Economic Policy in the Face of Severe Tail Events

William Nordhaus
May 17, 2010

DRAFT FOR DISCUSSION PURPOSES

Abstract

From time to time, something occurs which is outside the range of normal expectations. We will call these "tail events" in the sense that they are way out the tail of a probability distribution. I consider the question of the implications of tail events for economic policy and climate-change economics. This issue has been analyzed by Martin Weitzman, who proposed a Dismal Theorem. The general idea is that, under limited conditions concerning the structure of uncertainty and risk aversion, society has an indefinitely large expected loss from high-consequence, low-probability events. Under such conditions, standard economic tools such as cost-benefit analysis cannot be applied. The present study is intended to put the Dismal Theorem in context and examine the range of its relevance, with an application to catastrophic climate change. I conclude that tail events are sometimes of extreme importance, and we must be extremely careful to include them in situations of deep uncertainty. However, we conclude that no loaded gun of strong tail dominance has been uncovered to date.

I. Climate Policy in the Age of Tail Events

In an earlier era, climate policy was a straightforward exercise in weighing costs and benefits of mitigation in light of the economic costs of reducing emissions and the lower economic damages from reduced concentrations of greenhouse gases (GHGs). Most analyses in this framework called for modest near-term reductions in emissions, followed by increasingly tight controls in the decades to come. In this view, the science and policy could improve sharpen our estimates as there was no great urgency in terms of either the damages of sharp discontinuities.

Three developments in the natural and social sciences have upended the earlier relatively relaxed stance on climate policy. The first was the argument in *The Stern Review* that societies must take a longer view than was customary. This approach argued for a near-zero rate of pure time preference because of the very long duration of impacts of today's emissions.

A second development was the change in the view of climate change from one of gradual and smooth changes in the impacts to one with potentially abrupt climate change and sharp irreversible tipping points. This view emphasized the importance of potential thresholds or potentially catastrophic impacts. Important examples that have been discussed are reversal of the Atlantic thermohaline circulation, disintegration of the Greenland and West Antarctic Ice Sheets, shifts in monsoonal patterns, and die-off of the Amazon rain forests. Little economic analysis of these tipping points has taken place.

The third change has been the increasing concern with the potential of "tail events" in the impacts of climate change. This paper deals with the third of these issues. I focus in particular on the implications of the combination of outcomes that are potentially catastrophic in nature and have "fat tails." The combination of these two circumstances may lead to situations in which our standard analyses need to be modified or even break down. In the extreme case, the combination of fat tails and unlimited exposure implies that the expected loss from certain risks such as climate change is infinite.

The plan of the present study is as follows. We begin in the first section with a discussion of tail events along with a definition of tail dominance. We provide an overview of the basic analysis underlying Weitzman's Dismal Theorem, and then provide a simplified heuristic example in which the basic structure is easily seen. We put the analysis in the context of policy decisions and note that the usual version of the Dismal Theorem actually contains no policy decisions.

The subsequent section sketches the critical assumptions underlying strong tail dominance as exemplified by the Dismal Theorem. It shows that two central assumptions are unboundedness of both the uncertain variable and the marginal utility of consumption as consumption approaches zero.

The final section explores in greater depth the question of tail dominance in the important case of climate change. This example is motivated both because of its importance in current policy discussions and because it is the policy example that Weitzman explores in his analysis. We first examine the question of whether we can establish upper bounds on the critical uncertain variable, the temperature sensitivity coefficient (TSC) that Weitzman analyzes. We indicate that using paleoclimatic data, a secure although relatively high upper bound can be determined. We also explore the distribution of consumption declines using the analog of consumption "disasters" in economic history. Using two different databases, we find that consumption disasters appear not to have a tail that is sufficiently fat to trigger the Dismal Theorem.

We conclude that a loaded gun of strong tail dominance has not been discovered to date. At the same time, the results of the Dismal Theorem are sufficiently powerful to serve as a reminder that we must constantly be alert to this possibility. Perhaps climate change is not the Dismal event, but there is a sufficiently large number of other high-consequence, low-probability events that we need to pay careful attention to tail events.

A. Tail Events and the Dismal Theorem

A tail event is an outcome which, from the perspective of the frequency of historical events or perhaps only from intuition, should happen only once in a thousand or million or centillion years. Momentous tail events were the detonation of the first atomic weapon over Hiroshima in 1945, the sharp rise in oil prices in 1973, the 23 percent fall in stock prices in October 19, 1987, the destruction of the World Trade Towers in 2001, and the collapse of the world financial system in 2007-08.

Statisticians have known for a long time that events with "fat tailed" distributions may behave in an unintuitive way.[1] Relatively little work has

---

[1] There is no generally accepted definition of the term "fat tails," also sometimes called "heavy tails." (1) One set of definitions divides distributions into three classes. A thin-tailed distribution has a finite domain (such as the uniform), a medium-tailed distribution has exponentially declining tails (such as the normal), and a fat-tailed distribution has power-law tails (such as the Pareto distribution). See Eugene F. Schuster, "Classification of Probability Laws by Tail Behavior," *Journal of the American Statistical Association*, Vol. 79, No. 388, Dec., 1984, pp. 936-939. (2) Weitzman proposes a new definition, that a fat-tailed distribution is one whose moment generating function is infinite. As we will see below, this

examined the implications of fat tails for economic modeling and policy. In a recent series of papers, Martin Weitzman has proposed a dramatically different conclusion from standard analysis in what he has called the Dismal Theorem. This result holds that under certain circumstances the expected value of utility or marginal utility does not exist and we therefore cannot apply our standard analysis. The circumstances that Weitzman proposes are ones with fat tails and with strong risk aversion.

