

Symposium:**Fat Tails and the Economics of Climate Change****Fat Tails, Thin Tails, and Climate Change Policy**

Robert S. Pindyck*

Introduction

How much environmental damage would result from unabated water pollution, greenhouse gas (GHG) emissions, toxic waste disposal, and other potentially destructive activities? And whatever that environmental damage is expected to be, what economic and social cost will it have? In other words, what is the benefit of taking costly actions today or in the near future to reduce rates of pollution and emissions, thereby reducing damages in the future?

These questions are at the heart of environmental policy. What makes these questions interesting—and difficult—are the considerable uncertainties involved: over the underlying physical or ecological processes, over the economic impacts of environmental damage, and over technological change that might reduce those economic impacts and/or reduce the cost of limiting the environmental damage in the first place. These inherent uncertainties are especially pertinent for environmental damage that occurs or lasts over long time horizons, such as nuclear waste disposal, deforestation, and—my focus in this article—GHG emissions and climate change.

Uncertainty is often incorporated into the evaluation of climate change policy by applying Monte Carlo simulation methods to an integrated assessment model (IAM). Such models “integrate” a description of GHG emissions and their impact on temperature and other aspects of climate (a climate science model) with projections of current and future abatement costs and a description of how changes in climate affect output, consumption, and other economic variables (an economic model). An IAM might be compact and highly aggregated, or large, complex, and regionally disaggregated. But, it will always contain physical and economic relationships that are subject to uncertainty over functional form and parameter values.¹ In Monte Carlo simulations, the functional forms are usually assumed to be known with

*Massachusetts Institute of Technology, Cambridge, MA 02142; e-mail: rpindyck@mit.edu.

My thanks to Charles Kolstad and two referees for helpful comments and suggestions.

¹A well-known example of a relatively compact IAM is the Nordhaus (2008) DICE model. A much larger and complex one is the model developed by the MIT Joint Program on the Science and Policy of Global Change; see Webster et al. (2009). For a general discussion of the inherent uncertainties surrounding climate change policy, see Heal and Kriström (2002).

certainty, but parameter values for each individual simulation are drawn from probability distributions that might be estimated, otherwise inferred from data, or based on assumptions. By running hundreds or thousands of simulations, expected values and confidence intervals can be calculated for variables of interest. Adding some assumption about discount rates, one can compute and compare the present values of expected costs and benefits from some policy, along with confidence intervals.²

The validity of these types of approaches has been thrown into question by the “dismal theorem,” developed in a recent article by Weitzman (2009a). The basic idea behind the dismal theorem is straightforward. Suppose we are concerned with the increase in global mean temperature over the rest of this century, which I will denote by T . Suppose we believe that T is normally distributed with a known mean, m . Note that the normal distribution is *thin tailed*, that is, its upper tail, which reflects probabilities of very high values of T , declines to zero faster than exponentially. (I will say more about this later.) Finally, suppose we do not know the variance of the distribution and therefore estimate the variance using all available data, with Bayesian updating of our estimate as new data become available. In this case the posterior distribution for T (i.e., the distribution conditional on our estimation process for the variance) is necessarily *fat tailed*, meaning that its upper tail declines to zero more slowly than exponentially. To keep things clear, I will refer to this result as “Part 1” of the dismal theorem.

Before proceeding, it is important to note that there are other routes by which one could conclude that the distribution for T has a fat tail. For example, structural climate models with feedback loops can transform thin-tailed distributions for input variables into fat-tailed distributions for output variables such as temperature.³ Or, one might infer a fat-tailed distribution simply from observing distributions for T derived from existing climate science and economic studies.

Why does it matter whether or not the distribution for T is fat tailed? This brings us to what I will call “Part 2” of the dismal theorem. Suppose higher temperatures cause “damage” by directly causing a reduction in consumption, which for simplicity I will model as

$$C = \frac{C_0}{1 + T} \quad (1)$$

where C_0 is consumption in the absence of any warming. I will assume that a reduction in C directly reduces social welfare via a utility function $U(C)$, which I will take to have the widely used constant relative risk aversion (CRRA) form, that is,

$$U(C) = \frac{1}{1 - \eta} C^{1 - \eta} \quad (2)$$

²An alternative approach, used in Pindyck (2009, 2010), is to calibrate probability distributions for variables of interest (e.g., temperature in the year 2100) from estimates of expected values and confidence intervals derived from climate science and economic studies done by others. In work related to this article, Newbold and Daigneault (2009) explore how alternative probability distributions and damage functions affect willingness to pay to reduce emissions.

³See, for example, Roe and Baker (2007), Weitzman (2009b), and Mahadevan and Deutch (2010).

Thus marginal utility is $U'(C) = C^{-\eta}U'(C) = C^{-\eta}$, and η is the index of relative risk aversion.⁴ Note that this CRRA utility function implies that as consumption approaches zero, marginal utility becomes infinite.

Now consider what happens in the upper tail of the distribution for T . Very high values of T imply very low values for C , and thus very high values for marginal utility. If T has a thin-tailed distribution, the probabilities of extremely high values of T will be sufficiently small that the expected value of marginal utility will be finite. But if T has a fat-tailed distribution, those probabilities of extremely high values of T will be large enough to make expected marginal utility infinite. And what's wrong with that? It means that the expected gain from any policy that would reduce warming is unbounded. The reason is that with fat tails, the expected gain in utility from preventing or limiting increases in T will be infinite. This, in turn, has an alarming consequence: it means that society should be willing to sacrifice close to 100 percent of gross domestic product (GDP) to reduce GHG emissions and thereby limit warming.

