WORKSHOP PROCEEDINGS

Assessing the Benefits of Avoided Climate Change: Cost-Benefit Analysis and Beyond

Uncertainty and the Benefits of Climate Change Policies

Stephen C. Newbold

U.S. EPA, National Center for Environmental Economics

Adam Daigneault

U.S. EPA, Climate Change Division

May 2010

PEW CENTER ON Global CLIMATE CHANGE

This workshop was made possible through a generous grant from the Energy Foundation.

Energy Foundation
301 Battery St.
San Francisco, CA 94111

Workshop Speakers

David Anthoff, Eileen Claussen, Kristie Ebi, Chris Hope, Richard Howarth, Anthony Janetos, Dina Kruger, James Lester, Michael MacCracken, Michael Mastrandrea, Steve Newbold, Brian O'Neill, Jon O'Riordan, Christopher Pyke, Martha Roberts, Steve Rose, Joel Smith, Paul Watkiss, Gary Yohe

Project Directors

Steve Seidel
Janet Peace

Project Manager

Jay Gulledge

Production Editor

L. Jeremy Richardson

Content Editors

Jay Gulledge, L. Jeremy Richardson, Liwayway Adkins, Steve Seidel

Suggested Citation

Newbold, S. C. and A. Daigneault. 2010. "Uncertainty and the Benefits of Climate Change Policies." In *Assessing the Benefits of Avoided Climate Change: Cost-Benefit Analysis and Beyond.* Gulledge, J., L. J. Richardson, L. Adkins, and S. Seidel (eds.), Proceedings of Workshop on Assessing the Benefits of Avoided Climate Change, March 16–17, 2009. Pew Center on Global Climate Change: Arlington, VA. p. 183–197. Available at: http://www.pewclimate.org/events/2009/benefitsworkshop.

The complete workshop proceedings, including video of 17 expert presentations, this summary report, and individual off-prints of expert papers are available free of charge from the Pew Center on Global Climate Change at http://www.pewclimate.org/events/2009/benefitsworkshop.

May 2010

Pew Center on Global Climate Change
2101 Wilson Blvd., Suite 550
Arlington, VA 22201

Uncertainty and the Benefits of Climate Change Policies[1]

Stephen C. Newbold
U.S. EPA, National Center for Environmental Economics

Adam Daigneault
U.S. EPA, Climate Change Division

Abstract

This paper discusses the role of uncertainty in estimating the economic benefits of greenhouse gas emission reductions. First, we give a general overview of the range of approaches that analysts can use to account for uncertainty in benefit-cost analyses of climate change policies, and we discuss how to account for the "value of insurance" that a policy provides against potential climate catastrophes. A simple numerical example (given in an appendix) shows that uncertainty can in principle have a large influence on estimates of economic benefits. We then review some of the recent research by climate change economists that has begun to quantify this influence. We also give suggestions for short, medium, and longer term research. In the short and medium term, we recommend further synthesizing the recent research on the effects of risk and uncertainty on the benefits of climate policies and improving the currently available integrated assessment models to better account for these factors. In the longer term, we recommend expanding these models or developing new ones to incorporate the effects of learning, policy flexibility, and the value of additional information on the response of the climate system to greenhouse gas emissions and the economic consequences of the resulting climate changes.

[1] The views expressed in this paper are those of the authors and do not necessarily represent those of the U.S. EPA. No Agency endorsement should be inferred

Introduction

Virtually all public policy decisions must be made in the face of uncertainty, and—at the risk of understatement—regulations to address climate change are no exception. Uncertainty can take many forms and can have different implications for the optimal stringency and structure of a policy, depending on the specifics of each case. The net effect depends on a variety of factors, including the relative risks of acting now versus waiting for more information, the potential for and costs of learning more about the impacts of the policy over time, and irreversibilities associated with ecosystem thresholds or mitigation activities (i.e. sunk benefits and costs [Pindyck 2000]). In the final analysis, uncertainty may increase or decrease the optimal stringency of a policy or weigh more heavily in favor of one type of instrument over others (such as cap and trade versus taxes), depending on the balance of these sometimes competing factors.

