

THE CLIMATE POLICY DILEMMA^{*}

by

Robert S. Pindyck

Massachusetts Institute of Technology
Cambridge, MA 02142

This draft: November 5, 2012

Abstract: Climate policy poses a dilemma for environmental economists. The economic argument for stringent GHG abatement is far from clear. There is disagreement among both climate scientists and economists over the likelihood of alternative climate outcomes, over the nature and extent of the uncertainty over those outcomes, and over the framework that should be used to evaluate potential benefits from GHG abatement, including key policy parameters. I argue that the case for stringent abatement – if it can be made at all – cannot be based on the kinds of modeling exercises that have permeated the literature, but instead must be based on the possibility of a catastrophic outcome. I discuss how an analysis that incorporates such an outcome might be conducted.

JEL Classification Numbers: Q54; D81, Q51.

Keywords: Environmental policy, GHG abatement, uncertainty, outcome distributions, fat-tailed distributions, global warming, climate change, economic impact, catastrophic outcomes.

^{*}This paper was written for the *Review of Environmental Economics and Policy*, and is based on my Keynote Address at the 2012 Conference of the Association of Environmental and Resource Economists. My thanks to Adam Jaffe, Chris Knittel, Elizabeth Santorella, Richard Schmalensee, Rob Stavins and an anonymous referee for helpful comments and suggestions.

1. Introduction.

Most environmental problems – e.g., SO_x and NO_x emissions from coal-burning power plants – are amenable to standard cost-benefit analysis. There will be uncertainties over the costs and benefits of any candidate abatement policy, but the characteristics and extent of those uncertainties will usually be well-understood, and comparable in nature to the uncertainties involved in many other public and private policy or investment decisions. Of course economists can (and will) argue about the details of the analysis. But at a basic level, we're in well-charted territory, and we think we know what we're doing. If we come to the conclusion that a policy to reduce SO_x emissions by some amount is warranted, that conclusion will be seen – at least by most economists – as defensible and reasonable.

Not so with climate change. Climate policy poses a serious dilemma for environmental economists. Partly because of declining economic growth and rising unemployment in much of the world, there has been waning political enthusiasm for implementing stringent greenhouse gas (GHG) abatement policies, and climate change is taking a back seat to other environmental – and non-environmental – policy problems. More importantly, the economic argument for stringent GHG abatement is far from clear. There is disagreement among both climate scientists and economists over the likelihood of alternative climate outcomes, as well as the nature and extent of the uncertainty over those outcomes. There is also disagreement over the framework that should be used to evaluate the potential benefits from an abatement policy, including the social welfare function and the discount rate to be used to put future welfare benefits from abatement in present value terms. These disagreements make climate policy difficult to evaluate, and a hard sell for the public at large.

Given these disagreements and the limits to our current state of knowledge, should climate policy be a priority for environmental economists and policy makers? Many economists would support policies that result in gradual GHG abatement, and that might cost as much as 1% or 2% of GDP. But can a case be made for the *early adoption of a stringent GHG abatement policy* that would sharply reduce emissions and thereby limit the accumulation of GHGs in the atmosphere, and that would have an annual cost of more than 2% or 3% of GDP?¹ Put simply, is there a good economic argument for a stringent policy that is likely to be costly to implement

¹ To put these cost numbers in context, it has been estimated that the total direct cost of *all* environmental regulation in the U.S. has been about 2% of GDP on average. See Jaffe, Peterson, Portney and Stavins (1995).

and that would yield highly uncertain benefits only 50 or 100 years from now? This is what I call the climate policy dilemma.

Why is it so difficult to apply standard cost-benefit analysis to a GHG abatement policy? Compared to most other policy problems, the analysis is more complicated because of the very long time horizon involved, the very large uncertainties, and the difficulty of even characterizing those uncertainties. Even if there were no uncertainty, the time horizon by itself creates a problem by making the present value of future benefits extremely sensitive to the choice of discount rate – and there is considerable disagreement over what the “correct” discount rate should be. As for the uncertainties, they pertain to the extent of warming (and other aspects of climate change) under current and expected future GHG emissions, as well as the economic impact of any climate change that might occur. The impact of climate change is especially uncertain, in part because of the possibility of adaptation. We simply don’t know much about how worse off the world would be if by the end of the century the global mean temperature increased by 3° or 5°C. In fact, we may never be able to resolve these uncertainties (at least not over the next 50 years). It may be that the impact of higher temperatures is not just unknown, but also unknowable.

Over the last 20 years we have seen a proliferation of quantitative studies of climate policy, including a variety of integrated assessment models (IAMs), both large and small.² What conclusions can we draw from this large (and still growing) body of research? At the risk (or intent) of being provocative, I will argue that the case for a *stringent* GHG abatement policy cannot be based on “most likely” scenarios, i.e., climate and impact outcomes that are within our 90 or even 95% confidence range. Indeed, average estimates of the social cost of carbon (SCC) based on three widely cited IAMs range (based on the assumed discount rate) from \$5 to \$35 per ton in 2010, rising to \$16 to \$65 per ton in 2050 – values consistent with at most moderate abatement. Making the case for a stringent policy would require assumptions about costs, benefits and economic parameters that are well outside the consensus range.³

² 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).

³ For a survey of SCC estimates, see Greenstone, Kopits, and Wolverton (2011) and the Interagency Working Group (2010). Those SCC estimates were generated from three IAMs – DICE (Dynamic Integrated Climate and Economy), PAGE (Policy Analysis of the Greenhouse Effect), and FUND (Climate Framework for Uncertainty, Distribution, and Negotiation). The models are described in Nordhaus (2008) and Nordhaus and Boyer (2000), Hope (2006), and

If one accepts this point of view, then the question becomes whether an argument for stringent abatement might be based on the tail of the outcome distribution, i.e., based on the possibility of a *catastrophic climate outcome*. The kind of outcome I am referring to is not simply a very large increase in temperature (or change in other climate indicia), but rather a very large *impact*, in terms of a decline in human welfare, from whatever climate change occurs. If the likelihood and impact of such an outcome were sufficiently large, a net present value calculation *might* support a stringent policy. But addressing this question is not so simple. First, it may be that we can lay out possible catastrophic climate outcomes, but have little basis for assigning probabilities and/or a range of potential impacts. Second, once we consider the possibility of a catastrophe, we must also consider *other* potential catastrophes that could seriously threaten human welfare, and thus might be deserving of their own policy responses. Given that a maximum of 100% of GDP can be devoted to catastrophe prevention, the existence of other potential catastrophes affects the economics of policies targeting a climate catastrophe.