He summarizes the basic point as well as its application to climate change as follows: "The burden of proof in climate-change CBA {cost-benefit analysis} is presumptively upon whoever calculates expected discounted utilities without considering that structural uncertainty might matter more than discounting or pure risk. Such a middle-of-the-distribution modeler should be prepared to explain why the bad fat tail of the posterior-predictive PDF is not empirically relevant and does not play a very significant – perhaps even decisive – role in climate-change CBA."[2]

The presence of consequential tail events is potentially of great importance for both economic modeling and for economic analyses of climate change. The purpose of this note is to put tail events and the Dismal Theorem in context and analyze the range of their applicability. I conclude that Weitzman raises important issues about the selection of distributions in the analysis of decision-making under uncertainty. However, the assumptions underlying the theorems are very strong, so the broad claim to have reversed the burden of proof on the use of expected utility analysis needs to be qualified.

B. Some preliminaries on of tail dominance

The Dismal Theorem is basically about the way that tail events can dominate our analysis. We can describe the analysis here is slightly different terms that can be described as "tail dominance." We can classify problems into three classes:

---

is also the condition for the Dismal Theorem, so it is tautological. We will also see that within a class of distributions the condition will depend on incidental parameters such as the degrees of freedom.

[2] There are multiple iterations on the Dismal Theorem. These quotations are from Martin L. Weitzman, "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *The Review of Economics and Statistics,* February 2009, Vol. 91, No. 1: 1–19. I am grateful for comments on these issues from William Brainard, Gary Yohe, Richard Tol, and Martin Weitzman. Version is paris_051110.docx.

- *Tail irrelevance.* Certain problems are ones in which the distribution of the random variable makes no (or little) difference to the policy or the outcome. For example, some problems are characterized as ones in which certainty equivalence applies. In these, clearly the tails are irrelevant.

- *Weak tail dominance.* A second class of cases is one where the outcomes are strongly affected by the tails of the distributions. For example, it might be the case that the outcome is strongly affected by the tails of the distribution, but the answer converges as we continue to look at more unlikely events.

- *Strong tail dominance.* A final class of cases is one where there the outcome does not converge as we continue to look at more and more extreme events. In other words, the optimal policy or the outcome does not exist (say because it is unbounded).

I take a concrete example to show the point. Suppose we are looking at the earthquakes and need to consider what "size" earthquakes to plan for in our building codes. We cannot consider every possible outcome, so we might consider a rule of thumb that considers the "100-year earthquake." That is, we consider earthquakes that are likely to occur with at least a 1 percent per year probability, but we ignore the 100+ year earthquakes as too remote to worry about.

Suppose that an analysis reveals that this 100-year policy leads to a certain rule on setback from the fault – perhaps no structures would be allowed closer than 100 meters. The question is, how do our decisions change as we move out the tails? Do the tails dominate our decisions? To examine tail dominance, we decide to repeat the analysis with earthquake frequency cutoffs of 0.1 percent per year, 0.01 percent per year, 0.001 percent per year, and so on. In many cases, the analysis does not differ markedly as we move to more remote events. The reason is that, while earthquakes are more costly, they are not sufficiently costly to revise the outcomes substantially.

Under the Dismal Theorem, which I call strong tail dominance, the results would be completely different. As we move to higher cutoffs, our best policy continues to change. So perhaps we find that we should build further from the fault by 200 meters, 400 meters, 1000 meters, etc. We see that there is in reality no optimal policy because the policy continues to change as we move further out the tail. And we must consider the entire distribution, not just the most likely events.

Weaker tail dominance would occur when the tail have a significant impact on our policies. It might be that the optimal decisions converge as we move out the tails, but they converge slowly so that the tail is an important part of the decision outcome. So, weak tail dominance means that tails matter; strong tail dominance (as in the Dismal Theorem) means that there is no convergence as the optimal policy changes the further we move out the tail.

C. Analytics of Strong Tail Dominance and the Dismal Theorem

An early example of the implications of strong tail dominance was derived by John Geweke.[3] Geweke was concerned about the use of constant relative risk aversion (CRRA) utility in the context of Bayesian learning in economic-growth models. Recall that a CRRA utility function is of the form $U(c) = c^{1-\alpha} / (1-\alpha)$, for $\alpha \neq 1$, where c is a measure of consumption and α is the elasticity of the marginal utility of consumption [$U(c) = \ln(c)$, when $\alpha = 1$]. Weitzman usually takes the elasticity to be $\alpha > 1$, and I will follow that convention in this discussion. A central assumption in both Geweke's and in Weitzman's analyses is that consumption has a structural uncertainty that is lognormally distributed:

(1) $\ln(c) \approx \overline{c} + \varepsilon$

where $\varepsilon \approx N(0, \sigma^2)$, with mean μ and standard deviation σ.

Geweke provided a number of examples of expected utility in which expected utility would exist (be finite) or would be unbounded depending upon the value of α and the probability distribution of consumption. For example, if consumption is lognormally distributed with known mean and variance, then expected utility exists (is finite) for all α. A degenerate case comes when consumption is log-normally distributed with unknown mean and unknown variance, and when the parameters of the distribution are derived from Bayesian updating. In this case, the distribution of the parameters is a normal-gamma distribution and the expected utility is unbounded (negative infinity) for $\alpha \neq 1$. This example is of particular interest because the sampling distribution for the standard deviation of a normal distribution is a t-distribution, which is in the gamma family. The existence of expected utility is "fragile" with respect to changes in the distributions of random variables or changes in prior information. Fragile in this context denotes that with CRRA the expected utility exists with some distributions but not for others.

