible to learn their true probabilities of occurrence from finite samples alone. The underlying sampling-theory principle is that the rarer is an event the more unsure is our estimate of its probability.

This paper suggests that standard approaches to modeling the economics of climate change (even those that purport to treat risk by Monte Carlo simulations) very likely fail to account adequately for the implications of uncertain large consequences with small probabilities. The overarching general message is that from inductive experience alone one cannot acquire sufficiently accurate information about the probabilities of tail disasters to prevent the expected marginal utility of an extra sure unit of consumption from becoming unbounded for any utility function having strict relative risk aversion. Instead, to eliminate or reduce the possibilities of bad extremes one must rely on "extra" prior information *outside* the data-evidence – and subsequent expected-utility analysis will then depend critically on this exogenously-imposed prior information. In this sense, structural or deep uncertainty is potentially much more of a driving force than discounting or risk *per se* for cost-benefit applications of expected utility theory to open-ended situations with unlimited exposure. For such situations where there do not exist prior limits on damages (like climate change from greenhouse warming), expected present discounted utility analysis of costs and benefits is likely to be dominated by considerations and concepts related more to catastrophe insurance than to the consumption-smoothing consequences of long-term discounting at one or another particular interest rate.

2 How Bad Could It Possibly Get?

Let C be consumption. Suppose the existence of a representative agent having a smooth utility function $U(C)$ with $U'(C) > 0$ and $U''(C) < 0$. The elasticity of marginal utility

$$\gamma(C) \equiv -\frac{\frac{dU'(C)}{U'(C)}}{\frac{dC}{C}} = -\frac{CU''(C)}{U'(C)} \tag{1}$$

is a standard (local) measure of utility curvature (as a function of C). Another name for $\gamma(C)$ is the *coefficient of relative risk aversion*, because it quantifies the degree to which an agent dislikes uncertainty in proportionate consumption changes. Define

$$\eta \equiv \inf_{C>0} \gamma(C), \tag{2}$$

and say that a utility function displays *strict relative risk aversion* if $\eta > 0$.

In this paper I am concerned with the economics of catastrophic situations where C might be disastrously small with tiny probability. A key role in determining the expected-utility properties of such potentially-catastrophic situations will be played by the behavior of $\gamma(C)$ in the region where C is very small. Without loss of generality to anything of substance, it is convenient for ease of exposition in this paper to assume straightaway the isoelastic or CRRA (constant relative risk aversion) power utility function whose marginal utility is

$$U'(C) = C^{-\eta} \tag{3}$$

for some $\eta > 0$, which means that

$$U(C) = \frac{C^{1-\eta}}{1-\eta} \tag{4}$$

for $\eta \neq 1$ and $U(C) = \ln C$ for $\eta = 1$. The more general situation of a utility function where (2) holds in place of (3) involves more elaborate notation but has essentially identical consequences and conclusions.

In the model of this paper there are just two periods, the present and the future. (For applications to climate change, the future would be one or two centuries hence.) Instead of working directly with C, in this paper it is analytically more convenient to work with $\ln C$. If present consumption is normalized to unity, then the growth rate of consumption between the two periods is

$$Y \equiv \ln C. \tag{5}$$

In this model, Y is a random variable capturing *all* uncertainty that influences future values of $\ln C$. For the purposes of this paper Y includes not just economic growth narrowly defined. Much more essential to this paper's application of the model, Y includes the consumption-equivalent damages of adverse climate change. This paper is mostly concerned with the small probability of a large negative realization of Y. (Note: here the "bad" tail of Y is its *left* tail.) The "stochastic discount factor" or "pricing kernel" is an expression of the form

$$M(C) \;=\; \beta \frac{U'(C)}{U'(1)} \tag{6}$$

for discount factor β $(0 < \beta \leq 1)$. The amount of present consumption that the agent would give up in the present period to obtain one extra sure unit of consumption in the future period is here

$$E[M] \;=\; \beta E[\exp(-\eta Y)], \tag{7}$$

which is a kind of shadow price for discounting future costs and benefits in project analysis.

If the random variable Y has PDF $f(y)$, then (7) can be written as

$$E[M] \;=\; \beta \int_{-\infty}^{\infty} e^{-\eta y} f(y)\, dy, \tag{8}$$

which means that $E[M]$ is essentially the Laplace transform or moment-generating function of $f(y)$. This is helpful because the properties of the expected stochastic discount factor are the same as the properties of the moment-generating function of a probability distribution, about which a great deal is already understood. For example, if $Y \sim N(\mu, s^2)$ then plugging the usual formula for the expectation of a lognormal random variable into (8) gives the familiar expression

$$E[M] \;=\; \exp\left(-\delta - \eta\mu + \frac{1}{2}\eta^2 s^2\right), \tag{9}$$

where $\delta = -\ln \beta$ is the instantaneous rate of pure time preference. Expression (9) shows up in innumerable asset-pricing Euler-equation applications as the expected value of the stochastic discount factor or pricing kernel when consumption is lognormally distributed. Equation (9) is also the basis of the well-known generalized-Ramsey formula for the riskfree interest rate

$$r^f \;=\; \delta + \eta\mu - \frac{1}{2}\eta^2 s^2, \tag{10}$$

which (in its deterministic form, for the special case $s = 0$) plays a key role in recent debates about what social interest rate to use for intergenerational cost-benefit discounting of policies to mitigate GHG emissions. The intergenerational-discounting debate has mainly revolved around choosing "ethical" values of the rate of pure time preference δ, but this paper will demonstrate that for any $\eta > 0$ the δ-effect on r^f in formula (10) becomes overshadowed by the effect of an uncertain scaling parameter s.

