

An Analysis of the Dismal Theorem

William Nordhaus

January 16, 2009

Preliminary and for discussion purposes

Abstract

In a series of papers, Martin Weitzman has proposed a Dismal Theorem. The general idea is that, under limited conditions concerning the structure of uncertainty and preferences, society has an indefinitely large expected loss from high-consequence, low-probability events. Under such conditions, standard economic analysis cannot be applied. The present study is intended to put the Dismal Theorem in context and examine the range of its applicability, with an application to catastrophic climate change. I conclude that Weitzman makes an important point about selection of distributions in the analysis of decision-making under uncertainty. However, the conditions necessary for the Dismal Theorem to hold are limited and do not apply to a wide range of potential uncertain scenarios.

I. The Dismal Theorem

In a series of papers, Martin Weitzman has proposed what he calls a Dismal Theorem. He summarizes the theorem as follows: “[T]he catastrophe-insurance aspect of such a fat-tailed unlimited-exposure situation, which can never be fully learned away, can dominate the social-discounting aspect, the pure-risk aspect, and the consumption-smoothing aspect.” The general idea is that under limited conditions concerning the structure of uncertainty and preferences, the expected loss from certain risks such as climate change is infinite, and standard economic analysis cannot be applied.

He summarizes his 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.”¹

These points are potentially of great importance for both economic modeling and for economic analyses of climate change. The purpose of this note is to put the Dismal Theorem in context and analyze the range of its applicability. I conclude that Weitzman raises critical 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.

A. A Heuristic Version of the Dismal Theorem

The basic proposition under the Dismal Theorem is that with “fat tailed” distributions, expected utility analysis may behave in an unintuitive way. This arises

¹ 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,” July 7, 2008, forthcoming *The Review of Economics and Statistics*. I am grateful for comments on these issues from William Brainard, Gary Yohe, Richard Tol, and Martin Weitzman. Version is weitz_010909.docx.

because distributions with fat tails are ones for which the probabilities of rare events decline relatively slowly as the event moves far away from its central tendency.²

An early example of the difficulty, closely related to the Dismal Theorem, was derived by John Geweke in 2001.³ 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:

$$(1) \quad \ln(c) = \bar{c} + \varepsilon$$

where $\varepsilon \sim N(0, \sigma^2)$, 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

² 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 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.

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

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. In Geweke's language, 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.

Weitzman's Dismal Theorem is closely related to Geweke's. I interpret the Dismal Theorem as being about the potentially disastrous effects of taking or not taking policies, such as policies to slow global warming. An effective policy will be interpreted here as preventing climate change, so the policy variable is set at one ($P = 1$). An ineffective policy will allow climate change, so the policy variable is set at zero ($P = 0$). Using this convention, we can rewrite Weitzman's model as a variant of equation Geweke's equation (1) by adding the explicit policy variable:

$$(2) \quad \ln(c) = \bar{c} + \varepsilon + \mu(P - 1)$$

In Weitzman's climate-change analysis, μ is the critical uncertain parameter, which is a generalized temperature sensitivity coefficient (called TSC in the discussion below). If policy is effective, then $P = 1$ and $\mu(P-1) = 0$, while if policy is ineffective, then $P \neq 1$ and $\mu(P-1) \neq 0$. In this framework, P is the policy variable and μ is an uncertain policy multiplier. Weitzman assumes that the policy multiplier μ is uncertain and is distributed as $\mu \sim N(\bar{\mu}, \sigma_{\mu}^2)$. He then shows that expected utility is unbounded in a situation of Bayesian learning, although the same result would hold with classical sampling theory.

For this analysis, we can provisionally take $\varepsilon = 0$. Further assume that policy is a (0, 1) variable, such as ("weak climate-change policies, effective climate-change policies"). We can then translate the Weitzman model into the Geweke model in a straightforward way. Weitzman assumes that the policy multiplier μ is unknown and normally distributed, but observers must learn about it by sampling from history or models. To simplify, assume that we are operating in a classical framework. Then, 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 (although it is either medium-tailed or fat-tailed depending upon the degrees of freedom using our terminology of footnote 2).