[3] John Geweke, "A note on some limitations of CRRA utility," *Economic Letters*, 71, 2001, 341–345.

D. A simplified version of the Dismal Theorem

We can get similar results to the Geweke-Weitzman analysis by simplifying the analysis as follows.[4] For the *utility function,* we retain the CRRA utility function. Recall that in the CRRA framework, the utility function is $U(c) \approx c^{1-\alpha} / (1-\alpha)$ (we work always with $\alpha > 1$). A high value of α signifies high risk-aversion or inequality-aversion. For the *probability distribution,* we work with the Pareto distribution (as least in the tail), this being the archetypal fat-tailed distribution. We assume that *1/c* has a Pareto distribution; that is, we look at tail events in declines in consumption. For small *c*, this implies that $f(c) \approx c^k$, $k > 0$. Note in this context that a low value of *k* signifies a distribution with a fatter tail.

Define the *conditional utility* at consumption level *c* (which denotes the probability times utility) as $V(c) = f(c)U(c)$. For this specification, $V(c) = f(c)U(c) \approx -c^k c^{1-\alpha} = -c^{k+1-\alpha}$. The question is what happens to the conditional utility as *c* tends to zero. The expected utility [the integral of *V(c)*] over the interval between zero and some positive level of consumption, $\overline{c}$, converges to a finite number as $c \to 0$ if and only if $k + 2 - \alpha > 0$. (An alternative approach would be to consider approach to some catastrophic minimum consumption level, but that raises no new issues.)

Weitzman works with *conditional marginal utility.* The conditional marginal utility is defined as $CMU(c) = f(c)U'(c)$. Expected marginal utility [the integral of *CMU(c)*] is bounded as $c \to 0$ if and only if $k + 1 - \alpha > 0$.

We can take for illustrative purposes an example where $\alpha = 2.5$ and $k = 1.5$. In this case, the conditional utility is $V(c) \approx -c^{1 - 2.5 + 1.5} = -c^{-2}$. A minimal amount of calculation will show that this combination of parameters leads to bounded expected utility and bounded expected marginal utility. On the other hand, assume that $\alpha = 3.5$ and $k = 1$, in which case the conditional utility is $V(c) \approx c^{1 - 3.5 + 1} = c^{-1.5}$. For this case, both expected utility and expected marginal utility are unbounded.

The intuition behind these results is straightforward: The Dismal Theorem holds if the distribution is not only fat tailed but *very fait tailed* (meaning that *k* is

[4] Weitzman usually works with the expected value of marginal utility, while we focus on the expected value of utility. The parameter conditions for divergence are slightly different for the two, but the general insights are the same.

small), or if the utility function shows not only risk aversion but *very high risk-aversion* (meaning that α is large).

While this example simplifies the logic of the argument, it shows some important points. It shows that fat tails *per se* are not sufficient to lead to unbounded expected utility or expected marginal utility. Moreover, the question of boundedness depends upon both the parameters of the utility function and the parameters of the preference function. Note as well that this example involves the distribution of the level of consumption, whereas Weitzman's analysis involves the distribution of the log of consumption, so there is yet another important assumption involving what variable the fat-tailed distribution applies to.

E. The Role of Policy in Tail Events

Implicit in the Dismal Theorem is that the results of policies may have unbounded utility or expected utility. Recall Weitzman's argument that cost-benefit analyses "coming out of a thin-tail-based model remains under a very dark cloud" until the tail issues are resolved. In fact, there is no cost-benefit analysis in the analysis underlying the Dismal Theorem, and indeed there are no policies. How might we extend the analysis to the questions of policy?

Policy as a binary variable

One interpretation of the Dismal Theorem is that there are potentially disastrous effects of continuing business as usual in the face of global warming. To put this point analytically, assume that climate-change policy is represented by a policy variable Z. An effective policy will be interpreted as preventing climate change, so the policy variable is set at one (Z = 1). An ineffective policy will allow business as usual, so the policy variable is set at zero (Z = 0). Using this convention, we can rewrite Weitzman's model as a variant of equation (1) by adding the explicit policy variable:

(2) $$\ln(c) \approx \overline{c}(Z) + \varepsilon + \mu\,(Z-1)$$

In Weitzman's climate-change analysis, μ is the critical uncertain parameter, which is a generalized temperature sensitivity coefficient , TSC. If policy is effective, then $Z = 1$ and $\mu(Z-1) = 0$, while if policy is ineffective, then $Z \neq 1$ and $\mu(Z-1) \neq 0$. In this framework, Z is the policy variable and μ is an uncertain policy multiplier.

Weitzman assumes that μ is distributed as $\mu \approx N(\overline{\mu}\ , \sigma_{\mu}^{2})$. If the underlying distribution of μ is normal, the estimated policy multiplier (call it $\hat{\mu}$) has a t-distribution, which is fat-tailed in Weitzman's framework. This implies that that the expected utility for the CRRA utility function is unbounded (negative infinity). This arises because the policy multiplier μ has the t-distribution.