Although not phrased in this way, the existing literature already contains an example that can be interpreted as showing a sense in which expected-utility-maximizing agents are much more averse to structural uncertainty about the scaling-multiplier parameter s than they are to pure risk in the form of an equation like (10) with as-if-known s. The example used here to convey this basic idea is a relatively simple specification consisting of the workhorse isoelastic or constant-relative-risk-aversion (CRRA) utility function (3) along with familiar probability distributions: lognormal, Student-t, gamma.[3] The question will then arise at the end of this section of the paper whether the insight that structural uncertainty is potentially much more worrisome than pure risk is due here to the particular quirks of this relatively simple example or, alternatively, it represents a generic idea of significantly broader scope. The answer, given in the next section, is that the result is generic (or at least much broader than the example), and the relatively simple formulation of this section will at that point help to motivate the subsequent development of a more general theory of catastrophic change due to structural uncertainty about the true value of the relevant scale parameter.

Throughout this paper, the structural scale parameter controlling the tail spread of a probability distribution is the most critical unknown. For the super-simple example of this section, $Y \sim N(\mu, s^2)$ where the mean μ *is known* but the scale parameter s *is unknown*. In an extremely loose sense this unknown structural scale parameter s is a highly-stylized abstraction of the effect that is embodied in an uncertain climate-sensitivity-like amplifying multiplier which was discussed in the introduction. With this rough analogy in this super-simple example, $Y = A - B\Delta T$, where A and B are positive constants.

The structural uncertainty about s will be modeled here *as if* this scale parameter is a random variable (denoted S) whose distribution must be inferred by inductive reasoning from n observed data points. At a very high level of abstraction, these data points might be interpreted as outcomes from various economic-scientific studies very roughly akin to the eighteen climate-sensitivity studies discussed in the introduction. Let $\mathbf{y} = (y_1, ..., y_n)$ be a sample of n i.i.d. random draws from the data-generating process of the normal distribution

[3] An example with these particular functional forms leading to existence problems from indefinite expected-utility integrals blowing up was first articulated by Geweke (2001). It subsequently was rediscovered in a context of discounting by Weitzman (2007a), who further developed its meaning and implications for asset pricing. A similar point to the main theme of this paper (tail-thickening of posterior-predictive distributions under general conditions) was introduced earlier (than Geweke (2001) or Weitzman (2007a)) in the important pioneering contribution of Schwarz (1999), although for a different context than expected utility theory.

whose PDF is

$$h(y \mid s) = \frac{1}{\sqrt{2\pi s}} \exp\left(-\frac{(y-\mu)^2}{2s^2}\right). \tag{11}$$

The sample variance is

$$\nu \;\equiv\; \frac{1}{n}\sum_{j=1}^{n}(y_j-\mu)^2 \tag{12}$$

and the likelihood function for the random variable S here is

$$L(s;\mathbf{y}) \;\propto\; \frac{1}{s^n} \exp\left(-\frac{n\nu}{2s^2}\right). \tag{13}$$

In a Bayesian framework, deriving the agent's posterior distribution of S requires that some prior distribution be imposed on S. There is a standard tried-and-tested reference prior for this particular situation, which forms the Bayesian mirror image of classical linear-normal regression analysis. This traditional prior (which is explained in any textbook on Bayesian statistics) is a uniform distribution of $\ln S$ on $(0,\infty)$, meaning the prior PDF expressed in terms of S is

$$p_0(s) \;\propto\; \frac{1}{s}, \tag{14}$$

which represents an improper probability distribution because the density (14) does not have a positive normalizing constant that makes it integrate to one. The next section of the paper will work with a finite-positive-support generalization of (14) that constitutes a proper prior distribution (because proper convergence of such kind of indefinite integrals is a prerequisite for the analysis that will then follow).

The *precision* θ is commonly defined to be the reciprocal of the variance, so that here

$$\theta \;\equiv\; \frac{1}{s^2}. \tag{15}$$

The posterior probability density is proportional to the product of the prior $p_0(s)$ times the likelihood $L(s;\mathbf{y})$. When the change of variables (15) is plugged into the product of (14) times (13), the posterior becomes conveniently expressed in terms of θ by the gamma distribution

$$\Gamma(\theta) \;\propto\; \theta^{a-1}\exp(-b\theta), \tag{16}$$

where

$$a \;=\; \frac{n}{2} \tag{17}$$

and

$$b \;=\; \frac{n\nu}{2}. \tag{18}$$

The mean of the gamma distribution (16) is a/b while its variance is a/b^2, so that from (17) and (18),

$$E[\theta \mid \mathbf{y}] \;=\; \frac{1}{\nu} \tag{19}$$

and

$$V[\theta \mid \mathbf{y}] \;=\; \left(\frac{2}{\nu^2}\right)\frac{1}{n}. \tag{20}$$

After integrating out the precision from the conditional-normal distribution (11), the unconditional or marginal posterior-predictive PDF of the growth rate Y is

$$f(y) \;\propto\; \int_0^\infty \sqrt{\theta}\, \exp(-\,\theta(y-\mu)^2/2)\; \Gamma(\theta)\, d\theta\,. \tag{21}$$

Straightforward brute-force integration of (21) (for (16), (17), (18)) shows that $f(y)$ is the Student-t distribution with n degrees of freedom:

$$f(y) \;\propto\; \left(1+\frac{(y-\mu)^2}{n\,\nu}\right)^{-(n+1)/2}. \tag{22}$$

(Any Bayesian textbook shows that a normal with gamma precision becomes a Student-t). Note that, asymptotically, the limiting tail behavior of (22) is a fat-tailed power-law PDF whose exponent is $n+1$. In the next section of the paper, this power-law-tail result will be shown to hold for a generalization of the example in this section.

When the posterior-predictive distribution of Y is (22) (from s being unknown), then (8) becomes

$$E[M] \;=\; +\infty, \tag{23}$$

because the moment-generating function of a Student-t distribution is infinite. Something quite extraordinary seems to be happening here, which is crying out for further explanation! Thousands of applications of expected-utility theory in thousands of articles and books are based on formulas like (9) or (10). Yet as soon as it is acknowledged that s is unknown and its value in formula (9) or (10) must instead be inferred as if from a data sample that is *arbitrarily large* (but finite), expected marginal utility explodes in (23). The question then naturally arises: what is expected-utility theory trying to tell us when its conclusions for a host of important applications are so sensitive to merely recognizing that the distribution implied by the normal *conditioned on finite realized data* is the Student-t?