Given all this, we see that if the policy variable takes a value of $P = 0$ (which is ineffective policy and therefore implies climate change) then we have the result that the expected utility for the CRRA utility function is unbounded (negative infinity). This arises because the policy multiplier μ has the t-distribution. This is Weitzman's Dismal Theorem.

B. An simplified version of the Dismal Theorem

This idea behind the Dismal Theorem can be understood intuitively as follows.⁴ Recall that in the CRRA framework, the utility function is $U(c) \sim -c^{1-\alpha}$ (we work always with $\alpha > 1$). A high value of α signifies high risk-aversion or inequality-aversion. Consider a fat-tailed probability distribution such as a power law. For small c , this implies that $f(c) \sim c^k$, $k > 0$. 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 denoted the probability times utility) as $V(c) = f(c)U(c)$. For this specification, $V(c) = f(c)U(c) \sim -c^k c^{1-\alpha} = -c^{k+1-\alpha}$. The question is what happens to the conditional utility as c tends to zero. For the Dismal Theorem to hold, $V(c)$ should go to minus infinity quickly as c approaches zero. The expected utility [the integral of $V(c)$] over the interval between zero and some positive level of consumption, \bar{c} , converges to a finite number as $c \rightarrow 0$ if and only if $k + 2 - \alpha > 0$.

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

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.

conditional utility is $V(c) \sim c^{1-2.5+1} = c^{-0.5}$. 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 fat tailed* (meaning that k is 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 finiteness 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. Finally, in this example, the exponential distribution, which Weitzman identifies as fat-tailed but we have called "medium-tailed" (see footnote 3), leads to a finite expected utility or marginal utility in all cases.

C. Some key features of the Dismal Theorem

The Dismal Theorem 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 many other utility functions. 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.

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. I consider some possible bounds in the next section.

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 policy multiplier, perhaps, in Weitzman's example of the temperature sensitivity coefficient, from fundamental physics. Alternatively, the underlying distribution of the policy multiplier might be uniform or a distribution of μ that, with sampling, leads to a distribution of $\hat{\mu}$ that

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 $k \rightarrow \infty$. 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.

has 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).

The conclusion here is that the Dismal Theorem rests on two important assumptions. First, the use of the CRRA utility function leads to unbounded negative utility as consumption approaches zero, yet this assumption leads to unacceptable conclusions in a wide variety of other circumstances. Second, the analysis relies on a very special set of assumptions about probability distributions – the lognormal distribution of the policy multipliers with no prior bounds on its values and a value determined by statistical sampling. The Dismal Theorem would not generally hold without both of these very special assumptions.

II. Further Analysis of the Dismal Theorem

I next comment on Weitzman's numerical examples as well as his criticisms of existing economic models. I begin with a general observation. Weitzman's analysis 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.⁶

The question addressed here is, does the Dismal Theory apply 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:

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

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 $f(P)$. I interpret P as a climate-change policy variable in which $f(P) = 0$ when effective climate change policies are taken (perhaps zero net carbon emissions over the next two centuries), and $f(P) = 1$ 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 P , 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 at length:⁷

In this paper I am mostly concerned with the roughly 15% of those TSC₁ "values substantially higher than 4.5 °C" which "cannot be excluded" {by the IPCC Fourth

⁶ The example of stock prices is a useful one. Prices on U.S. stock markets fell approximately 23 percent on October 19, 1987. An estimate of the daily standard deviation of price change over the 1950-1986 period (assuming a finite variance) shows a standard deviation of 1 percent. A 23-sigma event has vanishingly small probability for a normal distribution.

⁷ Weitzman, op. cit., pp. 5, 7. Note that I have for convenience of exposition changed Weitzman's S₁ and S₂ to TSC₁ and TSC₂ to conform to the notation used here.

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 $P[TSC_1 > 7^\circ\text{C}] \approx 5\%$. 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 $P[TSC_1 > 10^\circ\text{C}] \approx 1\%$.

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 $P[TSC_2 > 10^\circ\text{C}] \approx 5\%$ and $P[TSC_2 > 20^\circ\text{C}] \approx 1\%$, implying that anthropogenic doubling of CO_2 -e eventually causes $P[\Delta T > 10^\circ\text{C}] \approx 5\%$ and $P[\Delta T > 20^\circ\text{C}] \approx 1\%$, 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 catastrophic 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. 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.