As a guide to policy, the conclusion that we should be willing to sacrifice close to 100 percent of GDP to reduce GHG emissions is not very useful, or even credible, and it is unlikely that one would interpret the dismal theorem in this way. A more useful interpretation—and the one that Weitzman (2010, 2011) appears to support—is that with fat tails, traditional benefit–cost analysis based on expected values (and this would include Monte Carlo simulation exercises with IAMs, no matter the number of simulations) can be very misleading, and in particular will underestimate the gains from abatement. It also implies that when evaluating or designing a climate policy, we need to pay much more attention to the likelihood and possible consequences of extreme outcomes.

While this interpretation makes sense, there is a problem with both the dismal theorem itself and the implications I have just outlined. As popular as it is among economists (largely because of the analytical tractability it provides), there is something not quite right about the CRRA utility function of equation 2 when it is applied to extreme events. What does it mean to say that marginal utility becomes infinite as consumption approaches zero? Marginal utility should indeed become very large when consumption approaches zero—after all, zero consumption usually implies death. But “very large” is quite different from infinite. Perhaps marginal utility should approach the value of a statistical life (VSL) or (because an environmental catastrophe so bad that it drives total consumption close to zero might also mean the end for future generations) some multiple—even a large multiple—of VSL. The point here is that if we put some upper limit on the CRRA utility function so that marginal utility remains finite as consumption approaches zero, then Part 2 of the dismal theorem no longer holds: even if T has a fat-tailed distribution, the expected gain from a policy that would reduce warming is no longer unbounded, and society should not be willing to spend close to 100 percent of GDP on such a policy.

We can call the part of the utility function that applies to very low values of C (corresponding to very high values of T) as the “tail” of the utility function. I would then argue that there are two kinds of “fat tails” that we need to consider. There is fat-tailed uncertainty of the kind that Weitzman (2009a, 2010) has focused on and there are fat-tailed damage or utility functions, such as the CRRA utility function discussed above, for which marginal

⁴The index of relative risk aversion is defined as $IRRA = -CU''(C)/U'(C)$, which for the utility function of equation 2 is $\eta C^{-\eta}/C^{-\eta} = \eta$.

utility approaches infinity as C becomes very small. This article will show that in terms of the implications for the economics of climate change, both kinds of “tails” are equally relevant.⁵

In the next section, I clarify some of the differences between fat-tailed and thin-tailed distributions and provide an example by comparing two particular probability distributions for temperature change—the fat-tailed Pareto distribution and the thin-tailed exponential distribution. In the following section, I combine these two distributions with a CRRA utility function that has been modified by removing the “fat” part of the tail so that marginal utility is bounded at some upper limit. This will help to elucidate the implications of uncertainty (fat-tailed or otherwise) for climate change policy.⁶ Next, I discuss environmental and other kinds of catastrophes more generally. I present some conclusions in the final section.

Fat-Tailed versus Thin-Tailed Uncertainty

A thin-tailed probability distribution is one for which the upper tail declines to zero exponentially or faster. Such a distribution has a moment generating function, and all moments exist. An example of a thin-tailed distribution that I use in this article is the exponential distribution. If the increase in temperature at some point in the future, T , is exponentially distributed, its probability density function, $g(T)$, for $T > 0$, is:

$$g(T) = \lambda e^{-\lambda T} \quad (3)$$

The k th moment is $E(T^k) = k!/\lambda^k$, so the mean is $1/\lambda$ and the variance around the mean is $1/\lambda^2$.

A fat-tailed probability distribution is one for which the upper tail declines toward zero more slowly than exponentially, so there is no moment generating function. The example I use in this article is the Pareto or power distribution:

$$f(T) = \alpha(1 + T)^{-\alpha-1} \quad (4)$$

where $\alpha > 0$ and $T \geq 0$. The “fatness” of this distribution is determined by the parameter α ; the k th moment of the distribution will exist only for $k < \alpha$. Thus, the smaller is the value of α , the “fatter” the distribution. For example, if $\alpha = 1/2$, the mean and variance will be infinite (and we might call the distribution extremely fat, or obese). If $\alpha = 3/2$, the mean of T is $1/(\alpha - 1) = 2$, but the variance and higher moments do not exist.

⁵This article is part of a symposium on Fat Tails and the Economics of Climate Change. The other articles in the symposium are Nordhaus (2011) and Weitzman (2011).

⁶One important aspect of uncertainty, which I do not discuss in this article, is its interaction with the irreversibilities inherent in climate change policy. Atmospheric GHG concentrations decay very slowly, so that the environmental impact of emissions is partly irreversible. But any policy to reduce emissions imposes sunk costs on society (e.g., to better insulate homes, improve automobile gas mileage), and these sunk costs are also an irreversibility. These two kinds of irreversibility have opposite implications for climate change policy. For a discussion of these effects, see Pindyck (2007), and for a more technical treatment, see Pindyck (2002).