I want to emphasize as emphatically as I can at this relatively early stage of the paper that the problem (both here *and throughout the rest of the paper*) is *not* that the illegitimate symbol $+\infty$ appears in formula (23), thereby temptingly offering a (illusory) way out of the dilemma that $E[M] \to +\infty$ by somehow discrediting this application of expected-utility theory on the narrow grounds that infinities are simply not allowed in a legitimate theory of choice under uncertainty. It is easy to put arbitrary bounds on utility functions, or to arbitrarily truncate probability distributions, or to introduce *ad hoc* priors that are arbitrarily cut off or otherwise severely dampened. Introducing any of these changes formally causes the symbol $+\infty$ to be replaced by an arbitrarily-large but finite number. (Indeed, the escape route of effectively placing arbitrary finite positive supports on θ or S is taken in the next section.) However, getting rid of the infinity symbol does not in any way eliminate (or even marginalize) the underlying problem, which will then come back to haunt in the form of an arbitrarily large expected stochastic discount factor, whose exact value will now depend hypersensitively upon obscure bounds, truncations, severely-dampened or cut-off prior PDFs, or whatever other tricks (of whose arbitrary meaning one has no idea) have been used to banish the $+\infty$ symbol. One can easily get rid of the $+\infty$ in formula (23) but one cannot banish the underlying economic problem that expected stochastic discount factors – which lie at the heart of cost-benefit discounting, asset-pricing theory, and many other important applications of expected utility theory – can become arbitrarily large just by making unobjectionable statistical inferences about limiting tail behavior. The take-away message here is that reasonable attempts to constrict the fatness or length of the "bad" tail of the multiplier still leaves us with uncomfortably big numbers.

The core underlying problem is that it is impossible to learn limiting tail behavior from finite data, and seemingly thin-tailed probability distributions (like the normal), which are actually only thin-tailed *conditional on known structural parameters* of the model, become tail-fattened (like the Student-t) after integrating out the scale-parameter uncertainty. There is no clean way to eliminate this core issue, and of necessity it must influence any utility function that is sensitive to low values of consumption. It is important to realize that the unboundedness potential for $E[M]$ in (23) comes essentially from the fattened "bad" left tail of the posterior-predictive distribution of Y given by (22), and that such tail-fattening must inevitably occur as a result of taking probability distributions over probability distributions of scale parameters. Utility isoelasticity *per se* is inessential to the reasoning (although it makes the argument easier to understand), because the expected stochastic discount factor $E[M]$ is $+\infty$ for *any* relatively-risk-averse utility function satisfying the limiting curvature requirement: $\lim_{C \to 0} [-CU''(C)/U'(C)] > 0$.

The Student-t child posterior-predictive density from a large number of observations looks almost exactly like its bell-shaped normal parent except that the probabilities are somewhat more stretched out, making the tails appear relatively fatter at the expense of a slightly-flatter center. Intuitively, a normal density "becomes" a Student-t from a tail-fattening spreading-apart of probabilities caused by the variance of the normal having itself a (inverted gamma) probability distribution. There is then no surprise from expected utility theory that people are more averse qualitatively to a relatively fat-tailed Student-t posterior-predictive child distribution than they are to the relatively thin-tailed normal parent which begets it. A perhaps more surprising consequence of expected utility theory is the *quantitative strength* of this endogenously-derived aversion to the effects of unknown tail-structure. The story behind this quantitative strength is that fattened-posterior bad tails represent structural or deep uncertainty about the possibility of rare unlimited-exposure disasters that – using colorful language – "scare" any agent having a utility function with strict relative risk aversion. With (2) holding, the only way to contain this scary "Student-t explosion" is to exclude it *a priori* by prior information via imposing some kind or another of a restrictive or strongly-dampening prior distribution.

The obvious next issue to be investigated is the extent to which this particular example generalizes. It turns out that such a scary fattened-posterior-tail effect essentially holds for *any* probability distribution characterized by having an uncertain scale parameter (not just the normal) and for *any* utility function having strict relative risk aversion (not just CRRA). When indeterminateness is compounded because probability distributions themselves have probability distributions, then posterior-predictive tails must inevitably become fattened, with potentially strong consequences for applications of expected utility theory.

The next section formalizes the idea that structural parameter uncertainty in a much more general model inevitably leads to a tail-fattened posterior-predictive distribution of growth rates that can cause expected marginal utility to blow up. It will turn out that there is a rigorous sense in which containing this explosion of the expected stochastic discount factor necessitates an unavoidable dependence upon some kind or another of exogenously-imposed subjective prior beliefs in the form of a restrictive prior distribution, which reflects critical "extra" knowledge of limited exposure that cannot be transmitted by finite data alone.

3 The General Model

In order to focus sharply on structural parameter uncertainty, the model of this section is patterned closely as a generalization of last section's example and is very sparse. To create families of probability distributions that are simultaneously fairly general and analytically tractable, the following generating mechanism is employed. Let Z represent a random variable normalized to have mean zero and variance one. Let $\phi(z)$ be *any* continuous PDF with

$$\int_{-\infty}^{\infty} z\,\phi(z)\,dz \;=\; 0 \tag{24}$$

and with

$$\int_{-\infty}^{\infty} z^2\,\phi(z)\,dz \;=\; 1. \tag{25}$$

It should be noted that the PDF $\phi(z)$ is allowed to be extremely general. For example, the distribution of Z might have finite support (which means that unbounded catastrophes will be ruled out *conditional* on the value of the finite lower support being known), or it might have unbounded range (like the normal, which will allow unbounded catastrophes to occur but assigns them a thin bad tail *conditional* on the standard deviation being known). Aside from continuity, the *only* restrictions placed on $\phi(z)$ are the two weak regularity conditions that $\phi(0) > 0$ and

$$\int_{-\infty}^{\infty} \exp(-\alpha z)\,\phi(z)\,dz \;<\; \infty \tag{26}$$

for all $\alpha > 0$, which is automatically satisfied if Z has finite support.