In the next section I discuss some of the uncertainties and areas of disagreement that complicate the evaluation of climate policy, including the framework for evaluating social welfare and key parameters. I then turn to the nature and extent of the inherent uncertainties – over temperature and other climate outcomes, and over the economic impact of those outcomes. Using a “willingness to pay” (WTP) framework, I show that alternative distributions for temperature (including fat-tailed distributions) do not provide much guidance for policy. I also explain that the key uncertainty is over economic impact, about which we know very little, and which may, in fact, be in the realm of the “unknowable.” In Section 4, I argue that the economic case for a stringent GHG abatement policy, if it is to be made at all, must be based on the possibility of a catastrophic outcome, and I discuss how an analysis that incorporates such outcomes might be conducted.

2. The Economic Evaluation of Climate Policy.

The standard economic approach to policy evaluation is to apply a net present value (NPV) calculation to the current and expected future costs and benefits for the policy. In the

Tol (2002) respectively. The Stern Review (2007), which argues for the immediate adoption of very stringent GHG abatement, is an exception, but as Nordhaus (2007), Weitzman (2007), Mendelsohn (2008) and others point out, the Stern analysis makes extreme assumptions about costs, benefits, and economic parameters.

case of climate policy, this typically involves five steps. First, projections are needed for future emissions of CO₂ and other GHGs under a “business as usual” (BAU) and one or more abatement scenarios, along with estimates of resulting atmospheric GHG concentrations. Second, these atmospheric GHG concentrations must be translated into global or regional temperature changes, along with other indicia of climate change. The third step involves economics – projecting the lost GDP and consumption (along with other measures of social welfare) likely to result from higher temperatures and other climate changes. Fourth, estimates are needed of the costs of abating GHG emissions by various amounts, i.e., the costs of the policy itself. Lastly, some assumptions about social utility and the social rate of time preference are needed so that lost GDP and consumption at different points in time can be translated into losses of social welfare, and so that these losses of social welfare (along with ongoing costs of a policy) can be put in present value terms.

These five steps are the essence of what makes up an integrated assessment model, and any IAM-based analysis of climate policy. What is important is that each of these five steps involves considerable amounts of uncertainty, disagreement among economists and climate scientists about the nature and extent of the uncertainties, and disagreement about the measurement of social welfare and the key behavioral or policy parameters that affect welfare. Given the vast amount of research that has been done by economists and climate scientists on each of these elements, why has it been so hard to reach a consensus, and what does the lack of consensus imply for climate policy?

2.1. The Discount Rate.

The disagreement and debate over the correct rate at which to discount the future benefits from GHG abatement is a good place to start to understand the climate policy dilemma. To keep things simple, let’s assume that everyone agrees (even though they don’t) that the damage from global warming and climate change generally occurs via a reduction in consumption, C , and that a reduction in C directly reduces social welfare via the widely used constant relative risk aversion (CRRA) utility function $U(C)$:

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

Here η is the index of relative risk aversion, which is also a measure of social aversion to consumption inequality across points in time. (The value of an extra unit of consumption, i.e., marginal utility, is $U'(C) = C^{-\eta}$, which declines as C grows.)

Now the question is how should we value the utility from some level of consumption 50 years from now relative to the *same amount of utility* enjoyed today? In other words, how should we discount future utility (*not* future consumption itself) so as to determine its present value? The discount rate used to do this, which I will denote by δ , is called the pure rate of time preference, but since we are looking at welfare for society as a whole, we will call it the social rate of time preference. What is the “correct” value for this discount rate?

We know from a broad range of studies that most individuals would prefer to receive a unit of consumption now rather than receive that same unit a month, a year, or 10 years from now. We also know that financial data reflecting investor behavior, as well as movements of macroeconomic aggregates reflecting consumer and firm behavior, suggest that δ is in the range of 2 to 5 percent.⁴ While a rate in this range might reflect the preferences of investors and consumers, should it also reflect intergenerational preferences and thus apply to time horizons of 50 or 100 years? In other words, should the welfare of our great-grandchildren be discounted relative to our own welfare, and if so, at what rate? The answer to this question is crucial for climate policy: a rate of even 2 percent would make the present value of future welfare gains from GHG abatement too small relative to the costs of abatement to justify almost any policy, even a quite moderate one. (With a discount rate of 2 percent, a \$1 benefit 100 years from now is worth less than 14 cents today.)

Unfortunately, economics has little to say about how we should make such intergenerational comparisons. Some economists (e.g., Stern (2007) and Heal (2009)) have argued that *on ethical grounds* the rate of time preference should be *zero* for such comparisons, i.e., that it is unethical to discount the welfare of future generations relative to our own welfare. But why is it unethical? Suppose John and Jane both have the same incomes. John saves 10 percent of his income every year in order to help finance the college educations of his (perhaps yet-to-be-born) grandchildren, while Jane prefers to spend all of her disposable income on sports cars, boats, and expensive wines. Does John’s concern for his grandchildren make him more

⁴ For an excellent survey of research on the rate of time preference, see Frederick, Loewenstein, and O’Donoghue (2002).

ethical than Jane? Putting aside their personal views, I don't think economists have much to say about that question.

It seems to me that the rate of time preference is a *policy parameter*, i.e., it simply reflects the values of policy makers, who in turn might or might not believe (or care) that their policy decisions reflect the values of voters. As a policy parameter, the rate of time preference might be positive, zero, or even negative. (Why negative? One could argue, perhaps based on simple altruism or a belief that human character is improving over time, that the welfare of our great-grandchildren should be valued *more* highly than our own.) The problem is that once we agree that the rate of time preference is somewhat of an arbitrary parameter, it becomes hard to make a clear case for (or against) a stringent climate policy. Put another way, as in other areas of economic policy, the case for a stringent climate policy should be reasonably robust, and not rely heavily on the value of a particular parameter (in this case the rate of time preference).

