

NBER WORKING PAPER SERIES

FAT TAILS, THIN TAILS, AND CLIMATE CHANGE POLICY

Robert S. Pindyck

Working Paper 16353

<http://www.nber.org/papers/w16353>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2010

This paper was written for the Review of Environmental Economics and Policy. My thanks to Charles Kolstad for helpful comments and suggestions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Robert S. Pindyck. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Fat Tails, Thin Tails, and Climate Change Policy

Robert S. Pindyck

NBER Working Paper No. 16353

September 2010

JEL No. D81,Q51,Q54

ABSTRACT

Climate policy is complicated by the considerable compounded uncertainties over the costs and benefits of abatement. We don't even know the probability distributions for future temperatures and impacts, making cost-benefit analysis based on expected values challenging to say the least. There are good reasons to think that those probability distributions are fat-tailed, which implies that if social welfare is based on the expectation of a CRRA utility function, we should be willing to sacrifice close to 100% of GDP to reduce GHG emissions. I argue that unbounded marginal utility makes little sense, and once we put a bound on marginal utility, this implication of fat tails goes away: Expected marginal utility will be finite even if the distribution for outcomes is fat-tailed. Furthermore, depending on the bound on marginal utility, the index of risk aversion, and the damage function, a thin-tailed distribution can yield a higher expected marginal utility (and thus a greater willingness to pay for abatement) than a fat-tailed one.

Robert S. Pindyck

MIT Sloan School of Management

100 Main Street, E62-522

Cambridge, MA 02142

and NBER

RPINDYCK@MIT.EDU

1. 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 value of taking costly actions today or in the near future to reduce rates of pollution and emissions, and thereby reduce future damages?

These questions are at the heart of environmental policy. What makes these questions interesting – and difficult – are the considerable compounded uncertainties involved: uncertainty 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. This is especially the case for environmental damage that occurs or lasts over long time horizons, such as nuclear waste disposal, deforestation, and – my focus in this paper – GHG emissions and climate change.

Incorporating uncertainty into the evaluation of climate change policy is often done 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 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

¹ A 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, see Heal and Kriström (2002).

confidence intervals.²

The validity of these approaches has been thrown into question by the “dismal theorem” developed in a recent paper 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, μ . Note that the normal distribution is *thin-tailed*, i.e., its upper tail (reflecting probabilities of very high values of T) declines to zero faster than exponentially. (I will say more about this later.) Finally, suppose we don’t know the variance of the distribution, and therefore we estimate the variance using all available data, with Bayesian updating of our estimate as new data becomes available. In that 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

$$C = \frac{C_0}{1+T} \tag{1}$$

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, i.e.,

² An 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 paper, Newbold and Daigneault (2009) explore how alternative probability distributions and damage functions affect willingness to pay to reduce emissions.

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

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

Thus marginal utility is $U'(C) = C^{-\eta}$, and η is the index of relative risk aversion.⁴ Note 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 $U'(C)$. 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 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) seems to be promoting – is that with fat tails, traditional cost-benefit 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 the dismal theorem itself, and with 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 eqn. (2) when 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

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

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 (2009, 2010) has focused on, and there are “fat-tailed” damage or utility functions, such as the CRRA utility function discussed above, for which marginal utility approaches infinity as C becomes very small. In terms of the implications for the economics of climate change, both kinds of “tails” are equally relevant.

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 Section 3, I combine each of these two distributions with a CRRA utility function that has been modified by removing the “fat” part of the tail. This will help to elucidate the implications of uncertainty (fat-tailed or otherwise) for climate change policy.⁵ In Section 4, I discuss environmental and other kinds of catastrophes more generally, and Section 5 concludes.

2. Fat-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 which I use in this paper is the exponential

⁵ One important aspect of uncertainty, which I do not discuss in this paper, 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, etc., 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).

distribution. If the increase in temperature at some point in the future, T , is exponentially distributed, its probability density function (for $T > 0$) is given by:

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

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 towards zero more slowly than exponentially, so there is no moment generating function. The example I use in this paper 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 α the “fatter” is the distribution. For example, if $\alpha = 1/2$, the mean, variance, etc. are all 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.

What difference does it make for policy purposes whether T follows an exponential versus 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 IPCC (2007)) is 10%. Thus I set $\lambda = .50$ 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, i.e., the probabilities of T exceeding various values, and can be compared to Table 1 in Weitzman (2010b).⁶ Note that the probabilities of temperatures exceeding 6°C or higher are much larger for the fat-tailed Pareto distribution. Weitzman (2010a,b) 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 to 10°C. For each distribution, the probability mass for temperature changes 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

⁶ Weitzman 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.

temperature changes in the range of 10 to 30°C, with the vertical scale magnified. Clearly the Pareto distribution falls to zero much more slowly than the thin-tailed exponential distribution.