With $s > 0$ given, make the change of variable of the linear form $y = sz + \mu$, which implies that the conditional PDF of y is

$$h(y \mid s) \;=\; \frac{1}{s}\,\phi\left(\frac{y-\mu}{s}\right), \tag{27}$$

where s and μ are structural parameters having the interpretations

$$\mu \;=\; E[Y \mid s] \;=\; \int_{-\infty}^{\infty} y\,h(y \mid s)\;dy \tag{28}$$

and

$$s^2 \;=\; V[Y \mid s] \;=\; \int_{-\infty}^{\infty} (y-\mu)^2\;h(y \mid s)\;dy. \tag{29}$$

For the ideas of this paper what matters most is structural uncertainty about the scale parameter controlling the tail spread of a probability distribution, which is the most critical unknown. The scale parameter s may be conceptualized extremely loosely as being a highly-stylized abstract generalization of a climate-sensitivity-like amplifying or scaling multiplier. In this crude analogy, $Z = \Delta \ln CO_2 / \ln 2$, $\Delta T = SZ$, $Y = A - B\Delta T$. Without significant loss of generality, it is assumed for ease of exposition that in (27) the mean μ *is known*, while the standard-deviation *scale parameter* s *is unknown*. The case where μ and s are *both* unknown involves more intricate notation but otherwise gives essentially identical results.

The point of departure here is that the conditional PDF of growth rates $h(y \mid s)$ given by (27) is known by the agent and, while the true value of s is unknown, the situation is *as if* there are available some finite number of i.i.d. observations on which to base an estimate of s by some process of inductive reasoning. Suppose that the agent has observed the random sample $\mathbf{y} = (y_1, ..., y_n)$ of growth-rate data realizations from n independent draws of the distribution $h(y \mid s)$ defined by (27) for some unknown fixed value of s. An example relevant to this paper is where the sample space represents the outcomes of various economic-scientific studies and the data $\mathbf{y} = (y_1, ..., y_n)$ are interpreted at a very high level of abstraction as the findings of n such studies so that n measures inductive knowledge.

From (27) the relevant likelihood function of s is

$$L(s; \mathbf{y}) \propto \prod_{j=1}^{n} h(y_j \mid s). \tag{30}$$

The prior PDF of s is taken to be a generalization of (14) of the form

$$p_0(s) \propto s^{-k} \tag{31}$$

for some number k. As k can be chosen to be *arbitrarily large*, it should be appreciated that the prior distribution (31) can be made to place *arbitrarily small* prior probability weight on big values of s. It should also be appreciated that *any invariant prior must be of the form* (31). Invariance (discussed in the Bayesian-statistical literature) is considered desirable as a description of a "noninformative" reference prior that favors no particular value of s over any other. For such a noninformative reference prior, it seems reasonable to impose a condition of scale invariance that might be justified by the following logic. If the action taken in a decision problem should not depend upon the unit of measurement used, then a plausible principle of rational invariance of a scale parameter would require that

$$p_0(s) \propto p_0(\lambda s), \tag{32}$$

and the only way that (32) can then hold for all $\lambda > 0$ and all $s > 0$ is when the prior is of the form (31).

In order to turn the improper prior (31) into a proper distribution that allows integrals to converge, arbitrary finite positive supports (ω, Ω) are imposed where $\omega > 0$ is arbitrarily small and $\Omega < \infty$ is arbitrarily large. Thus, when $k = 1$ then $p_0(s) = 1/(\ln\Omega - \ln\omega)s$, while with $k \neq 1$, (31) becomes

$$p_0(s) \;=\; \left[\frac{k-1}{\omega^{1-k} - \Omega^{1-k}}\right] s^{-k} \tag{33}$$

holding for $\omega < s < \Omega$ and $p_0(s) = 0$ elsewhere.

The posterior probability density $p_n(s\,|y)$ is proportional to the product of the prior $p_0(s)$ from (33) times the likelihood $L(s;\mathbf{y})$ from (30), which here yields

$$p_n(s \mid \mathbf{y}) \;\propto\; p_0(s) \prod_{j=1}^{n} h(y_j \mid s) \tag{34}$$

for $\omega < s < \Omega$ and $p_n(s\,|y) = 0$ elsewhere.

Integrating out the agent's uncertainty about s described by the probability density (34), the unconditional or marginal posterior-predictive density of the growth-rate random variable Y is

$$f(y \mid \Omega) \;=\; \int_{\omega}^{\Omega} h(y \mid s)\; p_n(s \mid \mathbf{y})\; ds, \tag{35}$$

and (8) then becomes

$$E[M \mid \Omega] \;=\; \beta \int_{-\infty}^{\infty} e^{-\eta y}\, f(y \mid \Omega)\, dy. \tag{36}$$

4 The Main Result

Section 2 gave a "Student-t example" where $E[M] \to +\infty$. The primary finding of this paper is that $E[M] \to +\infty$ actually holds under quite general structural uncertainty about the unknown scaling parameter.