2.2. The Index of Relative Risk Aversion (IRRA).

The IRRA, denoted by η , can also strongly affect the economic case for a climate policy. To see this, note that η affects expected future welfare in two ways. First, the larger is η , the faster the marginal utility of consumption will decline as consumption grows. Since (other things equal) consumption is expected to grow over time, the value of additional consumption in the future is smaller the larger is η . Second, η measures risk aversion; if future consumption is uncertain, future welfare will be smaller the larger is η . Thus a higher value of η has two opposing effects on the expected benefits from an abatement policy that reduces climate-induced losses of future consumption: (1) A higher η means the marginal utility of consumption in the future will be smaller, implying a smaller benefit from avoiding a climate-induced loss of consumption. (2) If there is uncertainty over the impact of GHG accumulation, a higher η means a larger loss of expected future welfare, implying a larger benefit from abatement. Most models show that unless risk aversion is extreme (e.g., η is above 4), the first effect dominates, so an increase in η (say from 1 to 4) will reduce the benefits from an abatement policy.

Then what is the "correct" value for η that should be used when evaluating a climate policy? Economists disagree. The answer depends in part on whether we view η as a *behavioral* parameter (i.e., reflecting the behavior of consumers, investors, and firms) or a *policy* parameter (i.e., reflecting the opinions and objectives of policy makers). As a behavioral parameter, the

consensus range, based on the macroeconomics and finance literatures, extends from about 1.5 to at least 4. As a policy parameter, we can consider the fact that η also reflects aversion to consumption inequality (in this case across generations). If a future generation is expected to have twice the income and consumption as the current generation, then the marginal utility of consumption for the future generation is $1/2^\eta$ as large as for the current generation, and would be weighted accordingly in any welfare calculation. Since values of η above 3 or 4 imply a relatively very small weight for the future generation, a policy maker might view smaller values of η as more appropriate. In that case, the reasonable range might be from about 1 to 3.⁵

Whether we view η as a behavioral or policy parameter, we are left with a wide range of reasonable values, and thus a wide range of estimates of the benefits of climate change mitigation. It is much harder to justify a costly climate policy if one believes that η is closer to 3 than to 1, but at the same time it is hard to make a clear case that η should be close to 1. As with the rate of time preference, our inability to pin down this crucial parameter makes it very difficult to argue clearly for (or against) a stringent climate policy.

2.3. Estimates of Willingness to Pay.

It simplifies matters somewhat to focus on the “demand” side of climate policy, and ask how large a sacrifice society should be willing to make to achieve a policy objective. For example, we could ask what is the largest fraction of consumption that society should be willing to give up, now and throughout the future, to prevent the global mean temperature from increasing by more than, say, 3°C by the end of the century. Denote the increase in temperature by T , and suppose the willingness to pay (WTP) to ensure that $T \leq 3^\circ\text{C}$ is 2% of consumption. It may or may not be feasible to limit the temperature increase to 3°C at a cost of 2% of consumption; the cost might be less than 2% of consumption (in which case the policy would yield a positive social surplus) or it might be more than 2% (in which case the policy could not be justified economically). The 2% is simply society’s reservation price for the policy objective, i.e., the most society would pay.

The advantage of working with WTP is that we can ignore the cost side of policy, and focus on the uncertainties involved in projecting temperature increases and their impact on GDP

⁵ For discussions of this point, see Dasgupta (2008) and Nordhaus (2011).

and consumption, as well as the implications of alternative values of the rate of time preference and the index of relative risk aversion. Suppose the policy objective is to reduce GHG emissions sufficiently that any increase in the global mean temperature by the year 2100 is *at most* 3°C. Climate science studies surveyed by the IPCC (2007) suggest that the *expected* temperature increase in 2100 under business as usual (BAU) is 3°C, so this seems like a reasonable objective.⁶ To make the case for a stringent abatement policy, we would want to show that the WTP for this objective is substantial, i.e., at least 2 or 3% of GDP, and is robust to a range of reasonable values for key parameters and to a range of reasonable probability distributions for the increase in temperature and for its impact.

I used this approach in a study of climate change policy that examined the implications of uncertainty over both temperature change and its impact on the growth rate of GDP.⁷ Using information on distributions for temperature change and economic impact from studies assembled by the IPCC and others, I fit displaced gamma distributions for these variables, which I argued roughly reflect the “state of knowledge” regarding the nature and extent of uncertainty over warming and its impact. Using these distributions, I calculated the WTP to limit the increase in temperature by the end of the century to 3°C. I found that for just about any values of η between 1 and 4, that WTP was less than 2% of GDP if the rate of time preference, δ , was 1% or greater. Setting δ to zero (the so-called “ethical” value), a WTP greater than 2% could only be obtained if η was less than 1.5. For values of η above 2, the WTP was less than 1%. These results are inconsistent with the immediate adoption of a stringent abatement policy.

Of course one might argue with the particular distributions I used (or the studies surveyed by the IPCC to which those distributions are calibrated), and with the range of parameters I considered. Stern, for example, set $\eta = 1$, the lowest end of the credible range. Research by climate scientists and economists may yield distributions for temperature change and its impact

⁶ This 3°C number is actually the expected value of *climate sensitivity* according to the IPCC (2007). Climate sensitivity is the temperature increase that would result from a doubling of the atmospheric concentration of CO₂-equivalent (CO₂e). But given the IPCC’s estimated trajectory for atmospheric CO₂e and the time it takes for higher temperatures to result, 3°C is, at least roughly, the expected value of the temperature increase.

⁷ See Pindyck (2012). Most IAMs and other economic studies of climate change posit a direct relationship between higher temperatures and the *level* of GDP, whereas (on theoretical and empirical grounds) I related higher temperatures to the *growth rate* of GDP. In Pindyck (2011b), I examined the extent to which my results might have been affected by relating temperature increases to the growth rate as opposed to the level of GDP, and found that it made little difference in the resulting estimates of WTP. For other studies using WTP, see, e.g., Heal and Kriström (2002) and Newbold and Daigneault (2010).

that are more pessimistic than those I (and others) have used, and might lead us to conclude that the “correct” distributions should have more weight in the tails. I turn to that possibility next.