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: Prob($T \geq T^*$) | .361 | .223 | .105 | .050 | .0067 | .00055 | .000045 | 2°C |
| Pareto: Prob($T \geq T^*$) | .230 | .161 | .103 | .075 | .041 | .025 | .017 | 3°C |

As these numbers and those in Weitzman (2010b) 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, there is no data of which I am aware that would allow us to test alternative distributional hypotheses, or directly estimate the parameters of some given distribution.⁷ Second, although one can construct theoretical models (or complicated IAMs) that transform distributions for inputs into distributions for outputs, there is no consensus on a single model, nor is there a consensus on input distributions, and in any case such models – depending on parameter values – can alternatively yield thin- and fat-tailed distributions.⁸

If the concern is a catastrophic outcome, then perhaps assuming that the relevant distribution is fat-tailed is more conservative. But before coming to that conclusion, we must address the question of what fat versus thin tails mean for expected losses and for policy. I turn to this question next.

⁷ In Pindyck (2009, 2010), I specify a (thin-tailed) displaced gamma distribution for T , and calibrate the parameters to fit the mean, 86% and 95% points based on studies compiled by IPCC (2007). But I do not do any statistical test of whether this is the “correct” distribution.

⁸ For example, Mahadevan and Deutch (2010) develop a theoretical model of warming which for some parameter values yields a thin-tailed distribution for temperature change and for others a fat-tailed one.

3. Implications of Fat vs. Thin Tails

To make this discussion as straightforward as possible, I will use an ultra-simple stripped down model that directly connects temperature to welfare. In particular, I will assume that higher temperatures reduce consumption according to eqn. (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 Table 3 in Weitzman (2010b). For example, the DICE model projects a 19% loss of GDP and consumption at a temperature of 10°C, while eqn. (1) projects a 91% loss of consumption at that temperature. Of course consumption itself is not the relevant quantity – we need some kind of social utility function to measure the welfare effect of a 19% or 91% loss of consumption. I will use the CRRA function given by eqn. (2). I will also assume zero discounting of utility and assume 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.⁹

Given eqn. (1) connecting consumption and temperature, marginal utility can be rewritten as a function of temperature in a very simple way: $MU(T) = (T + 1)^\eta$. Thus as T grows and consumption falls, the marginal utility of one more unit of consumption grows. I will set the index of risk aversion, η , equal to either 2 or 3, which is well within the consensus range. I can then calculate expected marginal utility using the Pareto and exponential distributions for T , which are given by eqns. (3) and (4).

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, and for $\eta = 2$.¹⁰ Observe 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, that is why expected marginal utility is infinite under the Pareto distribution.

⁹ In 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.

¹⁰ The graph shows $f(T)(T + 1)^2$ and $g(T)(T + 1)^2$, where $f(T)$ and $g(T)$ are given by eqns. (3) and (4).

For policy purposes, our concern is with expected marginal utility, because that is what determines the expected benefit from some 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 $E(\text{MU}) = 1 + 2/\lambda + 2/\lambda^2$, so that with $\lambda = 1/2$, $E(\text{MU}) = 13$. Under the Pareto distribution for T , however, expected marginal utility is infinite. This is why the Pareto, or any other fat-tailed distribution, implies a “willingness to pay” of 100 percent of GDP to limit T by even a small amount, and why the dismal theorem is so dismal.

But suppose we believe that there is some upper limit to marginal utility, so that no matter how high the temperature, marginal utility cannot be infinite. That upper limit might reflect the value of a unit of consumption when total consumption is only a small fraction of today’s consumption, it might reflect a fraction of the value of a human life (assuming that an environmental catastrophe 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 a roughly 90% loss of consumption (which certainly seems catastrophic to me).