Theorem 1 *For any given n and k,*

$$\lim_{\Omega \to \infty} E[M \mid \Omega] \;=\; +\infty. \tag{37}$$

Proof. Combining (27) with (33)-(36) gives

$$E[M \mid \Omega] \propto \int_{\omega}^{\Omega} \frac{1}{s^{k+n+1}} \prod_{j=1}^{n} \phi\left(\frac{y_j - \mu}{s}\right) \left[\int_{-\infty}^{\infty} e^{-\eta y}\, \phi\left(\frac{y-\mu}{s}\right) dy\right] ds. \tag{38}$$

Make the change of variable $z = (y - \mu)/s$ and reverse the order of integration to rewrite (38) as

$$E[M \mid \Omega] \propto \int_{-\infty}^{\infty} \phi(z) \left[\int_{\omega}^{\Omega} e^{-\eta z s} \frac{1}{s^{k+n+1}} \prod_{j=1}^{n} \phi\left(\frac{y_j - \mu}{s}\right) ds \right] dz. \tag{39}$$

Pick any value of z' for which simultaneously $z' < 0$ and $\phi(z') > 0$, and note here that

$$\lim_{\Omega\to\infty} \left[\int_{\omega}^{\Omega} \prod_{j=1}^{n} \phi\left(\frac{y_j - \mu}{s}\right) e^{-\eta z' s} \frac{1}{s^{k+n+1}} ds \right] \geq \phi(0)^n \lim_{\Omega\to\infty} \left[e^{-\eta z' \Omega} \frac{1}{\Omega^{k+n+1}} \right] = +\infty, \tag{40}$$

which, in conjunction with (39), concludes the proof. ■

In the model of this paper, people are averse to two types of indeterminacy – not knowing y *given* s, and not knowing s. Theorem 1 can be interpreted as providing a rigorous sense in which aversion to structural or deep uncertainty (in the form of not knowing the scale parameter s) is potentially far greater than aversion to the pure risk *per se* of not knowing the realized value of the growth-rate random variable y (when s is known). The underlying logic behind the strong result of Theorem 1 is described in (40) by the limiting behavior of $E[M \mid \Omega]$ for large values of Ω. For any given n and k, the *probability* of a disaster declines *polynomially* in the scale Ω of the disaster from (40), while the marginal-utility *consequence* of a disaster (when there is strict relative risk aversion) increases *exponentially* in the scale Ω of the disaster. Irrespective of the original parent distribution, the effect of an uncertain scale parameter fattens the tails of the posterior-predictive child distribution so that it behaves asymptotically like a power-law distribution whose power coefficient is $n+k$. In this sense power-law tails need not be postulated, because they are essentially unavoidable in posterior-predictive distributions. No matter how many observations there are, the race between a polynomially-contracting power-law probability times an exponentially-expanding marginal utility is won in the limit every time by marginal utility, so long as there is strict relative risk aversion (which might be considered a minimal assumption).

The interpretation of Theorem 1 is sensitive to a behind-the-scene tug of war between pointwise-but-nonuniform limiting behavior in Ω and pointwise-but-nonuniform limiting behavior in n. To see more clearly how the issue of risk vs. uncertainty plays out in determining $E[M]$ under such form of pointwise-but-nonuniform convergence, suppose that, unbeknownst to the agent, the true value of s is s^*. Since the prior $p_0(s)$ defined by (33) assigns positive probability to an open neighborhood around s^* (provided only that $\omega < s^* < \Omega$), the imposed specification has sufficient regularity for large-sample likelihood dominance to cause strong (almost sure) convergence of the posterior distribution (34) of S to its true data-generating value $s = s^*$. This in turn means that the posterior-predictive distribution of growth rates (35) converges a.s. to its true data-generating distribution $h(y \mid s^*)$. The latter further implies that for *any given* $\Omega < \infty$ (which via (26) places an upper bound on expected marginal utility), in the limit as $n \to \infty$, risk is more important than structural uncertainty in the sense that as full structural knowledge is approached $E[M]$ converges strongly to its true value:

$$n \to \infty \quad \Longrightarrow \quad E[M \mid \Omega] \underset{a.s.}{\longrightarrow} \int_{-\infty}^{\infty} e^{-\eta y} \frac{1}{s^*} \phi\left(\frac{y-\mu}{s^*}\right) dy. \tag{41}$$

A conventional pure-risk application of expected utility theory essentially corresponds here to a situation where there is enough inductive knowledge to identify the structure because n is reasonably large relative to the limited-exposure bound Ω. However, it is critical here to note that in (41) the a.s. convergence of $E[M \mid \Omega]$ to its true value is *pointwise but not uniform* in n. No matter how much data-evidence exists – or even can be *imagined* to exist – Theorem 1 says that $E[M \mid \Omega]$ is always exceedingly sensitive to Ω. In this sense, structural uncertainty always has the potential to trump pure risk for situations of potentially-unlimited exposure where no plausible bound $\Omega < \infty$ can legitimately be imposed by prior knowledge. The dominant statistical-economic truth behind Theorem 1 is that no finite sample can assess probabilities or magnitudes of the most extreme disasters lurking in the distant tails of distributions – and the expected utility of a Bayesian agent with strictly-positive relative risk aversion will be driven to an arbitrarily large extent by this unavoidable feature of learning-inference *unless* potential exposure is limited by prior information. Theorem 1 can therefore be interpreted as implying a sense in which it is unnecessary to append to the theory of decision making under uncertainty an *ad hoc* extra postulate of "ambiguity aversion" – because expected utility theory itself already tells us a precise way in which the "ambiguity" of structural-parameter uncertainty can be especially important and why people may be much more averse to it under some circumstances of limited information than to pure risk with known structure.

There is a general point being made by Theorem 1 and a particular application to the economics of climate change. The general point is that Theorem 1 embodies a very strong form of a "generalized precautionary principle." From experience alone one cannot acquire sufficiently accurate information about the probabilities of bad-tail disasters to prevent the expected marginal utility of an extra sure unit of consumption from becoming unbounded for any utility function having strict relative risk aversion, thereby potentially dominating cost-benefit applications of expected utility theory. The underlying problem that Theorem 1 is illustrating concerns a fundamental limitation on the ability of empirical learning or inductive knowledge to shed light on extreme events. Even in a stationary world, it is not possible to learn enough about tail events from finite samples alone to make expected stochastic discount factors be other than highly sensitive to possibly-ruinous disasters. A risk-averse expected-utility maximizer is potentially scared by structural uncertainty to a degree that is potentially beyond what any finite amount of empirical information can overcome. Of course in reality people are not infinitely scared by fat tails – due, presumably, to some combination of "high enough" effective n or k with "low enough" effective Ω. The point here is primarily to be alerted to the potentially-overwhelming consequences (on expected-utility applications) of high-Ω low-$(n+k)$ scary situations, such as climate change.