Table 1. Temperature probabilities for exponential distribution (with $\lambda = .50$) and Pareto distribution (with $\alpha = 4/3$)

$T^* =$	2°C	3°C	4.5°C	6°C	10°C	15°C	20°C	$E(T)$
Exponential: $\text{Prob}(T \geq T^*)$	0.361	0.223	0.105	0.050	0.0067	0.00055	0.000045	2°C
Pareto: $\text{Prob}(T \geq T^*)$	0.230	0.161	0.103	0.075	0.041	0.025	0.017	3°C

For policy purposes, what difference does it make whether T follows an exponential or a Pareto distribution? To address this question, I will choose α and λ so that for both distributions, the probability that T is greater than or equal to 4.5°C (the upper end of the “likely” range for temperature change by the end of the century according to the Intergovernmental Panel on Climate Change (IPCC) (2007) is 10 percent. Thus, I set $\lambda = 0.50$ and $\alpha = 4/3$. (The expected value of T is then 2°C for the exponential distribution and 3°C for the Pareto distribution.) Table 1 shows the upper tails for these distributions, that is, the probabilities of T exceeding various values, and can be compared to Table 1 in Weitzman (2011).⁷ Note that the probabilities of temperatures exceeding 6°C or higher are much larger for the fat-tailed Pareto distribution. Weitzman (2010, 2011) argues that there is indeed a sizeable probability of a very large outcome for T , an outcome that could be catastrophic.

These differences in the two distributions can also be seen graphically. Figure 1a shows the two distributions for temperature changes in the range of 0–10°C. For each distribution, the probability of a temperature change greater than 4.5°C is about 10 percent. Note that both functions drop off sharply for temperature changes above 6°C, and the tail weights appear to be about the same for these high temperatures. However, Figure 1b shows the two distributions for temperature changes in the range of 10–30°C, with the vertical scale magnified. Clearly the Pareto distribution falls to zero much more slowly than the thin-tailed exponential distribution.

As these numbers and those in Weitzman (2011) suggest, if our concern is with the likelihood of a catastrophic outcome—which we might associate with a temperature increase greater than 6°C—then the magnitude and behavior of the upper tail of the distribution seems critical. But how can we decide whether the Pareto, exponential, or some other (fat-or thin-tailed) probability distribution is the “correct” one for, say, the change in global mean temperature over the next century? As Weitzman (2009a, 2009b) has shown, one can argue that based on structural uncertainty, whatever the distribution, it should be fat tailed. But such arguments are hardly dispositive. First, I am not aware of any data that would allow us to test alternative distributional hypotheses or directly estimate the parameters of some given distribution.⁸ Second, although one can construct theoretical models (or complicated IAMs) that transform distributions for inputs into distributions for outputs,

⁷Weitzman compares a fat-tailed Pareto distribution to a thin-tailed normal distribution. Although we use different thin-tailed distributions, the basic comparison is similar—the Pareto distribution has much more mass than either thin-tailed distribution at temperatures of 6°C and higher.

⁸In Pindyck (2009, 2010), I specify a (thin-tailed) displaced gamma distribution for T and calibrate the parameters to fit the mean, 86 percent and 95 percent points based on studies compiled by IPCC (2007). But I do not do any statistical test of whether this is the “correct” distribution.