With this assumption and given our CRRA utility function, $\text{MU}(T) = (T + 1)^\eta$ for $T < T_m$ but $\text{MU}(T) = \mu(T_m + 1)^\eta$ for $T \geq T_m$. Thus 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), $C = \text{MU} = 1$. If $T = 15^\circ\text{C}$, $C = 1/16 = .06$, i.e., consumption falls by 94%, 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, for $\eta = 2$.¹¹ For the solid lines, $\mu = 1$, and for the dashed lines, $\mu = 3$. Note that for the Pareto distribution, expected marginal utility, $E[\text{MU}(T_m)]$, is always increasing in 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 $\mu > 1$, $E[\text{MU}(T_m)]$ first increases to a maximum and then decreases to an asymptotic value that is independent of μ .¹²

But most importantly, note in Figure 4 that for either value of μ , there is a range of T_m for which $E[\text{MU}(T_m)]$ 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 willingness to pay for that policy, is greater for the exponential than for the Pareto distribution. When $\mu = 1$, that range extends from 0 to nearly 10°C, and when $\mu = 3$, it extends from 0 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: (1) the probability distribution governing outcomes (e.g., the probability of a temperature change large enough to be “catastrophic”); and (2) 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.

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

¹¹ Expected 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 $\lambda = 1/2$.

¹² Using 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 $T_m = 5^\circ\text{C}$.

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 , 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 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 $E[\text{MU}(T_m)]$ is larger for the exponential distribution than for the Pareto.

In my ultra-simple model, I assumed that the only uncertainty was over 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, depend 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. I set the index of risk aversion, η , equal to 2, but the macroeconomics and finance literatures would put that parameter in the range of 1.5 to 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 goes as $(1 + T)^3$. However, there is still a range (although somewhat smaller) of T_m over which $E[\text{MU}(T_m)]$ is larger for the exponential distribution. Readers can experiment 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.¹³

¹³ The MATLAB program used to generate the results in this paper is available from the author on request.

4. 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” would lead to a catastrophic outcome – 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.¹⁴ As Weitzman (2010a) has argued, the case for an immediate stringent policy might then be justified as an “insurance policy” against a catastrophic outcome. Is such an insurance policy, which would be costly, indeed warranted?

4.1. Climate Catastrophes.

Doesn’t buying insurance against a catastrophic climate outcome make immediate 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 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 a 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 might even push this conclusion further, so that much of the analysis in studies such as Weitzman’s (2010a) might 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 $T > 10^{\circ}\text{C}$ would be catastrophic according to some generally agreed upon criteria, then clearly we should act quickly. We don’t need a complicated analysis or a debate about social utility functions to

¹⁴ An exception is the Stern Review (2007), 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.

come to this conclusion. If we face a near-existential and not totally improbable threat that we can do something to avert or at least reduce, 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% or greater than 4%? 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 *only* with these extreme outcomes, 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 was 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, will support the stronger abatement policy.

4.2. Other Catastrophes.

Let’s return to the question of whether 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 climate outcomes as well as the distribution for impacts of various outcomes.¹⁵ Then given a parameterized social utility function, we could in principle estimate the net benefits from various abatement policies and the willingness to pay to avoid extreme outcomes (i.e., the WTP for insurance to avoid a climate catastrophe). Suppose further that this willingness to pay turned out to be large – say 10% of GDP. If 10% of GDP were sufficient to pay for an abatement policy that would indeed avert an extreme outcome, shouldn’t we go ahead and buy this insurance?

¹⁵ Such distributions are calibrated to sets of studies done by others in Pindyck (2009, 2010), but that is a far cry from saying that we “know” the true distributions.

If a climate catastrophe were our only concern, then the answer is straightforward – yes, we should buy the insurance.¹⁶ 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 ones 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.¹⁷ These other potential catastrophes may be just as likely (or even more likely) to occur than a climate catastrophe, and could occur much sooner and with much less warning (and thus less time to adapt). And 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 willingness to pay to avoid a climate catastrophe is 10% of GDP. How will this WTP change once we take 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 is 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 reduce the expected growth rate of GDP, thereby reducing expected future GDP and increasing expected future marginal utility *before* a climate catastrophe occurs. 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 12 potential catastrophes, each with a WTP (when taken individually) of 10% of GDP. Spending

¹⁶ This answer is not quite right because if what we mean by a catastrophe is something that substantially reduced GDP, we would also have to account for general equilibrium effects, which are missing from standard cost-benefit analyses. See Pindyck and Wang (2010) for details.

¹⁷ For 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).

120% of GDP on catastrophe avoidance is clearly not feasible, so when taken as a group, the WTPs for each potential catastrophe would 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” (by spending more to inspect all goods that enter the U.S., gather more extensive intelligence, etc.) to reduce the likelihood of nuclear terrorism? As with climate 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.¹⁸ 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.