The part of the distribution of possible outcomes that can most readily be learned (from inductive information that comes in a form as if conveyed by data) concerns the most-likely outcomes in the central body of the distribution because, from previous experience, past observations, plausible extrapolations, and perhaps even the law of large numbers, there may be at least some modicum of confidence in being able to construct a reasonable approximation of the central regions of the PDF. As we move towards probabilities in the periphery of the distribution, however, we are increasingly moving into the unknown territory of subjective uncertainty where our probability estimates of the probability distributions themselves becomes increasingly diffuse because the frequencies of rare events in the tails cannot be pinned down by previous experiences or past observations. Climate change generally and climate sensitivity specifically are prototype examples of this general principle, because we are trying to extrapolate inductive knowledge far outside the range of our limited experience. The upshot of this deep uncertainty about tail uncertainties is that the reduced-form probability distribution of Y (after integrating out the probabilities of tail probabilities) has a fat left tail. The exact fatness of this bad left tail depends not only upon how bad a disaster might materialize and with what probabilities, but also upon how imprecise are our probability estimates of the probabilities of these disasters.

The degree to which the kind of "generalized precautionary principle" embodied in Theorem 1 is relevant for a particular application must be decided on a case-by-case basis. It depends generally upon the extent to which prior Ω-knowledge (and k-knowledge) combine with inductive-posterior n-knowledge in a particular case to fatten the bad tail. In the particular application to the economics of climate change, where there is so obviously limited data and limited experience about the catastrophic reach of climate extremes (and where what we *do* know about nonlinear dynamic stochastic climate predictions produces here a long fat right tail of ΔT), to ignore or suppress the significance of rare tail disasters is to ignore or suppress what economic theory is telling us loudly and clearly is potentially the most important part of the analysis.

Global climate change unfolds over a time scale of centuries and, through the power of compound interest, a standard cost-benefit analysis of what to do now to mitigate GHGs is hugely sensitive to the discount rate that is postulated. This has produced some sharp disagreements among economists about what is an "ethical" value of the rate of pure time preference δ to use for intergenerational discounting in the deterministic version ($s = 0$) of the Ramsey equation (10) that forms the analytical backbone for most analyses of the economics of climate change.[4] For the model of this paper, which is based on structural uncertainty, arguments about what value of δ to use in equation (9) or (10) translate into arguments about what value of β to use in the model's structural-uncertainty generalization of the Ramsey formula, which is the expectation formula (36). (A zero rate of pure time preference $\delta = 0$ in (10) corresponds to $\beta = 1$ in (36).) In this connection, Theorem 1 is saying that no matter what value of β is selected (so long as $\beta > 0$, which is equivalent to $\delta < \infty$), the outcome of any cost-benefit analysis of what to do now to mitigate GHGs is always potentially driven by structural uncertainty about s.

The unknown scale parameter s (whose uncertainty potentially drives the economic analysis) is an abstract generalized version of an amplifying multiplier whose role is very crudely analogous to the role of open-ended climate sensitivity. Therefore, it is no accident and no surprise that the PDF of climate sensitivity based on inductive data and studies has a fat tail. Expected utility theory is telling us here analytically that the debate about discounting may be secondary to a debate about the open-ended catastrophic reach of climate disasters. While it is always fair game to challenge the assumptions of a model, when theory provides a generic result (like "free trade is Pareto optimal" or "steady growth eventually outstrips one-time change") the burden of proof is commonly taken as residing with whomever wants to over-rule the theorem in a particular application. The take-away message here is that the burden of proof in the economics of climate change is presumptively upon whomever wants to model or conceptualize the expected present discounted utility of feasible trajectories under greenhouse warming *without* having structural uncertainty tending to matter much more than discounting or pure risk *per se*. Such a middle-of-the-distribution modeler needs to explain why the inescapably-fattened tails of the posterior-predictive distribution (for which the fat bad tail represents rare disasters under uncertain structure from an unknown scaling parameter) is *not* the primary focus of attention and does *not* play the decisive role in the

[4] While this contentious intergenerational-discounting issue has long existed (see, e.g., the various essays in Portney and Weyant (1999)), it was elevated to recent prominence by publication of the controversial *Stern Review of the Economics of Climate Change* (2007). The *Review* argued for a base case of preference-parameter values $\delta \approx 0$ and $\eta \approx 1$, on which its strong conclusions depended analytically. Alternative views of intergenerational discounting than *Stern*'s were provided in, e.g., Dasgupta (2007), Nordhaus (2007), and Weitzman (2007b). The latter account also contains a heuristic exposition of the contents of this paper, as well as giving *Stern* some credit for emphasizing the great uncertainties associated with climate change.

analysis.

5 Concluding Discussion

A common reaction to the likes of Theorem 1 is to acknowledge its mathematical logic but to wonder how it is to be used in practice for deciding what to do without effectively blocking all progress on account of the theoretical possibility of severe downside risks. After all, horror stories about theoretically-possible speculative disasters can be told for many situations without this aspect necessarily paralyzing decision-making by freezing the status quo. Dismal Theorem 1 gives an almost lexicographic or binary ordering, in the limit of which agents will pay virtually any price to eliminate deep uncertainty – an idea which is hard to wrap one's mind around. Is this an economics version of a kind of "impossibility theorem" that says there are limits to quantitative analysis when dealing with complex chaos-prone dynamic systems where knowledge is incomplete and deep uncertainty predominates? And in that case, what are we supposed to put in the place where ordinary cost-benefit calculations used to be before? Even if it were true that the dismal Theorem 1 represents a valid economic-statistical precautionary principle which theoretically dominates decision making, would not applying this "generalized precautionary principle" freeze all progress if taken too literally? What are we to do in practice with a situation of deep structural uncertainty? I don't know the answers to these kinds of questions. But I also don't think such questions can be ignored or even evaded in potentially-catastrophic situations where there is unlimited exposure to damages.