3. The Nature of Climate Change Uncertainty.

There are two key areas of uncertainty, and thus two distributions that we need to worry about – the extent of warming, and the economic impact from whatever warming occurs.⁸ Much of the discussion of uncertainty that has appeared in the literature has been in reference to the extent of warming and related climate changes. I believe the main reason for this discrepancy is that we know much more about the climate science than we do about economic impacts. Thus we can look at means or confidence intervals that climate scientists have arrived at for, say, temperature, and then consider the probability distributions that are consistent with those statistics. The treatment of economic impact in most IAMs is largely ad hoc. Typically a loss function $L(T)$ is specified such that $L(0) = 1$ and $L(T)$ declines as T increases, and the parameters of that loss function are chosen to yield moderate losses (e.g., 3% of GDP, so $L = .97$) from moderate values of T , e.g., 3° or 4°C.⁹ But there is little evidence on which to base the choice of parameters for the loss function, apart from yielding numbers (e.g., $L(3) = .97$) that seem “reasonable.” Furthermore, once we consider temperatures of 6°C and higher, determining the economic loss, or a distribution for that loss, become a matter of guesswork. One can plug high temperatures into IAM loss functions – see Figure 1 in Greenstone, Kopits, and Wolverton (2011) for an illustration – but the results are just extrapolations with no empirical or theoretical grounding.

Given our inability to even characterize the uncertainty over economic impact, I will focus in this section on uncertainty over T : If we do little or nothing to abate GHG emissions, how much might the temperature increase by 2100? To address this question, we would like to know what probability distribution best represents the likelihood of various outcomes for T . In summarizing 22 studies of climate sensitivity, the IPCC (2007) translated the implied distributions into a standardized form, and created graphs showing multiple distributions implied

⁸ I am assuming that all impacts of climate change, including health and social impacts, can be monetized and expressed in terms of lost GDP.

⁹ The Nordhaus (2008) DICE model uses the inverse-quadratic loss function $L(T) = 1 / (1 + \pi_1 T + \pi_2 T^2)$. Weitzman (2008) introduced the exponential-quadratic loss function, $L(T) = \exp(-\beta T^2)$, which allows for greater losses at high temperatures.

by groups of studies. I found that a displaced gamma distribution for T provided a good fit to these distributions, and in Pindyck (2011b, 2012), I calibrated the parameters of the gamma distribution to match the mean and critical 66% and 95% points from the IPCC summary. Using the calibrated distribution, I then calculated the WTP to keep T in the year 2100 at or below various levels, and obtained values for WTP that were generally small.¹⁰

A natural critique of this earlier work is that the distribution I used for T is overly optimistic. Some recent work suggests that the expected value for T is greater than 3°C, and/or that the 95% point should be at a higher temperature than the 7°C that I used. Recalibrating the distribution to a mean of 5°C indeed increases the WTP – I obtained a WTP above 3% if $\delta = 0$ and η is below 1.5. But this is hardly robust evidence for a high WTP; if $\delta = .01$ and/or η is above 1.5, the WTP again fall below 2%

Another possible critique is that the gamma distribution is thin-tailed, and I should have used a fat-tailed distribution.¹¹ This is the argument raised by Weitzman (2009, 2011) as a potential justification for stringent abatement. In Pindyck (2011a), I showed that a fat-tailed distribution by itself need not lead to larger estimates of WTP if marginal utility is bounded.¹² Nonetheless, one could argue that alternative distributions for T that are also consistent with the IPCC (2007) mean and critical points, whether fat- or thin-tailed, might result in higher values of WTP. I explore this possibility below.

3.1. Alternative Distributions for Temperature.

Figure 1 shows the gamma distribution I used in Pindyck (2011b, 2012), for temperatures up to 10°C, with parameters calibrated to match the mean and critical 66% and 95% points from the IPCC summary of climate sensitivity studies.¹³ This distribution implies a 2.9% probability

¹⁰ Given a value for T_{2100} , the temperature increase at each point in time, T_t , increases from zero to this value and then continues increasing, asymptotically approaching $2T_{2100}$. To be precise, the temperature change at time t is $T_t = 2T_{2100}[1 - (1/2)^{0.1t}]$. See Pindyck (2012) for details.

¹¹ A thin-tailed (fat-tailed) distribution for T declines to zero faster (more slowly) than exponentially as T increases.

¹² Eqn. (1) implies that marginal utility become infinite as consumption goes to zero. As a result, a fat-tailed distribution for T can yield a WTP of 100% because the high probability of a sufficiently high T and resulting high probability of a sufficiently low C can imply an expected marginal utility of consumption that is infinite. This result disappears if we bound marginal utility at some maximum value.

¹³ The distribution is

of a temperature *decline* as large as 1.1°C, which is consistent with the studies surveyed by the IPCC. In addition to an expected temperature change of 3°C, it implies a standard deviation of 2.1°C for the temperature change. Note that this distribution (as well as the distributions discussed below) applies to the temperature increase as of 2100.

Figure 1 also shows two other distributions for T – a Frechet (also called a Generalized Extreme Value, Type 2) distribution, and a distribution labeled Roe-Baker. This last one (not a standard probability distribution) was developed by Roe and Baker (2007) from a simple climate model with uncertain feedback effects. I chose parameter values for these two distributions to match the mean, standard deviation, and minimum point of the gamma distribution.¹⁴ (It appears from the graph that the Frechet and Roe-Baker distributions begin at values of T above 0, but in fact the distributions are positive, but extremely small, for values of T as low as -1° .) The Frechet distribution is “somewhat” fat-tailed; for the parameter values I have chosen, the distribution has a mean and variance but no higher moments. The Roe-Baker distribution is extremely fat-tailed; none of the moments exist. (I calibrated the Roe-Baker distribution to have a *calculated* mean of 3° and standard deviation of 2.1° for range of T up to 50° .)

All three of these distributions seem like credible descriptions of the likelihood of future temperature realizations, especially since we know very little about the likelihood of temperatures above 6 or 8°C. Depending on one’s priors, one distribution might seem more credible than others; the Frechet and Roe-Baker distributions, for example, have much more mass between 1° and 4° than does the gamma distribution, and the Roe-Baker distribution has a theoretical grounding. But note from Figure 2 that the tails of these distributions look quite

$$f(T; r, \lambda, \theta) = \frac{\lambda^r}{\Gamma(r)} (T - \theta)^{r-1} e^{-\lambda(T-\theta)}, T \geq \theta$$

where $\Gamma(r)$ is the gamma function, $r = 3.8$, $\lambda = 0.92$, and $\theta = -1.13$.