In this paper we address only a small part of this larger picture. Specifically, we focus on the effect of uncertainty on estimates of economic benefits of greenhouse gas emission reductions. First, we review the range of approaches that analysts can use to account for uncertainty in benefit-cost analysis. We perform some simple numerical calculations using a highly stylized model to show how uncertainty can influence the estimated benefits of climate policies, and we show how an expected utility framework can account for the "value of insurance" that a policy provides against potential climate catastrophes. Next, we review some recent research that has examined the effect of uncertainty on emissions reduction benefits, including our own work on climate response uncertainty and the shape of the damage function. We conclude with several recommendations for further research.

This paper is written for analysts, researchers, and especially managers and decision-makers who need to interpret and use the results of economic assessments in their deliberations over new climate change policies.

Tiers of Uncertainty Analysis

We begin by reviewing the range of approaches for addressing uncertainty in benefit-cost analysis in general and as applied to climate change policies in particular. This discussion loosely follows that in the Office of Management and Budget's Circular A-4 (OMB, 2003, p 41-42), though we elaborate further on the implications of conducting a formal uncertainty analysis in an expected utility framework.

The easiest approach for dealing with uncertainty is to simply describe it qualitatively, without addressing it explicitly in the quantitative analysis. More generously, we might say that the analyst can "average out the uncertainty" before estimating benefits and costs by plugging best-guess central point estimates of all uncertain parameters into the economic model. In doing this, the analyst tacitly accepts that the resulting point estimate of net benefits is only one among many possible outcomes. If the analysis ends here, the results are effectively treated as central best-guess estimates themselves. This approach is fairly common and will give an accurate estimate of the expected net benefits when the benefit function is (at least approximately) linear over the relevant ranges of all uncertain parameters.

If benefits are sufficiently non-linear over the relevant ranges, then the deterministic estimate may not be robust to the uncertainty in the input parameters. The next step, then, might be to conduct a sensitivity analysis, where the analyst varies each parameter over what are thought to be plausible ranges, based on the relevant scientific and economic research, and then records the effect of these variations on the net benefits. This provides a simple means of examining the response of the model to the key assumptions and often is useful for illustrating the importance of uncertainty to decision-makers and other consumers of the benefit-cost analysis. However, the more parameters that are varied at one time, the more difficult it is to interpret the results. Furthermore, the range of variations in model outputs illustrated in a sensitivity analysis may give little indication of their central tendency based on the relative likelihood of the many possible combinations of input parameters.

So the next logical step is to account for the uncertainty in all input parameters simultaneously. Known as Monte Carlo analysis, this can be done by specifying probability distributions for each parameter and then using computer simulation methods to construct a probability distribution for the estimated benefits.

It may seem that this is the final possible step in the progression of uncertainty analysis. However, it is possible to go further by framing the overall policy question in an expected utility framework. Under this approach, the analysis is structured to directly answer the question: Given all of the uncertainties regarding the input parameters and other assumptions of the model, what is the change in aggregate income with the policy that would make society just as well off as without the policy? In other words, what is the maximum amount of income society is willing to pay for the policy? In this approach, the analyst integrates over all sources of uncertainty within the economic model itself. The uncertainty is not "averaged out" before the parameters are plugged into the model, and the analyst does not simply construct a probability distribution for willingness to pay.

One key advantage of the expected utility approach is that it provides a natural way to account for potential low-probability high-impact outcomes. In effect, this framework can account for the value of the insurance that a policy would provide against the worst-case scenarios. This is an important consideration when analyzing greenhouse gas (GHG) emission reduction policies, since the potential for "climate catastrophes" is a key motivating factor for many citizens and decision-makers concerned about climate change (Keller et al., 2004; Hansen et al., 2007; Ramanathan and Feng, 2008). An evaluation framework that ignored this aspect of the problem would seem to be missing something essential.