Policy as a continuous variable

In most cases, and certainly for climate-change, policy is a continuous and even multi-dimensional variable. This genuinely complicates the analysis. We can write the expected value of policy as follows:

$$V(Z) = \int_{\mu} f[T(Z;\mu)]\,U\left[c(Z,T),T(Z;\mu)\right]\,d\mu \tag{3}$$

In this formulation, f is the distribution of climatic outcomes given policy and the uncertain parameter, and *U* is utility, which is a function of consumption and climatic outcomes. The cost-benefit optimum comes where (3) is maximized with respect to the policy Z, call it *V'(Z*)=0,* for the optimal policy *Z**. There is no particular relationship between tail dominance for the expected utility in *V(Z)* and tail dominance in the optimal policy equation, *V'(Z*)=0.* It might well be the case that the policy equation has weak tail dominance even though the expected value has strong tail dominance. For example, if the outcome of asteroids displays strong tail dominance but no policy exists that can prevent the very worst outcomes, then the policy equation shows tail irrelevance. If policy enters into (3) in a separable fashion, then the uncertainties about the TSC would cancel out in the cost-benefit analysis. In any case, when we introduce policies, the analysis underlying the Dismal Theorem no longer applies directly.

II. Some Key features of Tail Dominance

A. Key features of strong tail dominance

The Dismal Theorem of strong tail dominance depends upon some special assumptions. First, it is necessary that the value of the utility function tends to minus infinity (or to plus infinity for marginal utility) as consumption tends to zero. This first condition holds for all CRRA utility functions with $\alpha > 1$, but not for all utility functions with risk aversion. Second, it is necessary that the (posterior) probability distribution of consumption has fat tails. The fat tails for the distribution of consumption means that the probability associated with low values of consumption declines less rapidly than the marginal utility of consumption increases. We discuss these questions in turn.

B. Utility with near-zero consumption

We first discuss some problems that arise with CRRA for near-zero consumption. The CRRA functions that Weitzman analyzes (with $\alpha > 1$) assume that zero consumption has utility of minus negative infinity (and unbounded positive marginal utility) as consumption goes to zero. This has the unattractive and unrealistic feature that societies would pay unlimited amounts to prevent an infinitesimal probability of zero consumption. For example, assume that there is a very, very tiny probability that a killer asteroid might hit Earth, and further assume that we can deflect that asteroid for an expenditure of $10 trillion. The CRRA utility function implies that we would spend the $10 trillion *no matter how small was the probability.* Even if the probability were 10^{-10}, 10^{-20}, or even $10^{-1,000,000}$, we would spend a large fraction of world income to avoid these infinitesimally small outcomes (short of going extinct to prevent extinction).

An alternative would be to assume that near-zero consumption is extremely but not infinitely undesirable. This is analogous in the health literature to assuming that the value of avoiding an individual's statistical death is finite. To be realistic, societies tolerate a tiny probability of zero consumption if preventing zero consumption is ruinously expensive.

C. Fat tails and the distribution of parameters

The second crucial condition for the Dismal Theorem is that the probability distribution of consumption has "fat tails" as consumption approaches zero. Recalling equation (2) above, Weitzman derives this by assuming a very specific functional form for the distribution of consumption. The condition is that the structural distribution of consumption is lognormal, the uncertain policy multiplier is normally distributed, and knowledge about the distribution of the policy multiplier is attained through sampling or Bayesian learning.[5]

However, the results are not robust to minor changes in specifications. For example, a finite upper limit might be placed on the uncertain parameter, perhaps, in Weitzman's example of the temperature sensitivity coefficient, from fundamental

[5] Weitzman's analysis contains a discussion of a Bayesian analysis of the Dismal Theorem. He relies on the application of a "non-dogmatic prior distribution" in the form of a generalized power law, $p(\mu) \propto \mu^{-k}$ [using the notation of equation (2)] with a limiting argument as $k \to \infty$. I believe that the results can be obtained using an improper infinite uniform prior, which provides the same intuition as the classical discussion in the text.

physics.[6] Alternatively, the underlying distribution of the uncertain parameter might be a distribution that, with sampling, leads to a distribution of t has with thin tails. There is little reason to think that the particular distribution used in the analysis is the correct one.

The statistical approach in equation (2) proceeds in the absence of any prior information. This is not the way that most natural or social scientists derive their subjective distributions about the key parameters of important questions such as those regarding climate change, monetary policy, or tax policy. In doing statistical estimates of the radius of the universe, physicists might require that the parameter be non-negative. In the case of the temperature sensitivity, most of current knowledge comes from the application of physical principles, and until recently, *none* of scientists' judgments on the temperature sensitivity coefficient came from sampling of historical data. In general, subjective distributions on scientific parameters are derived from time series, expert opinion, statistical analyses, theory, and similar sources. There would seem little reason to force this complex process into the straightjacket of the model in equation (2).

III. Empirical Issues in the Application to Climate Change

The Dismal Theorem of strong tail dominance is a useful reminder that analysts should think carefully about the distribution functions of parameters when undertaking an analysis of uncertainties. In particular, the counterintuitive nature of fat-tailed distributions, where "23-sigma" events can happen in historical time, needs to be part of any serious analysis of risk. The events in financial markets of 1987, 1998, and 2008 are useful reminders of that important and oft-neglected point. The question addressed here is, how strong is the evidence for strong tail dominance as a general rule and in particular to climate change.

A. Estimates of the temperature-sensitivity coefficient (TSC)

The central example in Weitzman's exposition of the Dismal Theorem is the example of the temperature sensitivity coefficient. To begin with, he assumes that the TSC enters in a multiplicative way as shown in equation (2). For our purpose, we can rewrite equation (2) as:

(3) $$\ln(c_{2200}) = \ln(\bar{c}_{2200}) + TSC \times f(Z)$$

---

[6] This point is shown rigorously in Christopher J. Costello, Michael G. Neubert, Stephen A. Polasky, Andrew R. Solow, "Bounded Uncertainty and Climate Change Economics," *Proceedings of the National Academy of Sciences,* May 4, 2010, vol. 107, no. 18, pp. 8108-8110.