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 high end. If this were included, then under Weitzman's procedure, the 95th percentile would also be ∞ .

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 $Y \sim N(0,1)$, 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 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.

A final point 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.

A Simplified Meta-analysis of Studies of the TSC

One approach to examining the TSC would be to combine different studies of that parameter. I will illustrate a simple meta-analysis that relies upon the published studies reviewed in the IPCC-AR4. Of these, ten studies provide ranges from which the distributions can be calculated (Forest, Andronova, Knutti, Frame, Forster, Wigley, Hegerl, Schneider, Murphy, and Planì). The distributions are from Table 9.2. The average of the medians is 3.1 °C, and the average of the 95th percentiles is 8.4 °C.

If we use Weitzman's approach, we get close to his answer for the mean value of TSC_1 . However, if we take a simple meta-statistical procedure in which we assume that these are independent draws from the correct distribution, we estimate that the 95th percentile is 5.0 °C if we do not weight the different estimates, and 4.6 °C if we weight the studies according to their precision (measured as the square root of the log of the 5-95 range). This compares with Weitzman's estimates of 7 °C for the 95th percentile of TSC_1 and 7 °C for the 95th percentile of TSC_2 .

These estimates are meant to be illustrative only and not to provide the best meta-analysis of different studies of the TSC. A particularly difficult issue is the fact that different studies rely on similar data (such as the instrumental temperature record), so the estimates are not independent. However, the point of this section is not to propose a new distribution but only to show that Weitzman's estimates are based on an inappropriate technique.

B. Alternative Approaches to the Lower Bound on Utility

What does "zero consumption" actually mean?

The Dismal Theorem concerns evaluating situations where consumption approaches "zero." What exactly do we mean by zero consumption? Weitzman describes the zero level as the end of "statistical civilization as we know it, or perhaps even ... of statistical life on earth..." Note that, formally, the analysis does not actually analyze "zero consumption" but takes the limit as consumption approaches zero. Zero consumption is an ambiguous concept. Is zero consumption (1) declining average consumption of a fixed number of people, or (2) high average consumption of a declining number of people, or (3) high average consumption of thriving civilizations for a statistically declining period? I suspect most people would have a different view of the undesirability of these three alternative approaches to zero.

Some approaches to a lower bound

How can we think about societal valuation of "zero consumption?" Take the third of the possible approaches to "zero consumption" of the last section (an end of human civilizations as we know them). This is the number that Weitzman takes to be unboundedly negative.

Is this really the way people decide about catastrophic events, to put infinite disutility on them? Clearly not. This question has been contemplated from time to time. It arose about two decades ago in the context of “nuclear winter,” which was the theory that the detonation of a large number of nuclear weapons would lower global temperatures so much as to kill off most if not all of humanity.⁸ More recently, there has been a spirited debate about “strangelets” and black holes triggered by heavy ion collisions in large colliders.⁹ I believe that most knowledgeable scientists would regard these as events with positive (if very low) probabilities, and they are definitely catastrophic. The low probabilities of catastrophic outcomes were not, however, enough to induce people to dismantle all but a few nuclear weapons or to stop the experiments in colliders.

One example will illustrate the issue. I would judge that the most secure example of a well-defined catastrophic risk is killer asteroids. An asteroid such as the one at the K/T (Cretaceous/Tertiary) boundary, which had a diameter of around 10 km, would probably be sufficiently large to destroy human civilizations. These are estimated to have a probability of Earth collision of about 10^{-8} per year.¹⁰ Under Weitzman’s CRRA utility function, we should devote an unlimited fraction of our resources to reduce that probability by even a small amount. Yet we are at present spending only \$4 million per year to track hazardous asteroids. Some calculations indicate that with an outlay of \$1 billion per year we could reduce the probability of impact by at least 90 percent, but this sum is apparently not worth the avoided risk. (I do not make a case that these numbers are accurate but that they are probably

⁸ One of the most influential studies was R. P. Turco, O. B. Toon, T. P. Ackerman, J. B. Pollack and Carl Sagan, “Nuclear winter: global consequences of multiple nuclear explosions,” *Science*, 222, 23 December 1983, pp. 1283-1292. This was generally disregarded after further studies. However, recent work has done new modeling and found disturbing results, see Michael J. Mills, Owen B. Toon, Richard P. Turco, Douglas E. Kinnison, and Rolando R. Garcia, “Massive global ozone loss predicted following regional nuclear conflict,” *Proc. National Acad. Sci. (US)*, pp. 5307-5312, 2008.