A concrete illustration of the distinctions between the tiers of uncertainty analysis described above using a simple numerical example is provided in the appendix. The example shows that, under the typical assumption that climate change damages increase with temperature at an increasing rate, the deterministic analysis gives the lowest estimate of willingness to pay (WTP), the Monte Carlo approach gives a higher estimate of average WTP, and the expected utility approach gives the highest estimate. The magnitude of this "risk premium" will depend on both the level of uncertainty in the input parameters and the degree of risk aversion that is assumed.[2] Also, as emphasized by Weitzman (2009), the risk premium will depend crucially on the severity and probability of the worst-case outcomes.[3]

In light of the above, consider the Office of Management and Budget's (OMB's) guidelines. Circular A-4 indicates that the default assumption in a benefit-cost analysis should be one of risk neutrality. Specifically,

> "Emphasis on [expected values of benefits and costs] is appropriate as long as society is 'risk neutral' with respect to the regulatory alternatives. While this may not always be the case, you should in general assume 'risk neutrality' in your analysis. If you adopt a different assumption on risk preference, you should explain your reasons for doing so." (OMB, 2003, p 42).

This makes good sense for regulations that lead to small changes in risks. Even a risk averse individual would evaluate small risks based solely on their expected values as long as the risks are uncorrelated with the individual's income, since in this case the benefit function is approximately linear.[4] In contrast, the expected utility framework described in the preceding paragraphs and in the appendix explicitly assumes that society is not necessarily risk neutral with respect to climate change policies. The basic rationale is two-fold: 1) since the potential impacts of climate change are wide-spread—potentially global in scope, especially considering the worst-case catastrophic scenarios—the risks may be very large, and 2) the very high correlation among individual risks means that an effective risk-sharing arrangement is not possible (Arrow and Lind, 1970; Dasgupta and Heal, 1979, Ch 13). In other words, if the worst outcomes do come to pass then we may all be significantly impacted simultaneously, so there would be far less scope for spreading the risks. Furthermore, and on a more practical level, if the

[2] In this paper we use the term "risk premium" to refer to the difference between estimates of willingness to pay based on an expected utility framework that explicitly accounts for parameter uncertainty and risk aversion and analogous estimates of willingness to pay based on a deterministic model that ignores uncertainty and risk aversion. This should not be confused with the "risk premium" in the finance literature that refers to the interest rate mark-up associated with risky investments.

[3] We should note that, as in all integrated assessment models of which we are aware, both the simple example given in the appendix and the simulation experiments in our previous work (Newbold and Daigneault 2009) ignore any potentially catastrophic risks of *reducing* GHG emissions. Such risks could arise, for example, from the possibility that elevated atmospheric stocks of GHGs could forestall a natural trend of decreasing global temperatures and therefore another ice age in the future (e.g., Ruddiman 2005). While such a scenario may be highly unlikely (very low probability), it may not be completely implausible (zero probability). If so, and if the damages from such a scenario also could be catastrophic, then a complete uncertainty analysis would include these potentially countervailing risks as well.

[4] A person is risk neutral if they are indifferent between prospects with the same expected returns, regardless of the variance of the possible outcomes. A person is risk averse if, of multiple prospects with the same expected returns, they prefer the one with the lowest variance in the possible outcomes. The relevance of the correlation between the riskiness of the prospect and the individual's income can be understood by imagining a case where the prospect is more likely to pay off high when the individual's income is lower (higher) than normal. In this case, the individual would be willing to pay more (less) for the prospect, all else equal. See Dasgupta and Heal (1979 Ch 13) for a more complete exposition.

changes in risks are in fact small then the expected utility framework will collapse to the equivalent of a risk-neutral analysis anyway.[5]

So far we have argued—and the appendix has illustrated—that uncertainty *can in principle* have a strong influence on the estimates of benefits for climate change policies. Next, we discuss some recent research that has begun to quantify these effects using economic integrated assessment models (IAMs).