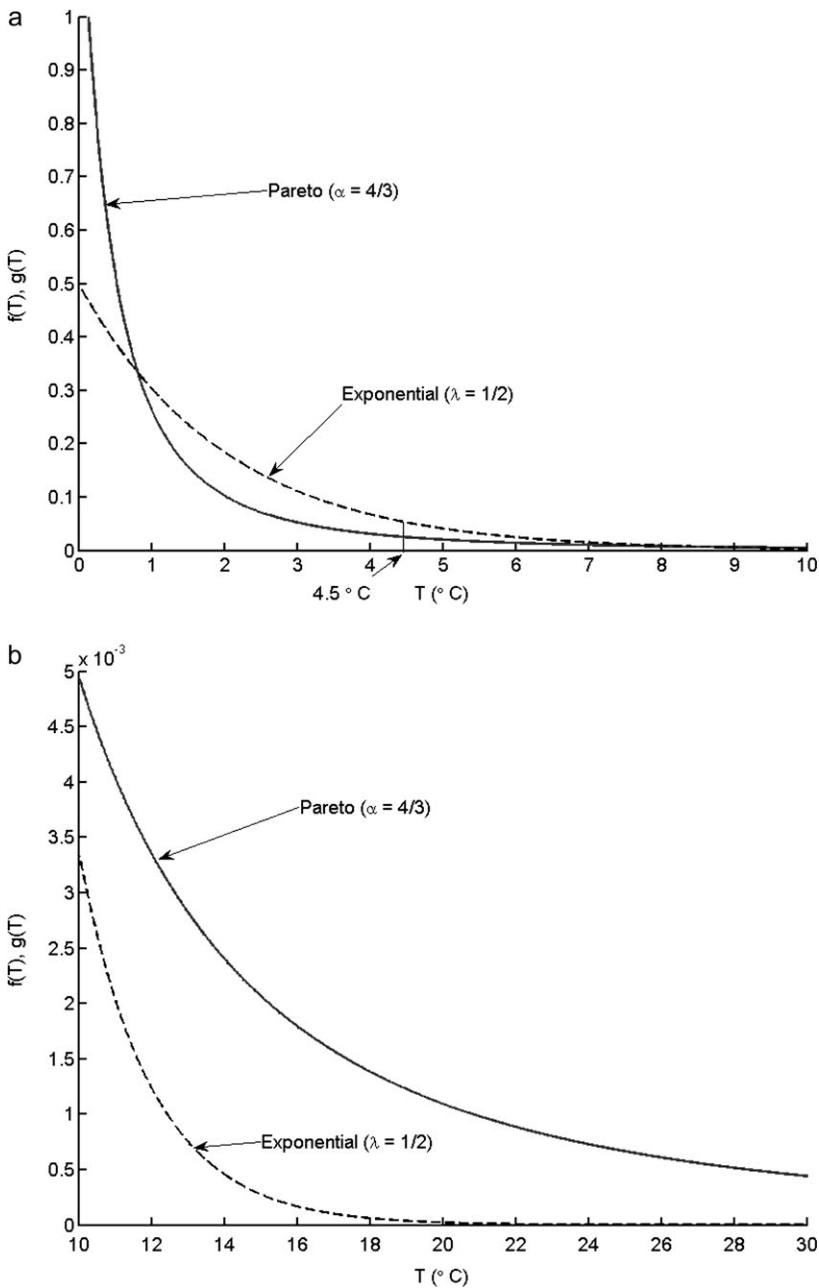


Figure I (a) Pareto and exponential distributions for increases in global mean temperature, T ; (b) Pareto and exponential distributions for large increases in global mean temperature, T

there is no consensus on a single model nor is there a consensus on input distributions. In any case, depending on parameter values, such models can yield either thin- or fat-tailed distributions.⁹

⁹For example, Mahadevan and Deutch (2010) developed a theoretical model of warming that yields a thin-tailed distribution for temperature change for some parameter values and a fat-tailed distribution for others.

If the concern is a catastrophic outcome, then perhaps it is more conservative to assume that the relevant distribution is fat-tailed. But we cannot draw such a conclusion without first addressing the implications of fat versus thin tails for expected losses and for policy. I turn to this issue in the next section.

Implications of Fat versus Thin Tails

To make this discussion as simple and straightforward as possible, I will use a stripped down model that directly connects temperature to welfare. In particular, I will assume that higher temperatures reduce consumption according to equation 1, and with no loss of generality I will set $C_0 = 1$. Note that this “damage function” leads to losses at high temperatures far worse than those projected by the Nordhaus (2008) DICE model and summarized by Weitzman (2011) in his table 3. For example, the DICE model projects a 19 percent loss of GDP and consumption with a rise in temperature of 10°C, while equation 1 projects a 91 percent loss of consumption for that temperature change. Of course, consumption itself is not the relevant variable—we need some kind of social utility function to measure the welfare effect of a 19 percent or 91 percent loss of consumption. I will use the CRRA function given by equation 2. I will also assume zero discounting of utility and zero economic growth in the absence of warming so that there is also no consumption discounting. Thus if T remains at zero, consumption and utility both remain constant over time.¹⁰

Calculating Marginal Utility

Based on equation 1, which directly connects consumption and temperature, marginal utility can be rewritten as a function of temperature in a very simple way: $MU(T) = (T + 1)^\eta$, where, again, η is the index of risk aversion. Thus as T grows and consumption falls, the marginal utility of one more unit of consumption grows. Setting η equal to either 2 or 3, which is well within the consensus range among economists, I can then calculate expected marginal utility using the Pareto and exponential distributions for T (given by equations 3 and 4 above).

Figure 2 shows the probability-weighted marginal utility as a function of the temperature increase, T , for probability weights given by the Pareto and exponential distributions, where $\eta = 2$.¹¹ Note that when weighted by the exponential distribution, marginal utility peaks at a temperature change of about 4°C and then drops rapidly to zero for high values of T . When weighted by the Pareto distribution, however, the probability weights for high temperatures are large enough so that marginal utility does not fall to zero—at any value of T . Indeed, this is why expected marginal utility is infinite under the Pareto distribution.

For policy purposes, our concern is with expected marginal utility because that is what determines the expected benefit from a policy that would reduce or limit T . Under the exponential distribution for T and assuming that the index of risk aversion $\eta = 2$, expected marginal utility is given by $E(MU) = 1 + 2/\lambda + 2/\lambda^2$, so that with $\lambda = 1/2$, $E(MU) = 13$. However, under the

¹⁰In the deterministic Ramsey growth model, the consumption discount rate is the rate of interest, which is given by $R = \delta + \eta g$, where δ is the rate of time preference (the rate at which utility is discounted, and that I assume is zero) and g is the real rate of growth of consumption.

¹¹The graph shows $f(T)(T + 1)^2$ and $g(T)(T + 1)^2$, where $f(T)$ and $g(T)$ are given by equations 3 and 4.




Figure 2 Probability-weighted marginal utility as a function of increase in temperature, T

Pareto distribution for T , expected marginal utility is infinite. This is why the Pareto, or any other fat-tailed distribution, implies a “willingness to pay” (WTP) of 100 percent of GDP to limit T by even a small amount, and why the dismal theorem is so dismal.

Putting an Upper Limit on Marginal Utility

But suppose we believe that there is some upper limit to marginal utility, so that no matter how high the temperature (and thus no matter how low is total consumption), marginal utility cannot be infinite. That upper limit might reflect the finite but very large value of a unit of consumption when total consumption is only a small fraction of today’s consumption. Or it might reflect finite but very large a fraction of the value of a human life (assuming that an environmental catastrophe causes a huge loss of consumption which in turn leads to the death of some fraction of the population), or it might be a multiple of the value of a human life (to reflect the fractional or total loss of future generations). I will assume that marginal utility reaches its maximum at some temperature T_m , and that for temperatures above T_m , marginal utility remains constant at that maximum level. For example, we might believe that any temperature change above 10°C would be catastrophic in that it would lead to roughly a 90 percent loss of consumption (which certainly seems catastrophic to me).