In seeking enlightenment about what it might mean to apply Theorem 1 to the economics of climate change, I don't believe the best route is via an abstract general discussion of tail probabilities (or theoretical considerations of when limited exposure to damages might, or might not, apply). Rather, I think it more useful to illustrate concretely what is involved by contrasting the specific situation regarding possible climate-change-induced environmental catastrophes with possible environmental catastrophes unleashed by another specific situation. The example I choose here is the widespread cultivation of crops based upon bioengineered genetically-modified organisms (GMOs).

At first glance, the two situations might appear casually similar. In both cases, there is deep unease about human-induced tinkering with the natural environment, which can generate frightening tales of a world turned upside down. Suppose for the sake of argument that in the case of GMOs the overarching fear of disaster concerns the possibility that widespread cultivation of so-called "Frankenfood" could somehow allow bioengineered genes to escape into the wild and wreak havoc with delicate ecosystems and native populations (including perhaps humans) fine-tuned by millions of years of natural selection. I think that the potential for environmental disaster with Frankenfood is quite different from the potential for environmental disaster with catastrophic climate change along the lines of the following reasoning.

Theorem 1 and the subsequent discussion of non-uniform convergence in Ω and n imply that deep uncertainty about the unknown scale of a disaster has the potential to dominate expected-utility cost-benefit calculations when the scope Ω of such disasters is large relative to the amount of prior-plus-inductive knowledge $k+n$. I think that in the case of Frankenfood interfering with wild organisms that have evolved by natural selection, there *is* a basic underlying principle that plausibly limits the extent of catastrophic jumping of artificial DNA from cultivars to landraces. After all, nature herself has already tried endless combinations of mutated DNA and genes over countless millions of years and what has evolved in the fierce battle for survival is only an infinitesimal subset of the very fittest permutations. In this regard there exists a basic prior argument making it fundamentally implausible that Frankenfood artificially selected for traits that humans find desirable will compete with or genetically alter the wild types that nature has selected via Darwinian survival of the fittest. Wild types have already experienced genetic mutations akin to human-induced artificial modifications – and these potential modifications have already been shown not to have survival value in the wild. I think that analogous arguments may also apply for invasive "superweeds," which (at least so far) represent a minor cultivation problem and have not displayed any ability to displace landraces. Besides all this, and importantly, safeguards in the form of so-called "terminator genes" can be inserted into the DNA of GMOs, which directly prevents GMO genes from reproducing themselves.

Contrast the above discussion about the limits of disaster (and the limits of a precautionary principle) in the case of Frankenfood with the lack of any such prior limitation on the magnitude of disasters possible from climate change. The climate-change "experiment," whose eventual outcome we are trying to infer now, concerns the planet's response to a geologically-instantaneous exogenous injection of GHGs. This planetary experiment of an exogenous injection of this much GHG this fast is probably unprecedented in Earth's history – even stretching back billions of years. Can anyone honestly say now from very limited information what are reasonable upper bounds on the eventual degree of global warming or climate change that we are currently trying to infer will be the outcome of such a first-ever planetary experiment? What we DO know about climate science and extreme tail probabilities is that chaotic dynamics cannot be ruled out and eighteen of the best current models of climate sensitivity are estimating on average that $P[S > 6°C] \approx 5\%$. To my mind this open-ended aspect makes GHG-induced global climate change vastly more worrisome than global cultivation of Frankenfood. I think this example shows that it is possible to make meaningful distinctions among situations where Theorem 1 might conceivably apply. While we cannot rule out a biotech disaster, I would say on the basis of this assessment that it seems *very very* unlikely (or maybe even *very very very* unlikely), whereas a climate disaster seems "only" *very* unlikely. In the language of this paper, biotech looks like a low-Ω, high-k, high-n situation whereas climate change feels much more like a high-Ω, low-k, low-n situation. I don't think there is a smoking gun in the biotech scenario – or in most other catastrophe scenarios – quite like average expert assessment in climate change being that $P[S > 8°\mathrm{C}] \approx 2\%$.

Many situations, especially those involving pure risk, have prior limited-exposure-like bounds on the possible damages that might materialize, for which the theory of this paper may be less relevant or perhaps even have no relevance. But a few real-world situations have potentially *unlimited* exposure due to structural uncertainty about their potentially open-ended catastrophic reach. This paper shows that the expected utility analysis of those relatively few deep-uncertainty situations with potentially catastrophic reach – like climate change – can give very different conclusions than what might emerge from a typical cost-benefit analysis of a more ordinary limited-exposure situation. The conclusions can even be very different from what would come out of a Monte Carlo simulation with a discrete grid or with a finite number of runs of a model whose core structure resembles the model of this paper. The theoretical outcome of Theorem 1 of this paper can be approached by a Monte Carlo simulation only as a double limit where the grid size and the number of runs both go to infinity.

Sampling based upon standard Monte Carlo simulations of any existing integrated-assessment model of the economics of climate change is liable to give a totally misleading picture of the expected-utility consequences of alternative GHG-mitigation policies.[5] The core underlying problem is that while it might be true in *expectations* that utility-equivalent damages of climate change are enormous, when chasing a fat tail this will not be true for the overwhelming bulk of Monte Carlo *realizations*. To see this most clearly in a crisp thought experiment, imagine what would happen to the simple stripped-down model of this paper in the hands of a Monte Carlo simulator. A finite grid may not reveal true expected utility in simulations of this model (even in the limit of an infinite number of runs) because the most extreme damages in the fattened tails will have been truncation-compressed into being evaluated at a single grid-point. A finite sample of simulations may not reveal true expected

[5] I am grateful to Richard Carson for suggesting the inclusion of an explicit discussion of why a Monte Carlo simulation may fail to account fully for the implications of uncertain large consequences with small probabilities.