¹⁴ The Frechet distribution is given by

$$f(T; k, \mu, \sigma) = (1/\sigma) \exp[-(1+kz)^{-1/k}] (1+kz)^{-1-1/k}$$

where $z = (T - \mu)/\sigma$, $k > 0$, and $T \geq \mu - \sigma/k$. The parameter values are $k = 0.28$, $\mu = 2.15$, and $\sigma = 0.915$. The Roe-Baker distribution is given by

$$g(T; \bar{f}, \sigma_f, \theta) = \frac{1}{\sigma_f \sqrt{2\pi z^2}} \exp \left[-\frac{1}{2} \left(\frac{1 - \bar{f} - 1/z}{\sigma_f} \right)^2 \right]$$

where $z = T + \theta$. The parameters values are $\bar{f} = 0.797$, $\sigma_f = .0441$, and $\theta = 2.13$. The feedback parameter in the model is normally distributed with mean and standard deviation $\bar{f}$ and σ_f respectively.

different. Compared to the gamma distribution, the Frechet and Roe-Baker distributions have much more mass at temperatures above 10° . But again, given how little we know about the likelihood of such high temperatures, it is hard to distinguish among these distributions and claim that one is more “correct” than the others.

3.2 Implications for Policy.

Might these three temperature distributions have very different policy implications? I address this using the framework in Pindyck (2012), but taking the parameter that relates T to the rate of GDP growth as fixed and equal to its mean value.¹⁵ This implies, consistent with IPCC (2007), a 3% loss of GDP in 2100 if T reaches 4°C . For each of the three temperature distributions, and for a range of values for the index of risk aversion, η , I calculate the WTP to ensure that T in 2100 will not exceed its mean value of 3°C .¹⁶ The results are shown in Figure 3.

Observe that the highest values of WTP are obtained when T is assumed to follow the (thin-tailed) gamma distribution. The Roe-Baker distribution, which is extremely fat-tailed, leads to a very low WTP for the entire range of η . The reason is evident from Figures 1 and 2. Compared to the gamma distribution, the Roe-Baker and Frechet distributions are “bunched up” in the range of 1° to 4° , with less mass in the higher temperature range of 5° to 8° . The standard deviations are the same because they have more mass at very high temperatures, but the probabilities for those very high temperatures are still very low. Thus fat tails need not imply a high WTP (a point discussed in detail in Pindyck (2011a)).

Perhaps the limited set of economic studies reviewed by the IPCC (2007) were overly optimistic regarding the impact of higher temperatures. To explore this, I recalculated WTP, but this time doubling the parameter that relates T to the rate of GDP growth (so that $T = 4^\circ\text{C}$ would result in a 6% loss of GDP in 2100, as opposed to a 3% loss). The results are shown in Figure 4. Not surprisingly, the WTP numbers are roughly double those in Figure 3. In fact for the gamma distribution, WTP is close to 10% of GDP if the index of risk aversion, η , is close to 1.

¹⁵ In Pindyck (2012), the real per capita GDP growth rate is $g_t = g_0 - \gamma T_t$, where γ is uncertain and follows a gamma distribution. In this exercise I focus on uncertainty over T and set γ equal to its mean value.

¹⁶ Of course I could have instead calculated the WTP to ensure T will remain zero or not exceed some other value. I chose 3°C for this exercise because a variety of studies focus on limiting temperature increases to 2°C or 3°C , and 3°C is the expected value for T in 2100.

What do the results in Figure 4 tell us? First, if we posit a large enough economic impact, for the right parameter values (and probability distribution for temperature) we can obtain a WTP that is large enough to be consistent with stringent abatement. But note that the parameter values needed to get this result are a bit extreme: a value of η below 1.5 and a rate of time preference, δ , equal to zero. (Although not shown in the figures, WTP falls dramatically if $\delta = .01$ or $.02$.) Unless one can make a strong argument for a robust case for $\delta = 0$ and $\eta < 1.5$, these results do not provide a good case for stringent abatement.

Perhaps we need to consider the possibility of even larger economic impacts, especially for temperature increases above 5°C , which, though unlikely, are certainly possible. The problem is that we know so little about the possible economic impacts of this degree of global warming. Unlike with temperature itself, it is difficult to even come up with probability distributions.

3.3 Impact of Climate Change: The Unknown and the Unknowable.

Why is it so difficult to estimate how climate change will affect the economy? One problem is that we have very little data on which to base empirical work. True, there have been a few empirical studies that made use of temperature and rainfall data for a large panel of countries over 50 or more years.¹⁷ Those studies have helped to confirm that the impact of higher temperatures is largely on the *growth rate* of GDP, as opposed to the *level* of GDP. But the size of the impact is imprecisely measured, and only applies to small changes in temperature, not the 5° or more of warming that many people worry about. The same is true for studies of the effects of temperature on agricultural output.

Second, there is little or nothing in the way of economic theory that could help us understand the potential impact of higher temperatures. We have some sense of how higher temperatures might affect agriculture, but we also know that losses of agricultural output in some regions of the world (e.g., near the Equator) might be matched by increased output in other regions (e.g., northern Canada and Russia). Furthermore, agriculture is a small fraction of total economic output: 1 to 2% of GDP for industrialized countries, 3 to 20% of GDP for developing countries. Beyond agriculture, it is difficult even at a heuristic level to explain how higher

¹⁷ See Dell, Jones, and Olken (2009, 2012), and Bansal and Ochoa (2011a, b).

temperatures will affect economic activity. That is why the loss functions that relate temperature to GDP in IAMs are essentially ad hoc, with little or no theoretical grounding.

Third, climate change will occur slowly, so that there is considerable potential for adaptation. In the case of agriculture, we have already seen substantial adaptation in the U.S. during the 19th century as settlers moved west and had to adapt crops to new and very different climatic conditions. (The recent book edited by Libecap and Steckel (2011) provides several detailed examples of this kind of adaptation.) Flooding is a potential hazard of climate change if sea levels rise substantially, but here, too, we have seen adaptation in the past (with the dikes of Holland perhaps the best-known example). This does not mean, however, that adaptation will eliminate the impact of climate change – it is simply another complicating factor that makes it very difficult to estimate any kind of loss function.