⁹ Arnon Dar, A. De Rújula, and Ulrich Heinz, “Will Relativistic Heavy Ion Colliders Destroy Our Planet?” A. Dar et al., *Physics Letters B*, 470 142-148 (1999); hep-ph/9910471 and R.L. Jaffea, W. Busza, J. Sandweiss, and F. Wilczek, “Review of Speculative ‘Disaster Scenarios’ at RHIC,” *Review of Modern Physics*, 72, pp. 1125-1140, 2000.

¹⁰ Clark R. Chapman and David Morrison, “Impacts on the Earth by asteroids and comets: assessing the hazard,” *Nature*, 367, 33 - 40, 06 January 1994.

within one or two orders of magnitude of the right numbers, which is sufficient for this discussion.) If we take this as an outcome of reasoned choice, these numbers would imply that the negative utility of human extinction is somewhere between $\$10^{15}$ and $\$10^{17}$. For reference purposes, the discounted value of world consumption is in the order of $\$10^{16}$. So to a first approximation the asteroid example looks more like the outcome of linear utility rather than of highly risk-averse utility.¹¹

To summarize, societies do not behave as if catastrophic outcomes have unbounded negative disutility. Perhaps, the Dismal Theorem is really a warning about applying the CRRA function to situations where consumption might be very small. If people are concerned about catastrophic impacts of climate change (as in near-zero consumption), then they should revisit their assumption about the utility function. This is one of the implications of the Dismal Theorem.

Where does the Dismal Theorem apply?

One interpretation of the Dismal Theorem is that we cannot be sure that consumption will not be zero or near-zero with 100 percent certainty. In fact, there is very little that we can rule out with 100 percent probability in most areas. However, if a positive probability of zero consumption were a worry for climate change, it would hold in a wide variety of circumstances in which we are highly uncertain (or not completely certain) about the technology or societal impacts. Example would include biotechnology, strangelets, runaway computer systems, nuclear proliferation, rogue weeds and bugs, nanotechnology, emerging tropical diseases, alien invaders, asteroids, enslavement by advanced robots, and so on. Like global warming, all of these have deep uncertainty – indeed, they may have greater uncertainty because there are fewer well-understood constants in the biological and technological world than in the geophysical world. So, if we accept the Dismal Theorem, we would probably dissolve in a sea of anxiety at the prospect of the infinity of infinitely bad outcomes.

¹¹ This is a very rough calculation as follows. Current world market consumption is in the order of $\$50 \times 10^{12}$, and we might increase this by a factor of 2 to include non-market consumption. With a growth rate of 1 percent per year and a real discount rate of 2 percent per year, this would be a present value of $\$10^{16}$.

Weitzman dismisses such anxieties. On the question of application to other areas, Weitzman says that these are “very very unlikely,” whereas catastrophic climate change is “very unlikely.” Other scientists come to very different conclusions. One example is Freeman Dyson, who believes that we are on the threshold of developing new technologies that can scrub carbon from the atmosphere at low cost.¹² To take another example, Ray Kurzweil argues strongly for the need to protect from the “GNR” (genetics, nanotechnology, robotics) revolution but believes that low-cost and clean energy will be attainable in two or three decades?¹³

What is “no policy”?

An additional question concerns how to interpret the “no policy” case. From a formal point of view, we can put this point in terms of Weitzman’s equation in equation (2). A critical part of his analysis is the implicit assumption that the policy variable is of the (0, 1) variety. In this framework, “zero” is interpreted as “no change,” as in no climate change, no research on biotechnology or nanotechnology, no international trade in dangerous viruses, no self-replicating robots, and the like. However, it is hard to think of anything that is truly a “zero” policy. And we would need to worry that slowing activities in any of these areas would take us in more dangerous directions, perhaps reducing the possibility of detecting nuclear proliferation or finding antidotes to lethal pandemics.