This equation relates the log of consumption two hundred years in the future (which is the date that Weitzman identifies) to a base value and the product of the *TSC* and *f(Z)*. I interpret Z as a climate-change policy variable in which *f(Z)* = 0 when effective climate change policies are taken (perhaps zero net carbon emissions over the next two centuries), and *f(Z)* = 1 for a business-as-usual case of rapid growth in carbon emissions over the next two centuries. Weitzman does not introduce an explicit policy variable such as *Z*, but it is implicit in the analysis and discussion of policy and models.

Weitzman's estimates of the temperature sensitivity coefficient (TSC)

The central empirical component of Weitzman's analysis is that the posterior distribution of *TSC* is extremely dispersed. I quote Weitzman's analysis of this issue at length:[7]

> In this paper I am mostly concerned with the roughly 15% of those TSC_1 "values substantially higher than 4.5 °C" which "cannot be excluded" {by the IPCC Fourth Assessment's Summary}. A grand total of twenty-two peer-reviewed studies of climate sensitivity published recently in reputable scientific journals and encompassing a wide variety of methodologies (along with 22 imputed PDFs of TSC_1) lie indirectly behind the above-quoted IPCC-AR4 (2007) summary statement. These 22 recent scientific studies cited by IPCC-AR4 are compiled in Table 9.3 and Box 10.2. It might be argued that these 22 studies are of uneven reliability and their complicatedly-related PDFs cannot easily be combined, but for the simplistic purposes of this illustrative example I do not perform any kind of formal Bayesian model-averaging or meta-analysis (or even engage in informal cherry picking). Instead I just naively assume that all 22 studies have equal credibility and for my purposes here their PDFs can be simplistically aggregated. The upper 5% probability level averaged over all 22 climate-sensitivity studies cited in IPCC-AR4 (2007) is °7 C while the median is 6.4 °C, which I take as signifying approximately that $P[TSC_1 > 7\ °C] \approx 5\%$. Glancing at Table 9.3 and Box 10.2 of IPCC-AR4, it is apparent that the upper tails of these 22 PDFs tend to be sufficiently long and fat that one is allowed from a simplistically-aggregated PDF of these 22 studies the rough approximation $P[TSC_1 > 10\ °C] \approx 1\%$.
>
> Instead of TSC_1, which stands for climate sensitivity narrowly defined, I work throughout the rest of this paper with TSC_2, which (abusing scientific terminology somewhat here) stands for a more abstract "generalized climate-sensitivity-like scaling parameter" that includes heat-induced feedbacks on the forcing from the above-mentioned releases of naturally-sequestered GHGs, increased respiration of soil microbes, climate-stressed forests, and other weakening of natural carbon sinks. Without further ado I just assume for purposes of this simplistic example that $P[TSC_2 >$

[7] Weitzman, op. cit., pp. 5, 7. Note that I have for convenience of exposition changed Weitzman's S_1 and S_2 to TSC_1 and TSC_2 to conform to the notation used here.

10 °C] ≈ 5% and P[TSC_2 > 20 °C] ≈ 1%, implying that anthropogenic doubling of CO_2-e eventually causes P[ΔT > 10 °C] ≈ 5% and P[ΔT > 20 °C] ≈ 1%, which I take as my base-case tail estimates in what follows.

Many people would agree that a 5 percent chance of a 10 °C change, or a 1 percent chance of a 20 °C change, would be a dangerous prospect for human societies. However, the procedures used to derive these numbers are flawed. I first review the technique used by Weitzman to derive the TSC and then show an alternative method.

Weitzman's estimates are in the spirit of a meta-analysis of existing statistical studies of the TSC. The problem with his procedure is the following. If we have studies with any statistical independence, then we would *never* take the average of the 95th or the 99th percentile as the appropriate estimate of those percentiles of the underlying distribution. Those numbers might be reasonable estimates of the 95th or the 99th percentile of the next study, but they are not good estimates of the percentiles of the underlying distribution. The appropriate procedure is to start with the underlying distributions, then combine them into a meta-distribution, and calculate the percentiles from the combined distribution. [8] The Weitzman procedure will be correct only if the studies are drawn from exactly the same data, so that the distributions have a perfect correlation. This is clearly not the case, as an examination of the sources, methods, and the distributions makes clear. [9]

An example will make the point. Suppose we want to estimate the 95th percentile of the estimated mean for a random normal variable, Y, for which we have 10,010 independent observations. We divide the observations into group A with the first 10 observations and group B with the next 10,000 observations. If we take 10,010 random draws of Y assuming $Y \approx N(0,1)$, then the 95th percentile of the estimated mean for the first group is 0.699, while the 95th percentile for the second group is 0.01956. Under the Weitzman procedure, we would average these to get an

---

[8] One key to the problem with this procedure is the treatment of the Gregory et al. study. That study reports a 95th percentile of ∞, which is probably because of low power at the upper end. If this were included, then under Weitzman's procedure, the 95th percentile would also be ∞.

[9] A problem involves Weitzman's procedure of moving from TSC_1 to TSC_2. Recall that the latter concept involves Weitzman's idea that the sensitivity may be much larger when other feedback mechanisms are included. While there can be little doubt that the current climate models do not capture all possible effects, Weitzman has provided no empirical foundation for his doubling of the TSC percentiles, nor has he considered the time scale on which these further feedbacks would occur.

overall standard deviation of 0.359. The correct answer is to combine the two, which yields a 95th percentile of the estimated mean of 0.01955.