It may be that the relationship between temperature and the economy is not just something we don't know, but something that we *cannot* know, at least for the time horizon relevant to the design and evaluation of climate policy. Some researchers have come to the conclusion that climate sensitivity is in this category of the “unknowable.”¹⁸ Yet, for the reasons given above, the impact of climate change is far less “knowable” than climate sensitivity. If so, then we will never (or at least over any relevant time period) reach a consensus on the question posed at the beginning of this paper: Should environmental economists push for the early adoption of a stringent GHG abatement policy?

Note that this does not mean the answer to this question is no, and that there is no case to be made for stringent GHG abatement. But I don't think the case can be made by applying Monte Carlo simulation methods to one or more integrated assessment models, or by calculating WTP based on “consensus” probability distributions as I have done earlier. The case would have to be based on the possibility of a *catastrophic outcome*, something that is far outside the realm of these models and probability distributions.

4. Catastrophic Climate Change.

For some environmentalists, without stringent abatement, a climate catastrophe is not just a possibility, but almost a sure thing. In a recent *New York Times* opinion piece, for example,

¹⁸ See, for example, Allen and Frame (2007), who show how the uncertain feedback effects that are central to the Roe-Baker model imply that we will never be able to determine an upper bound for climate sensitivity.

James Hansen (2012) claims that if Canada extracts the oil from its reserves of tar sands “and we do nothing, it will be game over for the climate.” Exactly what is meant by “game over” is unclear, but it sounds catastrophic. It is certainly not something embodied in any integrated assessment model I am aware of. While a catastrophic climate outcome is indeed a possibility (with or without Canadian tar sands oil), it is something we know very little about.

Is there some way to bring economic analysis to bear on the policy implications of possible catastrophic outcomes? For climate scientists, catastrophic outcomes almost always take the form of high temperature outcomes, e.g., a 7° or 8° increase by 2100. Putting aside the difficulty of estimating the probability of such a large temperature increase, what matters in the end is not the temperature increase itself, but rather the impact it would have. Would that impact be “catastrophic,” and if so, in what sense? And might a smaller (and more likely) temperature increase also have a catastrophic impact?

My calculations of WTP in the preceding section were based on full distributions of temperature outcomes, but for all three of the distributions I considered, the probabilities for very high temperatures are low. However, we saw in Figure 4 that by doubling the assumed impact of temperature on GDP growth, we obtained a higher WTP. This is no surprise; we could increase the impact even more (e.g., so 4° of warming would result in a 12% loss of GDP in 2100), and get a still higher WTP. The problem, as explained above, is that we have very little basis – empirical or theoretical – for determining the actual impact, or even the overall functional relationship between temperature and GDP or GDP growth.

Furthermore, it is difficult to see how our knowledge of the economic impact of rising temperatures is likely to improve in the coming years. Even more so than temperature change itself, economic impact may simply be in the realm of the “unknowable.” If so, it would make little sense to try to use an IAM-based analysis to make the case for stringent abatement. That case would have to be based on the likelihood – even if it is small – of a catastrophic outcome in which climate change is sufficiently extreme to cause a very substantial drop in welfare.

4.1. How Likely? How Extreme?

How should we think about catastrophic climate outcomes? Given how little we know, it seems to me that a very detailed and complex modeling exercise is unlikely to be helpful. (Even if we really believed the model accurately represented the relevant physical and economic

relationships, which is unlikely, we would have to come to agreement on key parameters, such as the rate of time preference.) Probably something simpler is needed – although not so simple as the type of claims made by Hansen (2012) and others that a catastrophe is inevitable. Perhaps the best we can do is come up with rough, subjective estimates of the probability of a climate change sufficiently large to have a catastrophic impact, and then some distribution for the size of that impact (in terms, say, of a reduction in GDP, or a reduction in the effective capital stock or the productivity of the capital stock).

This is the approach that has been used in recent studies of “consumption disasters,” defined as events that caused consumption to decline by some substantial amount (say, more than 10%). In much of that work, disasters are modeled as Poisson arrivals, the impact of which is a random percentage reduction in the capital stock (and thus in ongoing consumption), with the loss fraction given by a simple distribution such as a one-parameter power distribution. The mean Poisson arrival rate and the parameter of the impact distribution might be estimated from consumption data for a sample of countries over a century or more of time, or inferred from the behavior of macroeconomic and financial aggregates.¹⁹ For climate change, however, a catastrophic outcome has not yet occurred, so estimating impact parameters from panel data or from macroeconomic aggregates is not an option.²⁰

The problem is analogous to assessing the world’s greatest catastrophic risk during the Cold War – the possibility of a thermonuclear exchange between the U.S. and the Soviet Union. What was the likelihood of such an event? There were no data or reliable models that could yield good estimates. Instead, analyses had to be based on the *plausible*, i.e., on events that could reasonably be expected to play out, even if with low probability. The same approach had to be taken with respect to assessing the range of potential impacts of a thermonuclear exchange. Such analyses were useful because they helped evaluate the potential benefits (and risks) of arms control agreements.

It seems to me that the same approach can be used to assess climate change catastrophes. We could begin by asking what is a plausible range of catastrophic outcomes (under, for

¹⁹ Barro (2009) and Barro and Jin (2009) are examples of studies in which arrival rates and impact distributions are estimated from panel data. Pindyck and Wang (2012) estimate the relevant parameters (including the IRRA) as calibration outputs from a general equilibrium model.

²⁰ Although Dell, Jones, and Olken (2009, 2012), and Bansal and Ochoa (2011a, b) estimated the impact of temperature change on economic growth using historical data for a panel of countries, both the temperature variations and changes in growth rates were small and nothing close to what might be considered “catastrophic.”

example, BAU), as measured by a percentage declines in the stock of productive capital (thereby reducing future GDP over time). That range could be discrete (e.g., three or more potential outcomes) or continuous. Next, what are plausible probabilities? Here, “plausible” would mean acceptable to a range of economists and climate scientists. Given these plausible outcomes and probabilities, one can calculate the WTP to avert those outcomes, or to reduce the probabilities of their occurrence. That WTP will once again depend on preference parameters, i.e., the rate of time preference and index of risk aversion. But if the WTP is sufficiently large and is robust to reasonable ranges for those parameters, it might provide support for a stringent policy.