With this assumption and given our CRRA utility function, marginal utility is $MU(T) = (T + 1)^\eta$ for $T < T_m$, but $MU(T) = \mu(T_m + 1)^\eta$ for $T \geq T_m$, where μ is the multiplier on maximum marginal utility. This means that if $\mu = 1$, when $T \geq T_m$, marginal utility simply remains at the value it reaches at T_m , but if $\mu > 1$, marginal utility jumps to a multiple of its value at T_m and then remains at this level for any temperature above T_m . This is illustrated in Figure 3, which shows marginal utility as a function of temperature for $\eta = 2$, $\mu = 2$, and $T_m = 15^\circ\text{C}$. Note that if $T = 0$ (no warming), $C = MU = 1$. If $T = 15^\circ\text{C}$, $C = 1/16 = 0.06$, that is,




Figure 3 Marginal utility as a function of increase in temperature, T ($\eta = \mu = 2, T_m = 15^\circ\text{C}$).

consumption falls by 94 percent, and marginal utility would jump by a factor of about 500 to $2(16)^2 = 512$, which is shy of infinity but very large. (A reader who feels that these numbers are not sufficiently “catastrophic” can try other numbers for μ , etc.)

With this limit on marginal utility, expected marginal utility will be finite, even if T follows the Pareto (or any other fat-tailed) distribution. Figure 4 shows expected marginal utility as a function of the temperature T_m at which marginal utility reaches its maximum value of $\mu(T_m + 1)^\eta$, for both the exponential and Pareto distributions, and with $\eta = 2$.¹² For the solid lines, $\mu = 1$, and for the dashed lines, $\mu = 3$. Note that for the Pareto distribution, expected marginal utility is always increasing with T_m but is finite for any finite T_m and any finite maximum marginal utility. The fact that the tail of the Pareto distribution falls to zero more slowly than exponentially as T increases no longer matters because marginal utility no longer increases without bound. For the exponential distribution, if $\mu > 1$, expected marginal utility first increases to a maximum and then decreases to an asymptotic value that is independent of μ .¹³ (The reason the asymptotic value of expected marginal utility is independent of μ is that the exponential distribution declines to zero rapidly as the temperature change becomes large.)

¹²Expected marginal utility under the Pareto distribution is given by: $E[\text{MU}(T_m)] = \int_0^{T_m} \alpha(1+T)^{\eta-\alpha-1} dT + \mu(1+T_m)^\eta \int_{T_m}^\infty \alpha(1+T)^{-\alpha-1} dT = \frac{\alpha}{\eta-\alpha} [(1+T_m)^{\eta-\alpha} - 1] + \mu(1+T_m)^{\eta-\alpha}$. Figure 4 shows $E[\text{MU}(T_m)]$ for $\eta = 2$ and $\alpha = 4/3$. Expected marginal utility under the exponential distribution is: $E[\text{MU}(T_m)] = \int_0^{T_m} \lambda(1+T)^\eta e^{-\lambda T} dT + \mu(1+T_m)^\eta e^{-\lambda T_m}$. The integral on the right-hand side must be evaluated numerically. Figure 4 shows $E[\text{MU}(T_m)]$ for $\eta = 2$ and $\lambda = 1/2$.

¹³Using the equation in the previous footnote for $E[\text{MU}(T_m)]$ for the exponential distribution, take the derivative with respect to T_m , and note that for $\mu > 1$ that derivative is positive (negative) if $T_m < (>) \frac{\eta\mu}{\lambda(\mu-1)} - 1$. For $\eta = 2, \mu = 3$, and $\lambda = 1/2$, $E[\text{MU}(T_m)]$ reaches a maximum at $T_m = 5^\circ\text{C}$.

The most important result from Figure 4, however, is that it shows that for either value of μ , there is a range of T_m for which expected marginal utility is larger for the (thin-tailed) exponential distribution than for the (fat-tailed) Pareto distribution. Thus, there is a range of T_m for which the expected benefit of an abatement policy, and thus the WTP for that policy, is greater for the exponential than for the Pareto distribution. When $\mu = 1$, that range extends from 0°C to nearly 10°C , and when $\mu = 3$, it extends from 0°C to 6°C . These calculations illustrate a simple but important point. The value of an abatement policy to avoid (or insure against) a catastrophic climate outcome depends on two equally important factors: (a) the probability distribution governing outcomes (e.g., the probability of a temperature change large enough to be “catastrophic”); and (b) the impact of a catastrophic outcome, which might be summarized in the form of lost consumption and the resulting increase in the marginal utility of consumption.