5. Conclusions.

The design of climate change policy is complicated by the considerable compounded uncertainties over the costs and benefits of abatement. Even if we knew what atmospheric GHG will be over the coming century under alternative abatement policies (including no policy), we don’t know the temperature changes that will result, never mind the economic impact of any particular temperature change, and the welfare effect of that economic impact. Worse, we don’t even know the probability distributions for future temperatures and impacts, making any kind of cost-benefit 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% 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

¹⁸ Pindyck and Wang (2009) 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. They use a calibrated general equilibrium model to estimate the likelihood and expected impact of a catastrophe.

expected marginal utility infinite. I have argued, however, that the notion of an unbounded marginal utility makes little sense, and 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 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. (See Table 1 and Figure 1b.) Weitzman (2010a) 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. What matters 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, etc., climate insurance might turn out to be easier to justify with a thin-tailed distribution for outcomes.

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 similar magnitudes to a climate catastrophe. If catastrophes – climate or otherwise – would each reduce GDP and consumption by a substantial amount, then they cannot be treated individually. Potential non-climate catastrophes will affect the willingness to pay 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 paper 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, i.e., 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, *Nuclear Terrorism: The Ultimate Preventable Catastrophe*, Henry Holt & Company, 2004.
- Bostrom, Nick, and Milan Ćirković, eds., *Global Catastrophic Risks*, Oxford University Press, 2008.
- Heal, Geoffrey, and Bengt Kriström, “Uncertainty and Climate Change,” *Environmental and Resource Economics*, June 2002, **22**, 3—39.
- Intergovernmental Panel on Climate Change (IPCC), “Climate Change 2007: The Physical Science Basis,” Cambridge University Press, 2007.
- Mahadevan, L., and John M. Deutch, “Influence of Feedback on the Stochastic Evolution of Simple Climate Systems,” *Proceedings of the Royal Society A*, March 19, 2010, **466**, 393—1003.
- Newbold, Stephen C., and Adam Daigneault, “Climate Response Uncertainty and the Expected Benefits of Greenhouse Gas Emissions Reductions,” *Environmental and Resource Economics*, 2009, **44**, 351—377.
- Nordhaus, William D., *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, 2008.
- Pindyck, Robert S. “Optimal Timing Problems in Environmental Economics,” *Journal of Economic Dynamics and Control*, July 2002, **26**, 1677—1697.
- Pindyck, Robert S., “Uncertainty in Environmental Economics,” *Review of Environmental Economics and Policy*, Winter 2007.
- Pindyck, Robert S., “Uncertain Outcomes and Climate Change Policy,” NBER Working Paper No. 15259, August 2009.
- Pindyck, Robert S., “Modeling the Impact of Warming in Climate Change Economics,” NBER Working Paper No. 15692, January 2010.
- Pindyck, Robert S., and Neng Wang, “The Economic and Policy Consequences of Catastrophes,” NBER Working Paper No. 15373, Sept. 2009.
- Posner, Richard A., *Catastrophe*, New York: Oxford University Press, 2004.
- Roe, Gerard H., and Marcia B. Baker, “Why is Climate Sensitivity So Unpredictable?” *Science*, October 26, 2007, **318**, 629—632.

- Stern, Nicholas, *The Economics of Climate Change: The Stern Review*, Cambridge University Press, 2007.
- 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, “Analysis of Climate Policy Targets under Uncertainty,” MIT Joint Program on the Science and Policy of Global Change, Report No. 180, Sept. 2009.
- Weitzman, Martin L., “On Modeling and Interpreting the Economics of Climate Change,” *Review of Economics and Statistics*, February 2009a, **91**, 1—19.
- Weitzman, Martin L., “Additive Damages, Fat-Tailed Climate Dynamics, and Uncertain Discounting,” unpublished manuscript, April 2009b.
- Weitzman, Martin L., “GHG Targets as Insurance Against Catastrophic Climate Damages,” unpublished manuscript, May 2010a.
- Weitzman, Martin L., “Fat-Tailed Uncertainty and the Economics of Climate Change,” *Review of Environmental Economics and Policy*, forthcoming, 2010b.

Figure 1a: Pareto and Exponential Distributions for T






Figure 2: Probability-Weighted Marginal Utility





Figure 3: Marginal Utility, $\eta = \mu = 2$, $T_m = 15^\circ\text{C}$ 


Figure 4: Expected Marginal Utility, Maximum MU = $\mu[\text{MU}(T_m)]$ 


Figure 5: Expected Marginal Utility, Maximum MU = $\mu[\text{MU}(T_m)]$, $\lambda = 1/3$ 

Figure 6: Expected Marginal Utility, Maximum MU = $\mu[\text{MU}(T_m)]$, $\eta = 3$ 