This approach does not carry the perceived precision that is part of an IAM-based analysis. But that perceived precision is likely to be illusory. To the extent that we are dealing with unknowable quantities, it may be that the best we can do is rely on the “plausible.”

4.2. Multiple Potential Catastrophes.

So suppose an analysis of climate change catastrophes based on “plausible” outcomes and probabilities indicates a high WTP, something around 10% of GDP, for a reasonable range of preference parameters. Are we home yet? Would we then have what we need to make a strong case for a stringent abatement policy?

Maybe not. The problem is that we must consider other potential catastrophic events. A climate catastrophe is only one of a number of potential catastrophes that could cause major damage on a global scale. Different people have different nightmares, but my list would 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 (perhaps related to toxic waste or severe water shortages).²¹ 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.

The approach outlined above can also be used to assess other possible catastrophes and evaluate policies to avert them. An example is nuclear terrorism. Various studies have assessed

²¹ For those readers lacking imagination, see Posner (2004) and Bostrom and Ćirković (2008). For a sobering discussion of the likelihood and possible impact of nuclear terrorism, see Allison (2004).

the likelihood and potential impact of the detonation of one or more nuclear weapons (with the yield of the Hiroshima bomb) in major cities. “Plausible” estimates of the probability of this occurring in the next ten years vary. Allison (2004) puts the probability at 50%; others put it at 1 to 5%. What are plausible estimates of the impact? A million or more deaths, along with a shock to the capital stock and GDP resulting from reduced trade and economic activity worldwide, as vast resources are devoted to averting further events. (For a discussion and survey of studies of nuclear terrorism, see Ackerman and Potter (2008).)

Suppose the WTP to avert or reduce the probability of a nuclear catastrophe of this sort is 10% of GDP. How would this affect the analysis of a climate catastrophe? Suppose that with no other potential catastrophes, the WTP to avoid a climate catastrophe is also 10% of GDP. And suppose there are also five or six other potential catastrophes (a mega-virus, a severe water shortage, ... you can come up with your own list), each with a WTP of 10%. How will the WTP for averting a climate catastrophe change once we take into account the other potential catastrophes? It is likely to go down, as will the individual WTPs for the other catastrophes. The reason is that the WTPs are not additive; society would probably be unwilling to spend 60 or 70% of GDP to avert all of these catastrophes. Thus, when taken as a group, the WTP for each potential catastrophe (including climate) would fall.

To understand this, note that the inclusion of other potential catastrophes would affect the WTP for climate in two opposing ways. 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 increases the benefit of avoiding the climate catastrophe. On the other hand, assuming all of these potential catastrophes are equally threatening, 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.

Unfortunately, society faces a number of potentially catastrophic threats, and opinions will differ as to how large that number is, and which threats are worse than others. Climate change is probably one of those threats, but it is unclear where it would rank in a list that was ordered by likelihood and potential severity. The point is that when evaluating the expenditure of scarce resources to reduce the likelihood of a climate catastrophe, these other catastrophic threats should not be ignored.

5. Conclusions.

I began by asking whether environmental economists should push for the early adoption of a stringent GHG abatement policy. I have argued that the *economic* case for stringent GHG abatement cannot be made based on “most likely outcomes,” i.e., distributions for temperature change consistent with the IPCC (2007), and the economic impact functions used in most IAMs (and were also surveyed by the IPCC). Also, it doesn’t matter whether the temperature distributions are fat- or thin-tailed. Even if we assume that moderate temperature change (say 4°C) would have a large economic impact (say a 6% loss of GDP), the willingness to pay for a stringent policy will be small for most “reasonable” values of key parameters. For example, it is hard to justify even a moderate abatement policy if the rate of time preference is 1 or 2% rather than zero, or if the index of risk aversion is 2 or greater. Some have argued that “ethics” should dictate these parameter values; a rate of time preference of zero might reflect the ethical values of some economists, but there is no evidence I am aware of that it reflects the preferences of society as a whole.

I have also argued that the greatest area of uncertainty is over the economic impact of climate change. (Again, “economic impact” includes health and social impacts.) The economic loss functions that are part of most IAMs are essentially ad hoc. This is not surprising given how little we know – in terms of both theory and data – about the ways and extent to which changes in temperature and other climate variables are likely to affect the economy. In fact, the economic impact of climate change may well be in the realm of the “unknowable.” This in turn means that IAM-based analyses of climate change may not take us very far, and the models may be of very limited use as a policy tool.

By utilizing WTP, I have avoided having to deal with the cost (or “supply”) side of climate policy, and focus instead on the “demand” side (in that WTP is society’s reservation price for achieving a policy objective). But the cost side is indeed relevant. If it can be demonstrated that a policy objective (e.g., limiting T in 2100 to 3°C or less) can be achieved at low cost – say 1% of GDP – then a WTP of only 2% is more than sufficient to justify the policy. At this point, estimates of cost are not that small, but that may change as better technologies for reducing emissions or sequestering carbon become available.

The arguments raised in this paper imply that any case for stringent abatement must be based on the possibility of a catastrophic climate outcome. And a “catastrophic climate

outcome” does not simply mean a very high increase in temperature and rising sea levels; what matters is whether the economic impact of those physical changes would be catastrophic. Furthermore, simply stating that such an outcome is possible is not enough. We would need “plausible” estimates of probabilities of various temperature outcomes, and “plausible” estimates of large economic impacts from those temperature outcomes. The claims of likely catastrophic outcomes that I have seen (made largely by climate scientists, not economists) relate to temperature and other physical processes (such as climate variability and changes in sea levels). Usually absent is a clear analysis showing that the economic impact of these physical changes would be so large as to be catastrophic. Completely absent (to my knowledge) is any analysis that puts a climate catastrophe in the context of a set of potential global catastrophes.

So should environmental economists push for stringent GHG abatement? In the end, I have not answered that question. Instead I have tried to explain why the question is so inherently difficult, why an answer won’t come from the kinds of modeling exercises that have permeated the literature, and why the case for stringent abatement – if that case is to be made at all – must be based on an analysis of potential catastrophic outcomes.