In the case of climate change case, “no policy” is a troublesome concept. The estimated radiative forcings from long-lived greenhouse gases are already more than half of the way to a doubling of CO₂-equivalent forcings. If the expected utility of a doubling of CO₂ equivalent is infinite negative utility, then a reading of equation (2) would indicate that this would apply equally well to half of a CO₂ doubling, or a tenth of a doubling, or even a single molecule of CO₂. Moreover, since even doing nothing requires some affirmative steps in some dimension, perhaps everything we do would have negative infinite utility.

On learning

¹² See Freeman Dyson, “The Question of Global Warming,” *The New York Review of Books*, vol. 55, No. 10, June 12, 2008.

¹³ Ray Kurzweil, *The Singularity is Near*, Penguin, New York, 2005.

There is an important difference here among these many potentially catastrophic outcomes in the potential for learning. (Richard Tol and Gary Yohe have made this general point in an important article.¹⁴) For some catastrophes, we have no possibility of learning and mid-course corrections. Strangelets are in this category. There is no point in revising our views about strangelets in the microsecond after we discover that the calculations of the physicists are wrong. No mid-course correction would be possible. Rogue bugs may be in the same category as strangelets with respect to learning: once they have escaped, they cannot be contained in the lab. Edward Teller suggested that the Trinity Test of an atomic bomb in 1945 might generate enough heat to ignite the atmosphere – a situation that could only be definitively answered by the test.

Climate change, by contrast, is a situation where we can learn as we go along. Every theory that allows for a climate sensitivity of more than 8 °C would also predict that we should see a very large warming now, with a rapid gradient over the next half-century.¹⁵ So we can learn, and then act when we learn, and perhaps even do some geoengineering while we learn some more or get our abatement policies or low-carbon technologies in place. In other words, if the Dismal Theorem were to apply, it would apply primarily to areas where we have no reasonable chance of learning and taking mid-course corrections after learning that things are heading toward a catastrophic outcome.

III. What will produce catastrophic outcomes from climate change?

Weitzman raises the important issue of whether tail distributions may invalidate analyses of the economics of climate change, and in particular on cost-benefit analyses (CBA) in integrated-assessment models (IAMs) used to tie together

¹⁴ Gary Yohe and Richard Tol, “Precaution and a Dismal Theorem: Implications for Climate Policy and Climate Research,” in H. Geman, *Risk Management in Commodity Markets*, Wiley, New York, forthcoming.

¹⁵ IPCC Fourth Assessment, Science [2007]. Intergovernmental Panel on Climate Change, *Climate Change 2007: The Physical Science Basis*, Contribution of Working Group I to the Fourth Assessment Report of the IPCC, available online at <http://ipcc-wg1.ucar.edu/wg1/wg1-report.html>, Chapter 10.

economic and geophysical relationships. Weitzman makes powerful generic criticisms of integrated assessment models. He writes that “the artificial crispness conveyed by conventional IAM-based CBAs is especially and unusually misleading compared with more-ordinary non-climate-change CBA situations,” and that in this kind of analysis “an estimate might conceivably be arbitrarily inaccurate.”

There have been many studies using IAMs to investigate the implications of uncertainty. In my 2008 book, for example, I examined the effect on the DICE model of uncertainty for major parameters (including the temperature-sensitivity coefficient), and did not find anything approaching the catastrophic results that Weitzman predicted.¹⁶ Additionally, I tested for the sensitivity in the tails out to 6 subjective standard deviations, using both a normal distribution for the parameters as well as a t-distribution, but did not find any of Weitzman’s hypothesized tail-sensitivities.

However, while most integrated-assessment models do not display catastrophic outcomes, no law of nature or economics guarantees that outcomes of rapid climate change will not be catastrophic. This section investigates the conditions under which extreme parameter values might produce catastrophic outcomes for climate change using a standard IAM.¹⁷

¹⁶ William D. Nordhaus, *The Challenge of Global Warming: Economic Models and Environmental Policy*, Yale University Press, New Haven, CT, 2008, available at http://nordhaus.econ.yale.edu/Balance_prepub.pdf.