utility in this model either (even in the limit of an infinite grid) because the limited sample may not be able to go deep enough into the fattened tails where the most extreme damages are. Nor will typical sensitivity analysis of Monte Carlo simulations necessarily penetrate sufficiently far into the fattened-tail region to accurately represent disastrous damages. Instead of the existing emphasis on estimating or simulating the economic impacts of what are effectively the most plausible risk-like climate-change scenarios, to at least compensate partially for finite-sample bias the model of this paper calls for a dramatic *over*sampling (relative to probability of occurrence) of those stratified climate-change scenarios associated with the most adverse imaginable economic impacts. With limited sampling resources in big models, Monte Carlo analysis could be used much more creatively – not to defend a specific policy result, but to experiment seriously with what happens to the outcomes in the limit as the grid size and number of runs simultaneously increase. Of course this emphasis on sampling climate-change scenarios in proportion to marginal-utility-weighted probabilities of occurrence forces us to put marginal-utility weights on catastrophes (as well as subjective probabilities) – but that is the price we must be willing to pay for having a genuine economic analysis of potentially-catastrophic climate change.

In situations of potentially unlimited damage exposure – like climate change – a reframing of the focus of economic analysis might be called for, with more emphasis on a better treatment of the worst-case tail extremes (and what might be done about them, at what cost) relative to refining the calibration of merely-likely outcomes or arguing over discount rates. Perhaps it might be possible to reason somewhat by analogy with insurance for extreme events and to compare the cost to the world economy of buying an insurance policy going some way towards lowering the extreme high temperatures with, say, a homeowner's cost of buying fire insurance or a young adult's cost of buying life insurance. On a U.S. national level, rough comparisons might be made with the potentially-huge payoffs, small probabilities, and large costs involved in building anti-ballistic missile shields or eliminating hostile dictatorships that could be harboring weapons of mass destruction. A crude natural metric for calibrating cost estimations of climate-change environmental-insurance policies might be that the U.S. already spends approximately 3% of national income on the overall cost of a clean environment.[6]

The point is that economic analysis is not completely helpless in the presence of deep structural uncertainty. The analysis is more frustrating, more subjective, and less conclusive-looking because it requires some form of speculation (masquerading as an "assessment") about the extreme bad-tail probabilities and utilities. Compared with the thin-tailed case,

[6]U.S. Environmental Protection Agency (1990), executive summary projections for 2000 updated and extrapolated by me to 2007.

cost-benefit analysis of potential catastrophes is inclined to favor paying a *lot* more attention to learning *now* how fat the bad tail might be and – if the tail is discovered to be too thick for comfort after the learning process – is a *lot* more prone to investing in mitigation measures to slim it down. This slimming-down of overweight tails is likely to be a perennial theme in the economic analysis of catastrophes. The key economic issues here are: what is the cost of such a tail-slimming weight-loss program and how much of the bad fat does it remove from the overweight tail? A clear implication of this paper is that more research effort targeted at describing and estimating the risk-like central tendencies of what we already know fairly well is largely wasted – relative to understanding even slightly better the deep uncertainty (which potentially dominates the economic analysis) about what is in the bad fat tail and what to do about it. This being said, the bind we find ourselves in now on climate change starts from a high-Ω, low-k situation to begin with, and has convergence of inductive knowledge towards resolving the deep uncertainties being extremely slow in n relative to the lags and irreversibilities from *not* acting before structure is fully identified.

It is painfully apparent that the likes of Theorem 1 makes economic analysis trickier and more open-ended in the presence of deep structural uncertainty. The economics of fat-tailed catastrophes raises difficult conceptual issues which cause the analysis to appear less scientifically conclusive and to appear more contentiously subjective than what comes out of the analysis of thin-tailed situations. But if that is the way things are with fat tails, then that is the way things are – and it is a fact to be lived with rather than a fact to be evaded just because it looks less scientifically-objective in cost-benefit applications.

The contribution of this paper is to phrase precisely and to prove rigorously a basic theoretical principle that must hold under a standard assumption of strict positive relative risk aversion. In principle, what might be called the "catastrophe-insurance aspect" of a fat-tailed unlimited-exposure situation dominates the social-discounting aspect, the pure-risk aspect, or the consumption-smoothing aspect. Even if this principle in and of itself does not provide an easy answer to questions about how much catastrophe insurance to buy (or even an easy answer in practical terms to the question of what exactly *is* catastrophe insurance buying), I believe it still might provide a useful way of framing the economic analysis of catastrophes.

References

[1] **Dasgupta, Partha.** "Commentary: the Stern Review's Economics of Climate Change." *National Institute Economic Review*, 2007, 199, pp. 4-7.

[2] **Geweke, John.** "A note on some limitations of CRRA utility." *Economics Letters*, 2001, 71, pp. 341-345.

[3] **IPCC4.** *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press, 2007 (available online at http://www.ipcc.ch).

[4] **Nordhaus, William D.** "The Stern Review on the Economics of Climate Change." *Journal of Economic Literature*, forthcoming September 2007.

[5] **Parson, Edward A.** "The Big One: A Review of Richard Posner's *Catastrophe: Risk and Response*." *Journal of Economic Literature*, 2007, XLV (March), pp. 147-164.

[6] **Portney, Paul R. and Weyant, John P, eds.** *Discounting and Intergenerational Equity*. Washington, DC: Resources for the Future, 1999.

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

[8] **Schwarz, Michael.** "Decision Making Under Extreme Uncertainty." Stanford University Graduate School of Business: Ph.D. Dissertation, 1999.

[9] **Stern, Nicholas et al.** *The Economics of Climate Change*. Cambridge University Press, 2007.

[10] **Weitzman, Martin L.** "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, forthcoming September 2007a.

[11] **Weitzman, Martin L.** "The Stern Review of the Economics of Climate Change." *Journal of Economic Literature*, forthcoming September 2007b.

[12] **U.S. Environmental Protection Agency.** *Environmental Investments: The Cost of a Clean Environment.* Washington, DC: U.S. Government Printing Office, 1990.