A Misplaced Focus on Fat versus Thin Tails

I do not mean to downplay the importance of the probability distribution governing outcomes. As Figure 4 shows, if $\mu = 3$ and marginal utility happens to reach its maximum value at, say, $T_m = 15^\circ\text{C}$, the Pareto distribution will yield a much larger value for expected marginal utility than will the exponential distribution. Thus, it is important to determine (as best as we can) what distribution is most realistic. However, the focus on whether that distribution is fat or thin tailed is misplaced. For example, by changing the parameter λ in equation 3, one can obtain an exponential distribution that would yield a very high expected marginal utility at $T_m = 15^\circ\text{C}$. This can be seen in Figure 5, which is the same as Figure 4 except that the




Figure 4 Expected marginal utility as a function of the increase in temperature, T_m , at which marginal utility reaches its maximum value




Figure 5 Expected marginal utility as a function of the increase in temperature, T_m , at which marginal utility reaches its maximum value, with $\lambda = 1/3$

parameter λ in the exponential distribution has been reduced from $1/2$ to $1/3$ (so that both the mean temperature and standard deviation are now 3°C). Note that this small change in λ greatly increases the range of T_m over which expected marginal utility is larger for the exponential distribution than for the Pareto distribution.

In my stripped down, simple model, I assumed that the only uncertainty was about T , and that given T , we can precisely determine C and the resulting marginal utility. In reality, there is considerable uncertainty over the relationship between temperature and economic variables such as consumption (probably more uncertainty than there is over temperature itself). There is also uncertainty over the measurement of total welfare, and the use of a simple CRRA utility function is clearly an oversimplification. I could have introduced additional uncertainties and made the model more complicated, but the basic results would still hold: Expected marginal utility, and thus the expected benefit from abatement, depends not only on the probability distribution governing climate outcomes but also on the relationship between those outcomes and consumption and welfare. Furthermore, whether the probability distribution happens to be fat or thin tailed is not by itself the determining factor.

These results are also quite robust to the choice of parameters. In my analysis above, I set the index of risk aversion, η , equal to 2. However, the macroeconomics and finance literatures would put this parameter in the range of 1.5–4. Figure 6 is the same as Figure 4 (λ is again $1/2$), except that η has been increased to 3. Note that expected marginal utility rises more rapidly under the Pareto distribution than it did before, because now marginal utility is calculated as $(1 + T)^3$. However, there is still a range (although somewhat smaller) of T_m over which expected marginal utility is larger for the exponential distribution. Readers can experiment




Figure 6 Expected marginal utility as a function of the increase in temperature, T_m , at which marginal utility reaches its maximum value, with $\eta = 3$

with other parameter values for the probability distributions (α and λ), the index of risk aversion η , and the multiplier μ on maximum marginal utility.¹⁴

Catastrophic Outcomes

Many environmental economists would agree that our central concern with respect to climate change policy should be the possibility that “business as usual” will lead to a catastrophic outcome, that is, warming to such a degree, and with such a large impact, that welfare (as measured by some function of GDP or more broadly) will fall substantially and irreversibly. It is difficult to justify the immediate imposition of a very stringent abatement policy (something much more stringent than, say, the emission reductions specified in the Kyoto Accord) based on “likely” scenarios for GHG emissions, temperature change, economic impacts, and abatement costs.¹⁵ As Weitzman (2011) has argued, the case for an immediate stringent policy might then be justified as an “insurance policy” against a catastrophic outcome. But is such an insurance policy, which would be costly, indeed warranted?

Climate Catastrophes

¹⁴The MATLAB program used to generate the results in this paper is available from the author on request.

¹⁵An exception is the Stern (2007) Review, but as several authors have pointed out, that study makes assumptions about outcomes, abatement costs, and discount rates that are well outside the consensus range.

Does buying insurance now against a catastrophic climate outcome make sense? It may or may not. As with any insurance policy, the answer depends on the cost of the insurance and the likelihood and impact of a catastrophe. The cost of the insurance might indeed be warranted if the probability of a catastrophe is sufficiently large and the likely impact is sufficiently catastrophic. But note that we don't need a fat-tailed probability distribution to determine that "climate insurance" is economically justified. All we need is a significant (and it can be small) probability of a catastrophe, combined with a large benefit from averting or reducing the impact of that catastrophic outcome. As shown in the previous section, depending on parameter values, the specific damage function, and the welfare measure, the justification for "climate insurance" could well be based on a probability distribution for climate outcomes that is thin tailed.

We could push this conclusion even further so that much of the analysis in studies such as Weitzman (2010, 2011) could be bypassed altogether. If there is a significant probability (whether based on a fat- or thin-tailed distribution) of $T > 10^{\circ}\text{C}$, and if the outcome that $T > 10^{\circ}\text{C}$ would be catastrophic according to some generally agreed upon criteria, then clearly we should act quickly. We do not need a complicated analysis or a debate about social utility functions to come to this conclusion. If we face a near existential and not totally improbable threat, and we can do something to avert or at least reduce it, then we should do something about it.

Of course, determining the probability of a catastrophic outcome and its impact is no easy matter. We have very little useful data and a very limited understanding of both the climate science and the related economics. Referring back to Table 1, is the probability of $T > 10^{\circ}\text{C}$ less than 1 percent or greater than 4 percent? If we believe that T follows the fat-tailed Pareto distribution (because of "structural uncertainty" or because of feedback loops in the climate system), then the larger probability would apply. And if we are concerned with *only* these extreme outcomes, then the fat-tailed distribution implies a much stronger policy response.

However, if we are evaluating climate policies with a concern for *all* possible outcomes, then the fat-tailed distribution need not imply a stronger abatement policy. As illustrated in the previous section, once we bound the damages from warming (or more precisely, the welfare effects of those damages), it is no longer clear a priori which distribution, fat tailed or thin tailed, supports the stronger abatement policy.