References

- Ackerman, Gary, and William C. Potter, "Catastrophic Nuclear Terrorism: A Preventable Peril," in *Global Catastrophic Risks*, N. Bostrom and M. Ćirković, eds., Oxford University Press, 2008.
- Allen, Myles R., and David J. Frame, "Call Off the Quest," *Science*, October 26, 2007, **318**, 582—583.
- Allison, Graham, *Nuclear Terrorism: The Ultimate Preventable Catastrophe*, Henry Holt & Company, 2004.
- Bansal, Ravi, and Marcelo Ochoa, "Welfare Costs of Long-Run Temperature Shifts," NBER Working Paper 17574, Nov. 2011a.
- Bansal, Ravi, and Marcelo Ochoa, "Temperature, Aggregate Risk, and Expected Returns," NBER Working Paper 17575, Nov. 2011b.
- Barro, Robert J., "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, March 2009, **99**, 243--64.
- Barro, Robert J., and Tao Jin, "On the Size Distribution of Macroeconomic Disasters," NBER Working Paper 15247, August 2009.
- Bostrom, Nick, and Milan Ćirković, eds., *Global Catastrophic Risks*, Oxford University Press, 2008.
- Dasgupta, Partha, "Discounting Climate Change," *Journal of Risk and Uncertainty*, 2008, **37**, 141—169.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates," *American Economic Review*, May 2009, **99**, 198—204.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "Climate Change and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics*, 2012, **4**, 66—95.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, June 2002, **40**, 351—401.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton, "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation," NBER Working Paper No. 16913, March 2011.

- Hansen, James, "Game Over for the Climate," *New York Times*, May 9, 2012.
- Heal, Geoffrey, "Climate Economics: A Meta-Review and Some Suggestions," *Review of Environmental Economics and Policy*, 2009, **3**, 4—21.
- Heal, Geoffrey, and Bengt Kriström, "Uncertainty and Climate Change," *Environmental and Resource Economics*, June 2002, **22**, 3—39.
- Hope, C.W., "The Marginal Impact of CO₂ from PAGE 2002: An Integrated Assessment Model Incorporating the IPCC's Five Reasons for Concern," *Integrated Assessment*, 2006, **6**, 1—16.
- Interagency Working Group on Social Cost of Carbon, "Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis," February 2010.
- Intergovernmental Panel on Climate Change (IPCC), "Climate Change 2007: The Physical Science Basis," Cambridge University Press, 2007.
- Jaffe, Adam B., Steven R. Peterson, Paul R. Portney, and Robert N. Stavins, "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us?" *Journal of Economic Literature*, March 1995, **33**, 132—163.
- Libecap, Gary D., and Richard H. Steckel, editors, *The Economics of Climate Change: Adaptations Past and Present*, University of Chicago Press, 2011.
- Mendelsohn, Robert, "Is the *Stern Review* an Economic Analysis?" *Review of Environmental Economics and Policy*, Winter 2008, **2**, 45—60.
- Newbold, Stephen C., and Adam Daigneault, "Climate Response Uncertainty and the Benefits of Greenhouse Gas Emissions Reductions," *Environmental and Resource Economics*, 2010, **44**, 351—377.
- Nordhaus, William D., "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, Sept. 2007, 686—702.
- Nordhaus, William D., *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, 2008.
- Nordhaus, William, D., and J.G. Boyer, *Warming the World: Economic Models of Global Warming*, M.I.T. Press, 2000.
- Nordhaus, William D., "Estimates of the Social Cost of Carbon: Background and Results from the RICE-2011 Model," NBER Working Paper 17540, October 2011.

- Pindyck, Robert S., "Uncertainty in Environmental Economics," *Review of Environmental Economics and Policy*, Winter 2007.
- Pindyck, Robert S., "Fat Tails, Thin Tails, and Climate Change Policy" *Review of Environmental Economics and Policy*, Summer 2011a.
- Pindyck, Robert S., "Modeling the Impact of Warming in Climate Change Economics," in G. Libecap and R. Steckel, Eds., *The Economics of Climate Change: Adaptations Past and Present*, University of Chicago Press, 2011b.
- Pindyck, Robert S., "Uncertain Outcomes and Climate Change Policy," *Journal of Environmental Economics and Management*, February 2012.
- Pindyck, Robert S., and Neng Wang, "The Economic and Policy Consequences of Catastrophes," NBER Working Paper No. 15373, revised March 2012.
- 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.
- Tol, R., "Estimates of the Damage Costs of Climate Change, Part I: Benchmark Estimates," *Environmental and Resource Economics*, 2002, **22**, 47—73.
- Weitzman, Martin L., "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature*, Sept. 2007, 703—724.
- Weitzman, Martin L., "Some Dynamic Economic Consequences of the Climate Sensitivity Inference Dilemma," unpublished manuscript, February 2008.
- Weitzman, Martin L., "On Modeling and Interpreting the Economics of Climate Change," *Review of Economics and Statistics*, February 2009, **91**, 1—19.
- Weitzman, Martin L., "Fat-Tailed Uncertainty and the Economics of Climate Change," *Review of Environmental Economics and Policy*, Summer 2011.




Figure 1: Three Temperature Distributions. All three distributions were calibrated to have the same mean (3°), standard deviation (2.1°), and minimum possible temperature change (-1.1°).




Figure 2: Upper Tails of Temperature Distributions.




Figure 3: WTP to Keep T in 2100 Below 3°C . The three distributions (gamma, Frechet, and Roe-Baker) are calibrated to have the same mean (3°), standard deviation (2.1°), and minimum possible temperature change (-1.1°). Initial GDP growth rate, g_0 , is .02, the rate of time preference, δ , is zero, and the growth impact parameter, γ , equals its mean value of .000136.




Figure 4: WTP to Keep T in 2100 Below 3°C , with Growth Impact Doubled. The three distributions have the same mean (3°), standard deviation (2.1°), and minimum possible temperature change (-1.1°). Initial GDP growth rate, g_0 , is .02, and rate of time preference, δ , is zero. The growth impact parameter, γ , equals twice its mean value, i.e., $\gamma = .000272$.