Previous Research

A few recent studies have used Monte Carlo analysis or similar methods to account for uncertainty in economic climate assessment models, but so far the results have been decidedly mixed. For example, Roughgarden and Schneider (1999) constructed probability distributions over parameters of the damage function in DICE using results from a survey of experts and found that the average optimal carbon tax from a Monte Carlo simulation was around eight times higher than the point estimate from the standard DICE model. Pizer (1999) used a modified version of DICE and found that accounting for parameter uncertainty increased the estimated welfare gain from an optimal tax rate policy by roughly 25 percent compared to its deterministic counterpart. Tol (2003) used the FUND model and found that when accounting for uncertainty "the net present marginal benefits of greenhouse gas emission reduction becomes very large" and in one scenario appeared to be unbounded. Ceronsky et al. (2005) also used FUND and found "that incorporating [potential climate catastrophes] can increase the social cost of carbon [SCC] by a factor of 20." Hope (2006) used Monte Carlo analysis and found that the 5th percentile, mean, and 95th percentile of the probability distribution for the SCC were, respectively, $4, $19, and $51 per ton of carbon. Uncertainty in the climate sensitivity parameter made the largest contribution to the variance of the SCC estimates. Nordhaus (2008) conducted an uncertainty analysis using the DICE model and concluded that "the best-guess policy is a good approximation to the expected-value policy." Weitzman (2009) showed that if the climate sensitivity distribution has a "fat-tail"—in other words, if the probability of ever higher temperature changes does not decline faster than the rate at which damages increase with temperature—then there is no bound on the willingness to pay for emissions reductions. And finally, Pindyck (2009) used a thin-tailed gamma distribution, including some versions with a significant right skew, but in most cases found only a modest risk premium.

In our own recent research, we focused on the effect of climate response uncertainty on estimates of economic benefits of GHG emissions reductions (Newbold and Daigneault, 2009). Specifically, we used Bayesian updating and model averaging to construct alternative probability distributions over the climate sensitivity parameter, which determines the equilibrium change in average global temperature to a doubling of the atmospheric greenhouse gas concentration (Andronova et al., 2007). We combined 28 confidence intervals for the climate sensitivity

[5] The reader can use the model in the appendix to confirm this by assuming that the probability of a 1 degree temperature change is 100 percent. In that case, the estimates of willingness to pay differ by 0.1 percent or less.

parameter reported in 21 studies. Figure 1 shows the estimated probability distributions from each study and the two alternative composite distributions that we used to calculate willingness to pay in both a deterministic model and an expected utility model that incorporated uncertainty.

Figure 1. Roe and Baker (2007) probability distributions constructed from the 5th and 95th percentiles for the climate sensitivity parameter from 21 different studies (light dotted lines), the Bayesian model-averaged probability distribution function based on the average of the distributions (heavy solid line), and the Bayseian updated pdf based on the product of the distributions (heavy dotted line). From Newbold and Daigneault (2009).