Other Catastrophes

Let's return to the question of whether strong action (i.e., stringent abatement) can be justified as an insurance policy against a climate catastrophe. As explained above, answering this question is difficult because we know so little about the probability and likely impact of climate catastrophe. But that is not the only difficulty. Suppose we could somehow determine the probability distribution for various climate outcomes as well as the distribution for the impacts of those outcomes.¹⁶ Then, given a parameterized social utility function, we could in principle estimate the net benefits from various abatement policies and the WTP to avoid

¹⁶Pindyck (2009, 2010) calibrates such distributions to sets of studies done by others, but that is a far cry from saying that we "know" the true distributions.

extreme outcomes (i.e., the WTP for insurance to avoid a climate catastrophe). Suppose further that this WTP turned out to be large—say 10 percent of GDP. If 10 percent of GDP were sufficient to pay for an abatement policy that would indeed avert an extreme outcome, should we not go ahead and buy this insurance?

If a climate catastrophe were our only concern, then the answer would be straightforward—yes, we should buy the insurance.¹⁷ But matters are more complicated because a climate catastrophe is only one of a number of potential catastrophes that could cause major damage on a global scale. Readers can use their imaginations to come up with their own examples, but a few that come to my mind include a nuclear or biological terrorist attack (far worse than 9/11), a highly contagious “mega-virus” that spreads uncontrollably, or an environmental catastrophe unrelated to GHG emissions and climate change.¹⁸ These other potential catastrophes may be just as likely (or even more likely) to occur as a climate catastrophe, and could occur much sooner and with much less warning (and thus less time to adapt). Just as with climate, the likelihood and/or impact of these catastrophes could be reduced by taking costly action now.

Suppose that with no other potential catastrophes, the WTP to avoid a climate catastrophe is 10 percent of GDP. How would this WTP change once we took into account the other potential catastrophes? First, suppose that all potential catastrophes were equally likely and were “homogenous” in the sense that the likelihood, impact, and cost of reducing the likelihood and/or impact were the same for any one of them. Then, the WTP for climate would be affected in two ways, depending on the total number of potential catastrophes, their likelihood and expected impact, and the social utility function. On the one hand, the non-climate potential catastrophes would reduce the expected growth rate of GDP, thereby reducing expected future GDP and increasing expected future marginal utility *before* a climate catastrophe occurred. This in turn would increase the benefit of avoiding the further reduction of GDP that would result from a climate catastrophe. On the other hand, because all of these potential catastrophes are equally threatening, the WTP to avoid each one must be the same, which implies a large fraction of GDP would be needed to keep us safe. This “income effect” would reduce the WTP for climate. Unless the number of potential catastrophes is small, this “income effect” will dominate, so that the WTP for climate will fall. To see why, consider an extreme example in which there are twelve potential catastrophes, each with a WTP (when taken individually) of 10 percent of GDP. Spending 120 percent of GDP on catastrophe avoidance is clearly not feasible, so when taken as a group, the WTP for each potential catastrophe would have to fall.

Making matters more complicated, potential catastrophes are not homogenous and, as with climate change, are subject to considerable uncertainties (and disagreement) over their likelihood, impact, and costs of avoidance and mitigation. For example, should we buy “insurance” to reduce the likelihood of nuclear terrorism (by spending more to inspect all goods that enter the United States, by gathering more extensive intelligence, etc.)? As with climate

¹⁷The answer is actually not quite so straightforward, because if what we mean by a catastrophe is something that substantially reduces GDP, we would also have to account for general equilibrium effects, which are missing from standard benefit–cost analyses. See Pindyck and Wang (2010) for details.

¹⁸For additional examples, see Posner (2004) and Bostrom and Ćirković (2008). For a sobering discussion of the likelihood and possible impact of nuclear terrorism, see Allison (2004).

change, it depends on the expected costs and benefits of that insurance. But it also depends on the other potential catastrophes that we face and might insure against and the probability distributions governing their occurrence and impact.¹⁹ For some or all of these potential catastrophes, one could argue that there are structural uncertainties that would make the probability distributions fat tailed. However, if social welfare is bounded so that expected marginal benefits cannot be infinite, the fat-tailed versus thin-tailed distinction by itself gives us little guidance for policy.

Conclusions

The design of climate change policy is complicated by the considerable uncertainties over the benefits and costs of abatement. Even if we knew what atmospheric GHG concentrations would be over the coming century under alternative abatement policies (including no policy), we do not know the temperature changes that would result, never mind the economic impact of any particular temperature change, and the welfare effect of that economic impact. Worse, we do not even know the probability distributions for future temperatures and impacts, making any kind of benefit–cost analysis based on expected values challenging to say the least.

As Weitzman (2009a) and others have shown, there are good reasons to think that those probability distributions are fat tailed, which has the “dismal” implication that if social welfare is measured using the expectation of a CRRA utility function, we should be willing to sacrifice close to 100 percent of GDP to reduce GHG emissions and limit temperature increases. The reason is that as temperature increases without limit, so does marginal utility, and with a fat-tailed distribution the probabilities of extremely high values of T will be large enough to make expected marginal utility infinite. I have argued here, however, that the notion of an unbounded marginal utility makes little sense and that once we put a bound on marginal utility, the “dismal” implication of fat tails goes away: expected marginal utility will be finite no matter whether the distribution for T is fat or thin tailed. Furthermore, depending on the bound on marginal utility, the index of risk aversion, and the damage function, a thin-tailed distribution can actually yield a higher expected marginal utility than a fat-tailed one.