Alternative Estimates of the TSC

If we were concerned about the distribution of the TSC, we would first look to alternative sources of estimates. There are three sources that provide independent and scientifically respectable sources: models, time-series estimates, and paleoclimatic data.

1. The approach that is most commonly used is to examine the ensemble of climate models. Table 1 shows the results of a compilation of 18 ocean atmospheric general circulation models (OAGCMs) as compiled by the IPCC-AR4. These have similar analytical structures but have many differences in resolution and parameterization. This approach yields a mean TSC of 3.2 per CO_2-e doubling with a standard deviation across models of 0.7 °C.

2. A second approach uses time series data from the historical record. The second group of estimates in Table 1 show estimates of the transient TSC (the temperature increase when CO_2 concentrations double after a 70-year steady increase). These are consistent with the climate models.

3. A completely different approach is to put an upper bound on the TSC. This approach has apparently not been used, so I will sketch the method. We have relatively plentiful data from ice cores on historical CO_2 concentrations and temperature proxies (generally oxygen isotopes). It is well known that these show a very strong correlation, and this has sometimes been interpreted as proving the validity of the greenhouse effect. Unfortunately, there are positive feedbacks from temperature to CO_2 emissions and concentrations, so it is not possible to infer the structural relationship without further data. However, we can use the data to infer an upper bound on the TSC. Assume that the temperature- CO_2 relationship is $T(t) = \lambda + \alpha \ln[CO2(t)] + u(t)$, while CO_2 is also driven by temperature, with a relationship of the form $\ln[CO2(t)] = \mu + \gamma T(t) + \varepsilon(t)$. We are concerned about the distribution of α, which is the linearized TSC. Some algebra shows that a least squares estimate of the temperature equation yields a coefficient which is given by

(4)

$$E(\hat{\alpha}) = \frac{\alpha\sigma^2(\varepsilon) + \gamma\sigma^2(u)}{\sigma^2(\varepsilon) + \gamma^2\sigma^2(u)} = \alpha + \frac{\sigma^2(u)\gamma[1 - \alpha\gamma]}{\sigma^2(\varepsilon) + \gamma^2\sigma^2(u)} > \alpha;$$

$$[1 - \alpha\gamma] > 0 \text{ for stability.}$$

This will provide an unbiased estimate of the TSC if there is no feedback $(\gamma = 0)$, or if all the shocks come in the CO_2 equation. The important point here is that all the terms in the equation are presumptively positive, so the OLS estimate of α will be larger than the structural parameter. The last row of Table 1 shows an estimate of the coefficient in the OLS using data from the last 400,000 years from the Vostok ice core. The estimated coefficient is 12.95 (± 0.40). This applies to Vostok and is likely to have high-latitude amplification, I have reduced the estimate by a factor of 1.2 for the number in Table 1 based on the results of AR4 (p. 764).

The estimates in Table 1 show that there is indeed great uncertainty about the TSC. The upper bound from the econometric Vostok estimate is uncomfortably high, but it is very well determined statistically and would seem a secure upper bound in the context of the Dismal Theorem. So whatever distribution we impose on the TSC, we should probably [sic] truncate it at around a value of 12 °C per CO_2-e doubling. Drawing on the results in Costello et al., this then rules out strong tail dominance.

B. The Distribution of Consumption Declines

The Dismal Theorem concerns evaluating situations where consumption approaches "zero." From an empirical point of view, it would seem extremely difficult to determine a reliable probability distribution for the decline in consumption. In the case of climate change, for example, there are severe difficulties in estimating even the central values. Impact analyses are extremely crude and generally apply to high-income countries and to global mean temperature increases up to 2½ °C, with few studies of impacts above that level.[10] Determining tail behavior is even more difficult because it requires understanding the distribution of extreme outcomes for climate change as well as economic response and would involve projecting these well outside the range of current experience or estimates.

---

[10] See the survey by Tol which covers most comprehensive surveys (Richard Tol, "The Economic Effects of Climate Change," *Journal of Economic Perspectives,* 23(2): 29–51). The recent IPCC report did not even attempt to provide a comprehensive estimate of impacts.

From the point of view of determining the importance of tail events and the applicability of the Dismal Theorem, the key issue is whether the tail distribution of consumption changes is fat or thin. One useful approach would be to examine the history of extreme consumption changes in economic history. Output shocks come from a wide variety of sources including international and civil wars, trade shocks, revolutions, regime changes, droughts, and similar events. While climate shocks are different from other extreme shocks, they raise similar issues of social response, adaptability to extreme stress, availability of international trade and aid, and resort to conflict under conditions of economic stress. These are also useful analog because it is difficult to see how climate-induced economic impacts could compete with World War II, genocides, and prolonged civil strife in its economic stress.

For this purpose, I examined two data sources to determine both the frequency of extreme output shocks and the shape of the tail of the distribution. The first source of data is the study by Emi Nakamura, Jon Steinsson, Robert Barro, and Jose Ursua (NSBU) on the frequency of extreme events. This study looks at the prevalence of extreme consumption shocks of 24 relatively developed countries over the last century.[11] Their definition of disaster is relatively complicated, but it involves an average cumulative consumption decline of 30 percent from peak to trough. They find a probability of entering a disaster is 1.7% per year and that disasters last on average for 6.5 years.

A more useful source for the climate-change analog is the history of all countries. For this purpose, I collected data for 188 countries for the period 1950-2007 from the Penn World Table 6.3. I then examined declines in per capita output over different 10-year periods (although different lengths of time produced similar results). The idea is that a major shock from war, depression, drought, or other calamities would typically last in the order of a decade. This exercise produced 6543 (overlapping) observations.