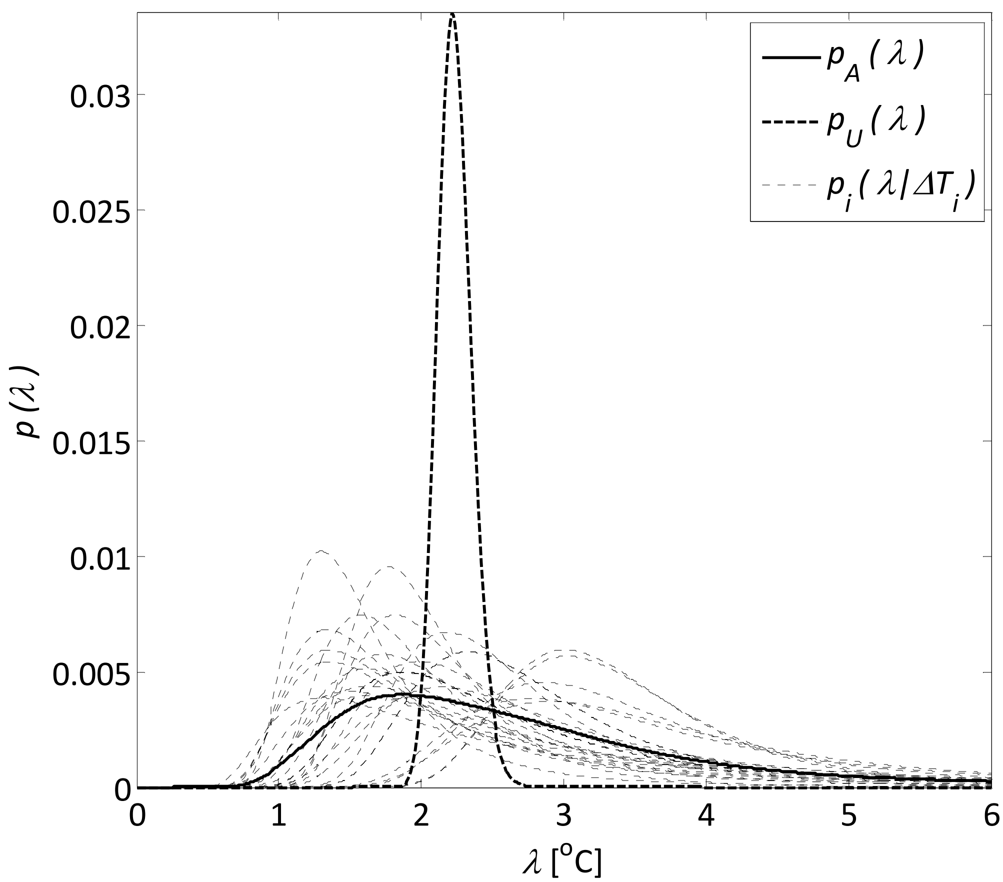

The distribution constructed using Bayesian updating is centered around 2.2°C and is very narrow, so carrying this through the expected utility model gives results very close to the deterministic estimates of willingness to pay. However, this composite distribution is based on what seems like an overly-optimistic view of the climate science literature. It effectively assumes that the studies we combined can be treated as independent estimates using *new data* but the *same underlying model* of how the climate system works. So we also considered an alternative assumption, that these studies effectively used the *same underlying data* but a *different model* of how the climate system works, i.e., we combined the estimates using a "model averaging" approach. This assumption gives a much wider distribution for the climate sensitivity parameter.[6] Carrying this distribution through the expected utility model can give very large risk premiums, depending on the other parameter values. For example, we found that by using the Bayesian model-averaged composite distribution and an exponential damage function, the risk-adjusted willingness to pay for emissions reductions consistent with the optimal path from the DICE model (Nordhaus, 2008) was nearly five times larger than the deterministic willingness to pay.

One important take-home message from this research is the following: because IAMs that account for uncertainty can produce such a wide range of benefits estimates, it is crucial for decision-makers to understand the key ingredients of any integrated assessment model when interpreting its results. Until recently, much of the discussion in the literature on the economics of climate policy has focused on the "usual suspects," namely the discount rate and the expected damages at the central estimates of future temperatures. However, the simulation experiments described in detail in our previous work (Newbold and Daigneault, 2009) and in short form in the appendix suggest that part of the explanation for the divergent results summarized above may lie in the (possibly subtle) differences between the way each study characterized the climate response uncertainty and the shape of the damage function at high temperatures. Specifically, in addition to the usual suspects, we would emphasize the coefficient of relative risk aversion (see the appendix) and the magnitude and probability of the worst-case scenarios as important members of the short list of parameters likely to have the largest influence on the benefits estimates.[7]

Conclusions and Recommendations

In this final section we respond directly to the stated objectives of the Pew Center workshop that provided the occasion for this paper. Those objectives were "to develop a set of practical recommendations that decision makers can employ in the near-term, and to outline a research path to improve decision making tools over time." The recommendations we offer below are aimed mainly at researchers and analysts who develop and use integrated assessment models for the purpose of informing decision-makers in their deliberations over climate policies. These recommendations are based on our own current (and perhaps idiosyncratic) understanding of both the state of the art of climate policy benefits assessment and the needs of decision-makers. We will offer our suggestions in the form of short, medium, and longer term recommendations.