¹⁷ An alternative approach is to examine this question from a theoretical perspective. For example, Geoffrey Heal reviewed several studies and concluded, “There are several combinations of assumptions that justify strong action [on climate change], depending on choices of the [pure rate of social time preference], the elasticity of marginal utility, the costs of climate change, the nature of uncertainty, and the way in which we react to this.” (Geoffrey Heal, “Climate Economics: A Meta-Review and Some Suggestions,” National Bureau of Economic Research, Working Paper 13927, April 2008, p. 22 and *Review of Environmental Economics and Policy*, Advance Access, published online on September 24, 2008.) However, most theoretical approaches require parameterization of various functions, so pure theory is unlikely to provide clean answers to the question of necessary and sufficient conditions. One particularly difficult problem in estimating the potential for catastrophic results is the complicated set of dynamics in geophysical systems, which involve lags of decades if not centuries.

We need to begin with a definition of “catastrophic climate change.” Much discussion of this topic focuses on environmental or ecological outcomes, for example, those concerning sea-level rise or species losses. While these are of great concern, we focus here on economic outcomes (including non-market values). In an earlier study, I studied high-consequence outcomes, where a high-consequence outcome meant a 25 percent loss in global income relative to the baseline consumption, sustained indefinitely.¹⁸ For the present estimates, I define a *catastrophic outcome as one in which world per capita consumption declines at least fifty percent below current levels*. Since output is generally estimated to grow rapidly over the coming century, such a decline is generally at least 90 percent below a reference or no-damage level.

It should be noted that such an outcome is well outside the range of most current studies. The most extreme scenario examined by the *Stern Review* – “market impacts plus risk of catastrophe plus nonmarket impacts” –represents a 32 percent decline in output relative to the baseline in 2200. Since per capita output is estimated to grow by a factor of 13 over this period, this most extreme *Stern Review* outcome still has a per capita consumption about nine times the level in 2000. So by catastrophic, we mean far beyond what is envisioned in the direst of current modeling runs.

Some possible catastrophic scenarios

After experimentation with different assumptions, we settled on the following three conditions as important ingredients for leading to extreme outcomes. First, it would be necessary either that scientists fail to understand the nature of the climate-society system in a timely fashion, or that societies fail to take steps to reduce the threat of catastrophic climate change. If the threat is understood, then there seems little doubt that it is technologically and economically possible to reduce emissions to essentially zero in a short time period at costs that might be large but are not ruinous.

A second condition is that the economic and geophysical systems lead to large climatic changes in the absence of effective policy measures. There are many

¹⁸ William Nordhaus, “Expert Opinion on Climatic Change,” *American Scientist*, January-February 1994, vol. 82, pp. 45-51.

combinations of parameters that could lead to the rapid climate change. In our simulations below, we will examine only a high temperature sensitivity coefficient as an example of unfavorable climatic conditions.

A final ingredient is economic or societal damages that are catastrophic at levels of climate change that might arise from the first two conditions. Most damage functions in the climate-change literature would not lead to catastrophic estimated damages as defined here for large temperature changes. A damage function that has sharp threshold effects would be required to lead to the catastrophic outcomes.

We undertake this approach using the DICE-2007 model. After some experiments, and based on other modelers' results, we identify three important parameters that map into the three ingredients discussed above:

1. The temperature sensitivity coefficient ("TSC")
2. The convexity of the damage function ("convex") at a relatively low tipping point
3. The ability of polities to recognize future consequences and take actions that will reduce emissions ("policy")

Additionally, we will examine the role of the discount rate (the pure rate of social time preference), although that would not appear to be a critical part of the answer.

For each of the parameters, we consider a "base value," which is the one used in the standard DICE model, along with an "extreme value," which represents what might happen in an extremely dire outcome. We do not attach any probabilities to the extreme outcomes, although our earlier uncertainty estimates would put them at the highly improbable level. Table 1 shows the parameters considered in the runs below. We make runs for 600 years with different combination of parameters and policy assumptions. Technical details on the runs are provided in the Appendix.

The results for salient variables are shown in Table 2. The first numerical row shows the social cost of carbon (SCC) for 2015. This is a useful indicator of the overall social cost of current carbon emissions. The first five columns show the results of taking each of the extreme values of the parameters *with policy*. The SCC ranges from \$42 per ton of carbon (\$/tC) in the standard case to \$350 in the most unfavorable case. The impact on economic welfare is large but not catastrophic, with a decline of around 2 percent of welfare or consumption annuity in the worst case.