If we define disasters in a similar way to NSBU as ones with peak-to-trough declines of at least 15 percent, this produces a mean output decline of 30 percent in output disasters. By this definition, we find that the number of episodes is 11 percent of all periods for the sample of countries. That is to say, 11 percent of years of all countries were ones that belonged to periods in which cumulative output declines were at least 15 percent. This number is almost ten times higher than those in NSBU because the larger sample includes developing countries. Moreover, major economic declines of more than 50 percent occurred in 1.1 percent of country-years.

---

11 Emi Nakamura, Jon Steinsson, Robert Barro, Jose Ursua, "Crises and Recoveries in an Empirical Model of Consumption Disasters," January 31, 2010.

Table 2 shows the largest output disasters in the sample along with a rough attribution to the major cause of the disaster. Wars and commodity price turmoil are heavily represented in the table of top disasters. Table 3 shows 25 of the largest natural disasters of the same period and shows how they fit into the rank of economic disasters. Only 1 of the natural disasters is in a top economic disaster (Iran), but that is likely to be a result of revolution.

We also examined the distribution of output declines to measure the tail properties of the distribution. Figure 1 shows the tail of the distribution (for the observations with output declines more than 15 logarithmic percent). The Pareto estimates from this sample have an estimated exponent of 3.7 to 4.5 depending upon the lower threshold of the consumption decline. For the entire disaster sample, the Pareto parameter is 4.22 (± 0.16).[12] For "super-disasters" with output declines of more than 70 percent (which include the wiggly tail on the tail in Figure 1), the Pareto parameter is 2.36 (± 0.75).[13]

We used two different procedures to determine the robustness of the estimates of the Pareto parameter. Using the PWT sample just described, we examined a set of disasters that look at the largest declines for countries for all periods between 5 and 25 years. The Pareto parameter for this sample looking at the largest 30 disasters is 2.31 (± 0.43). We also looked at the largest 50 disasters in the NSBU data set. These produced a Pareto parameter of 3.87 (± 0.55).

This excursion into the incidence of economic disasters suggests that they are actually relatively frequent in recent history, particularly in poor countries. The distribution of these events does not suggest that the tail is sufficiently fat to satisfy the Dismal Theorem, however. We suggested above that the condition for the expected marginal utility to be bounded is $k+1>\alpha$. With a Pareto parameter of k = 3, unbounded expected marginal utility would require a coefficient of relative risk aversion of at least 4. While nothing can be ruled out, this seems to require a risk aversion that is beyond the level normally observed.

---

[12] Estimates of the Pareto parameter are based on Mette Rytgaard, "Estimation In The Pareto Distribution, *Astin Bulletin*, Vol. 20, No. 2, pp. 201-216.

[13] Note that this approach will underestimate the standard errors because we use overlapping samples of years. The observations in little tail to the left are all data for Liberia during its economic disaster.

IV. Tentative Conclusions

We have examined the conditions under which tail behavior is likely to dominate economic outcomes or policies. Tail dominance occurs when outcomes or policies are not robust to including an increasingly large domain of uncertain outcomes. In the extreme case contemplated by Weitzman's Dismal Theorem, tail dominance is so strong that no outcome or policy exists in the sense that the expected value of output does not converge as the domain is increased.

I have examined a simplified version of the conditions for strong tail dominance in which preferences are constant relative risk aversion and consumption is distributed as a Pareto distribution. I showed that the Dismal Theorem of strong tail dominance holds when the rate of relative risk aversion is greater than one plus the Pareto parameter. This also requires that consumption must approach zero and that the marginal utility at zero consumption is unbounded.

The cases analyzed above do not correspond directly to a particular case, but they suggest that strong tail dominance as described by the Dismal Theorem requires very strong assumptions. In the case of climate change, it would appear that we can bound the temperature sensitivity coefficient on the basis of paleoclimatic data. The distribution of economic catastrophes over the last six decades indicates that there are indeed severe and frequent output declines, but the tail of the declines is not sufficiently fat to trigger strong tail dominance.

Even though the loaded gun of strong tail dominance has not been uncovered to date, the results of the Dismal Theorem are sufficiently powerful to serve as a reminder that we must constantly be alert to this possibility. Perhaps climate change is not the Dismal event, but there are a sufficiently large number of other possibilities from exotic events such asteroids and robotic enslavement to more mundane events such as earthquakes and financial meltdowns to motivate careful attention to tail events.

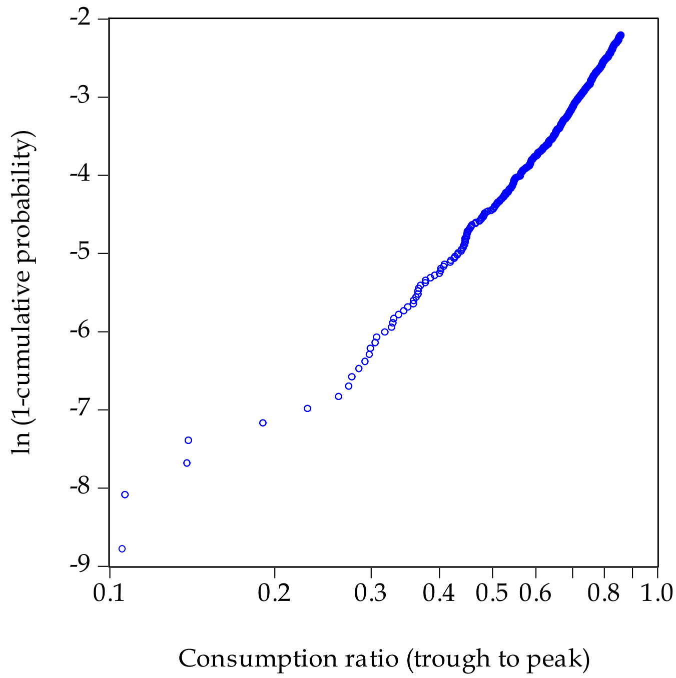

Figure 1. Cumulative distribution of consumption declines

Consumption ratio is logarithmic scale. The slope is the Pareto parameter described in the text. Sample is only consumption disasters. The Pareto parameter for this sample is 4.22 (± 0.16).