Of course, a fat-tailed distribution for temperature will have . . . fat tails, making the probability of an extreme outcome larger than it would be under a thin-tailed distribution (Table 1, Figure 1b). Weitzman (2010, 2011) suggests that this in turn justifies stringent abatement as an “insurance policy” against an extreme outcome. If our only concern is with avoiding an extreme outcome, then a fat-tailed distribution makes such an insurance policy much easier to justify. But as with any insurance policy, what matters for climate insurance is the cost of the insurance (in this case the cost of abatement) and its expected benefit, in terms of how it will shift the distribution for possible outcomes. Thus, what is important here is the entire distribution for outcomes and not necessarily whether that distribution has fat or thin tails. Once again, depending on the damage function, parameter values, and so on, climate insurance might turn out to be easier to justify with a thin-tailed distribution for outcomes.

¹⁹Pindyck and Wang (2010) estimate the WTP (in terms of a permanent tax on consumption) to reduce the likelihood or expected impact of a generic catastrophe that could occur repeatedly and would reduce the useable capital stock by a random amount. Using a calibrated general equilibrium model, they estimate the likelihood and expected impact of such a catastrophe.

The case for climate insurance is made more complicated (and harder to justify) by the fact that we face other potential catastrophes that could have impacts of magnitudes that are similar to those of a climate catastrophe. If catastrophes—climate or otherwise—would each reduce GDP and consumption by a substantial amount, then they cannot be treated independently. That is, potential nonclimate catastrophes will affect the WTP to avert or reduce a climate catastrophe and affect the economics of “climate insurance.”

So where does this leave us? The points raised in this article do *not* imply that we can dismiss the possibility of an extreme outcome (a climate catastrophe), or that a stringent abatement policy (i.e., purchasing “climate insurance”) is unwarranted. On the contrary, the possibility of an extreme outcome is central to the design and evaluation of a climate policy. We need to assess as best we can the probability distributions for climate outcomes and their impact, with an emphasis on the more extreme outcomes. We also need to better understand the cost of shifting those distributions, that is, the cost of “climate insurance.” And all of this needs to be done in the context of budget constraints and other societal needs, including schools, highways, and defense, as well as the cost of “insurance” against other potential catastrophes.

References

- Allison, Graham. 2004. *Nuclear terrorism: The ultimate preventable catastrophe*. New York: Henry Holt & Company.
- Bostrom, Nick, and Milan Ćirković, eds. 2008. *Global catastrophic risks*. Oxford: Oxford University Press.
- Heal, Geoffrey, and Bengt Kriström, 2002. Uncertainty and climate change. *Environmental and Resource Economics* 22: 3–39.
- Intergovernmental Panel on Climate Change (IPCC). 2007. *Climate change 2007: The physical science basis*. Cambridge, UK: Cambridge University Press.
- Mahadevan, L., and John M. Deutch. 2010. Influence of feedback on the stochastic evolution of simple climate systems. *Proceedings of the Royal Society A* 466: 393–1003.
- Newbold, Stephen C., and Adam Daigneault. 2009. Climate response uncertainty and the expected benefits of greenhouse gas emissions reductions. *Environmental and Resource Economics* 44: 351–377.
- Nordhaus, William D. 2008. *A question of balance: Weighing the options on global warming policies*. New Haven, CT: Yale University Press.
- . 2011. The economics of tail events with an application to climate change. *Review of Environmental Economics and Policy*, 5 (2): 240–57. 10.1093/reep/rr004.
- Pindyck, Robert S. 2002. Optimal timing problems in environmental economics. *Journal of Economic Dynamics and Control* 26: 1677–97.
- . 2007. Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1: 45–65.
- . 2009. Uncertain outcomes and climate change policy. NBER Working Paper No. 15259.
- . 2010. Modeling the Impact of Warming in Climate Change Economics. NBER Working Paper No. 15692.
- Pindyck, Robert S., and Neng Wang. 2010. The economic and policy consequences of catastrophes. NBER Working Paper No. 15373.
- Posner, Richard A. 2004. *Catastrophe*. New York: Oxford University Press.
- Roe, Gerard H., and Marcia B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318: 629–32.
- Stern, Nicholas. 2007. *The economics of climate change: The Stern review*. Cambridge University Press.
- Webster, Mort, Andrei Sokolov, John Reilly, Chris Forest, Sergey Paltsev, Adam Schlosser, Chien Wang, David Kicklighter, Marcus Sarofim, Jerry Melillo, Ronald Prinn, and Henry Jacoby. 2009. Analysis of climate policy targets under

uncertainty. MIT Joint Program on the Science and Policy of Global Change, Report No. 180.

Weitzman, Martin L. 2009a. On modeling and interpreting the economics of climate change. *Review of Economics and Statistics* 91: 1–19.

———. 2009b. Additive damages, fat-tailed climate dynamics, and uncertain discounting.

Economics: The Open-Access, Open-Assessment E-Journal 3: 39.

———. 2010. GHG targets as insurance against catastrophic climate damages. Unpublished manuscript.

———. 2011. Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*. 5 (2): 274–292. 10.1093/reep/rer006.