(The consumption annuity is the constant level of per capita consumption that gives the same level of utility as the case in question.)

The cases *without policy* are shown in the last four columns of Table 2. A high TSC or steep damage plus no policy are not sufficient to lead to the catastrophic results. High damages plus no policy (with a tipping point of 3 degrees C) does lead to a very steep loss. However, to get genuinely catastrophic results, in the sense used here, requires all three conditions: high TSC, extremely convex damage function, and no policy, as shown in the last column. When all three of these conditions are met, the consumption annuity declines 96.4 percent relative to the baseline. The catastrophic nature of the extreme values is signaled by an initial SCC of around \$5100 per ton C. (Note that using this as a shadow price on CO₂ emissions would produce a subtraction from “green GDP” that is virtually equal to global net output.)

An important comparison is the column labeled “1+3+4” with “2+3+4.” This shows the importance of policy to avoid the catastrophic outcomes where all parameters take their extreme value. Note as well that according to the DICE model structure, the world is not yet *irreversibly* on course for a catastrophic outcome even with the most unfavorable parameters. In all cases examined, a vigorous mitigation policy is able to prevent the world from going over the catastrophic threshold.¹⁹

Summary of catastrophic simulations

We can summarize the results as follows: First, None of the extreme parameter values taken singly produces catastrophic outcomes. The reason is that adding only one extreme value is insufficiently catastrophic. Second, as long as mitigation policies are taken quickly and sharply for the catastrophic cases, no combination of extreme values is sufficient to lead to catastrophic outcomes. The reason is that if nature deals a terrible combination of parameters, then policy moves to shut down emissions or even remove CO₂ from the atmosphere.

¹⁹ An interesting question is “How much time do we have?” In the model runs undertaken here, with all parameters at their extreme values, it is necessary to move to 100 percent emissions reductions within the first eight decades to avoid catastrophic decline. This timing is clearly sensitive to the exact details of the threshold, emissions trajectory, TSC, and other model parameters.

Third, discounting is a second-order issue in the context of catastrophic outcomes. A high discount rate will slow mitigation, but it does not by itself produce policies that would lead to future catastrophes. If the future outlook is indeed catastrophic, that is understood, and policies are taken, the discount rate has little effect on the estimate of the social cost of carbon or to the optimal mitigation policy.

This leads to the fourth and major finding of our investigation: *all of the three extreme conditions must hold to obtain the catastrophic outcome*. That is to say, there must be high temperature sensitivity *plus* catastrophic damages *plus* no policy. The intuition is that a high TSC produces a steep temperature trajectory. The steep temperature trajectory produces catastrophic damages when the damage function is extremely convex. But to these we must add that economies do not take steps to prevent the chain of catastrophic events.

In the end, the major result is the importance of “policy.” As long as policy is not shut down, the world economy can avoid catastrophic outcomes. We should not think of policy in a mechanical fashion as simply turning an emissions-control dial to the appropriate level and then going about our business. Rather, policy involves a series of difficult steps. It requires understanding the complicated geophysical and socioeconomic dynamics of climate change and economic growth over many decades; it requires solving the global public goods problem by gathering most nations together to take collective action; and it means designing a mechanism for ensuring that emissions-control policies are reasonably efficient and effective. None of these is easily accomplished, but taken together they are sufficient to overcome a set of outcomes that would otherwise be catastrophic for the human condition.

IV. Summary

Martin Weitzman’s Dismal Theorem holds that, under limited conditions concerning the structure of uncertainty and preferences, society has an indefinitely large expected loss from high-consequence, low-probability events. Under such conditions, standard economic analysis cannot be applied. The analysis in the present study concludes that Weitzman makes an important point about selection of distributions in the analysis of decision-making under uncertainty. However, the

conditions necessary for the Dismal Theorem to hold are limited and do not apply to a wide range of potential uncertain scenarios.

The results of the Dismal Theorem are important in emphasizing that we must always be cautious in our assumptions about specific functional forms in empirical research - whether those concern the utility functions or the probability distributions. There are indeed deep uncertainties about virtually every aspect of the natural and social sciences of climate change. But these uncertainties can only be resolved by continued careful analysis of data and theories.