Approach	Transient	Equilibrium
Climate models (1)		
Mean	1.77	3.20
Standard deviation	0.37	0.73
Time series estimates		
20th century instrumental record (2)		
Mean	1.3 - 2.0	
Paleoclimate of Vostok CO_2 and temperature record		
Upper limit of TSC (3)		
Mean		10.79
Standard deviation		0.33

(1) 18 climate models reviewed in IPCC, Fourth Assessment Report, *Science* (FAR), p. 631.

(2) Model estimates based on 20th century observed temperatures (FAR, p. 724)

(3) This is upper bound based on reduced form estimate. For explanation, see text.

Table 1. Alternative estimates of Transient and Equilibrium Temperature Sensitivity Coefficient (TSC)

Note that the transient TSC is the response in mean temperature increase for a 1 percent per year increase in CO2 concentrations centered on the year at which doubling occurs. (See IPCC FAR, Science, p. 629)

Country	Output decline	Peak year	Trough year	Source
Liberia	89.4	1984	1994	Wars
Dem. Rep. Congo	72.3	1990	2000	Wars
Iraq	70.6	1981	1991	Wars
Afghanistan	70.0	1984	1994	Wars
Kuwait	67.0	1972	1982	Oil market
Libya	66.2	1980	1990	Oil market
Sierra Leone	60.7	1989	1999	Wars
Lebanon	56.0	1987	1997	War
Equatorial Guinea	55.3	1977	1987	Wars
Iran	54.6	1976	1986	Revolution
Saudi Arabia	54.4	1977	1987	Oil market
Brunei	51.5	1979	1989	Oil market
United Arab Emirates	50.2	1977	1987	Oil market
Guyana	49.0	1976	1986	Economic plagues
Cambodia	48.8	1970	1980	Wars
Qatar	47.3	1976	1986	Oil market
Nicaragua	47.3	1983	1993	Economic plagues
Zambia	46.4	1974	1984	Wars
Montenegro	46.2	1990	2000	Wars
Bahrain	43.6	1977	1987	Oil market

Table 2. Top economic disasters among all countries, 1950-2008

This shows the top 20 economic disasters in the post-World-War II period. An economic disaster is defined as a period in which per capita output fell by more than 15 logarithmic percentage points in a ten-year period. Note that the output-decline figures are arithmetic declines in real GDP per capita. Data from Penn World Tables, version 6.3. Sources are attributed by the author.

Country	Year	Disaster	Killed	10-year growth rate	Rank of all periods	Percentile of economic disasters [low is worst]
China	1959	Flood	2,000,000	0.06	1860	28.4
Ethiopia	1972	Famine	600,000	-0.16	658	10.1
India	1967	Drought	500,000	0.21	3484	53.2
India	1966	Drought	500,000	0.24	3878	59.3
India	1965	Drought	500,000	0.18	3153	48.2
Ethiopia	1984	Drought	300,000	-0.09	945	14.4
Bangladesh	1970	Cycl.Hurr.Typh	300,000	-0.10	867	13.3
China	1976	Earthquake	242,000	0.58	6081	92.9
Ethiopia	1974	Drought	200,000	-0.16	665	10.2
Sudan	1984	Drought	150,000	0.10	2254	34.4
Bangladesh	1991	Cycl.Hurr.Typh	138,866	0.12	2468	37.7
Mozambique	1985	Drought	100,000	-0.08	968	14.8
Ethiopia	1973	Drought	100,000	-0.07	1029	15.7
Peru	1970	Earthquake	66,794	0.38	5176	79.1
Sudan	1974	Drought	62,500	0.06	1866	28.5
Sudan	1973	Drought	62,500	-0.27	400	6.1
Ethiopia	1972	Drought	62,500	-0.08	995	15.2
Iran	1990	Earthquake	36,000	0.35	4961	75.8
Bangladesh	1965	Cycl.Hurr.Typh	36,000	-0.11	861	13.2
China	1954	Flood	30,000	0.19	3185	48.7
Bangladesh	1974	Flood	28,700	0.12	2482	37.9
Guatemala	1976	Earthquake	23,000	-0.05	1125	17.2
Colombia	1985	Volcano	21,800	0.22	3587	54.8
Iran	1978	Earthquake	20,000	-0.48	168	2.6
China	1974	Earthquake	20,000	0.54	5951	91.0

Table 3. Top natural disasters of 1950-2007 and their economic consequence

The table shows the top 25 natural disasters ranked by the number of fatalities. The last three columns show the 10-year growth rate following the disaster, the rank out of 6543 country-periods (with a low number being the worst), and the percentile rank (with low being worst). Only 1 natural disaster was associated with a top economic disaster, but that one (Iran) was due to revolution and oil-market disturbances. None of the natural disasters led to major economic disasters in this test. (Source: EM-DAT International Disaster Data Base and other.)