<u>Parameters</u>	<u>Base value</u>	<u>Extreme value</u>
TSC	3	10
Convex damage component		
Intercept	0	0.1
Exponent	0	6
Tipping point (°C)	none	3
Policy begins	2015	2255
Pure time discount rate	0.015	0.001

Notes:

"TSC" is the equilibrium response of global mean temperature to a doubling of atmospheric CO₂ concentrations (°C)

"Convex damage component" is a term added to the DICE damage function that has "tipping point" at specified temperature increase

"Policy begins" indicates that there are no controls until that date, then controls are optimized after that date.

"Discount rate" is pure rate of social time preference per year.

Table 1. Parameters in standard DICE runs and extreme values

Variable	1	1+3	1+4	1+3+4	1+5	2	2+3	2+4	2+3+4
	Optimal	TSC=10 with policy	High damage with policy	All extreme with policy	Base parameters with low discounting with policy	Base parameter s with no policy	TSC=10 with no policy	High damage with no policy	All extreme with no policy
Social cost of carbon, 2015 (\$/tC)	42	92	80	350	102	44	105	551	5,100
Per capita consumption (2000 \$)									
Average, 2000-2200	50,338	48,898	50,373	47,534	50,752	48,872	43,254	26,091	5,966
Minimum 2000-2200	6,801	6,799	6,799	6,796	6,799	6,800	6,800	6,800	179
Consumption annuity per capita (c)									
Thousands, 2000 \$	17765	17641	17723	17441	(b)	17718	17422	15803	634
Percent decline		0.7%	0.2%	1.8%	(b)	0.3%	1.9%	11.0%	96.4%
Objective function									
Trillions, 2005 prices	1391.1	1381.3	1387.8	1365.2	(b)	1387.4	1363.7	1218.3	(a)
Difference from optimal		-9.8	-3.3	-25.8	(b)	-3.7	-27.3	-172.8	(a)
Percent decline		0.7%	0.2%	1.9%	(b)	0.3%	2.0%	12.4%	(a)

(a) This value is a large negative number because of non-linear objective function. Refer to consumption annuity.

(b) This value is not comparable to other runs because the discount rate is different from standard cases.

(c) The consumption annuity is the level of constant consumption that yields the same discounted utility as the case under considera

Cases:

- 1: Optimal policy from 2015
- 2: Hotelling rents on carbon until 2255, then optimal policy
- 3: TSC = 10°C per CO2 doubling
- 4: Catastrophic damages at tipping point of 3 °C
- 5: Social discount rate at 0.1 % per year

Units:


- Social cost of carbon in \$ per ton carbon, 2000 US \$
- Per capita consumption in 2000 U.S. international dollars

Table 2. Results of alternative extreme values of parameters in DICE model

Appendix. Technical Background on Catastrophic Scenarios

This appendix provides details on the runs used for the “catastrophic” scenarios. The base model is online at the author’s home page at http://www.econ.yale.edu/~nordhaus/homepage/DICE2007_short.gms. The “TSC” changes that parameter (“T2XCO2”) in the model parameters from 3 to 10. The “No Policy” runs set the emissions-control rate (“miu”) at the Hotelling scarcity rent on carbon fuels for 25 periods (250 years), and then allow optimization after that time. The extreme value for the discount rate sets the rate of time preference (“B_PRSTP”) at 0.001 per year instead of 0.015 per year.

The damage convexity is slightly more complicated but is particularly critical for the results. For the extreme case we add a sixth-order term to the damage function. The specification is somewhat speculative because there is no evidence to support such an extreme damage function. The basic idea is that after some “tipping point,” damages become very steep. We assume the tipping point is 3 °C above pre-industrial levels. The exact term that is added is “coefcat*(tatm(t)/3)**expcat” where coefcat = 0.1 and expcat = 6. Further note that this term must be added to three equations [for dameq(t), yneteq(t), and yy(t)]. The following graph shows the ratio of post-damage output to pre-damage output as a function of the temperature increase.



Note that the GAMS model was unable to solve the most extreme runs because of scaling problems. We therefore used an Excel version of the model. The results for the most extreme cases are therefore only approximate, but they are so extreme that it is hard to believe that current social and political systems would survive, so that the model would probably not apply in this situation for other reasons.