---

[6] If this assumption is overly pessimistic, it is perhaps only modestly so since it is broadly consistent with the summary provided in the latest IPCC report (Hegerl et al. 2007).

[7] Importantly, the magnitude and probability must be considered simultaneously. See Sunstein (2007) for a discussion of the errors in public decision-making that can arise from placing undue attention on worst-case scenarios or paying no attention to them at all. Therein, Sunstein proposes a "Catastrophic Harm Precautionary Principle," which calls for close attention to both the magnitude and probability of a harm and allows for a "margin of safety for certain large-scale harms... akin to a purchase of insurance. Whether the margin is worthwhile depends on what is lost and what is gained by insisting on it." In the expected utility framework, the coefficient of relative risk aversion is the extra ingredient that allows for a systematic determination of the margin of safety (and of course the cost side of the ledger, not addressed in this paper, accounts for what is lost).

First, in the very short term (within one year or so), two useful tasks would be (1) to further synthesize previous research on the social cost of carbon (SCC), along the lines of meta-analyses conducted by Tol (2005, 2008), and (2) to construct a simple and transparent model for calculating the global and domestic SCC. Our experience has been that one of the first hurdles in discussing the economics of climate change with decision-makers is merely explaining the meaning of the SCC itself. A simple model constructed from first principles could be used as a tool for communication with decision-makers—in particular, helping to explain the proper interpretation and use of SCC estimates in a policy setting. It also could be used to produce rough estimates of the SCC and conduct sensitivity analyses and bounding exercises given any range of input assumptions that the user deems plausible. (We have in mind something similar to the simple model created by Tol and Yohe (2009) to examine *The Stern Review*.)

Second, for the medium term (between one and two years or so), a useful task would be to develop an improved IAM suitable for regulatory analysis alongside the standard models that federal agencies such as the Environmental Protection Agency typically use to estimate the costs of climate policies (e.g., ADAGE[8] and IGEM[9]). Such a model should build on existing IAMs that have been widely used in the climate economics literature (e.g., DICE, FUND, PAGE), but it also should add extensions and elaborations as dictated by the evolving demands of decision-makers. These might include adding currently omitted categories of benefits, a probabilistic structure that is suitable for uncertainty analysis (as in PAGE), and a capacity to incorporate risk aversion explicitly. In the process of building such a model, clear documentation should be developed simultaneously. In our experience, the more the model looks like a "black box," the less weight decision-makers are able to place on its results.

Our principal motivation for recommending that IAMs be extended to account for uncertainty and risk aversion is that, as discussed above, making fuller use of the expected utility framework provides a natural way to account for the high-impact, low-probability outcomes that are of primary concern to many citizens and decision-makers. Importantly, this approach forces us to bring these issues into the analysis in an explicit way while maintaining an ability to weigh the trade offs between the costs and benefits of incrementally more or less stringent policies. Partly because of the large uncertainties involved, some have recommended that economists should abandon their attempts to quantify the benefits of climate policies and rely mainly on cost-effectiveness analysis instead (e.g., Ackerman et al., 2009). We agree that cost-effectiveness analysis is useful in its own right for helping to identify the most affordable ways to meet different targets, but we also believe that IAMs can be expanded to account for uncertainty in such a way that they also can inform the choice of the target itself.

Third, for the longer term (on the order of three years and more), useful tasks would include (1) continuing to support basic research on the science of climate change and its potential impacts, and (2) continuing to improve IAMs by incorporating learning and policy flexibility.

[8] http://www.rti.org/page.cfm?objectid=DDC06637-7973-4B0F-AC46B3C69E09ADA9
[9] http://www.hks.harvard.edu/m-rcbg/ptep/IGEM.htm

The importance of the first longer-term task is obvious. The weight of the scientific evidence on climate change, as summarized in the IPCC[10] and more recently the CCSP[11] reports, is substantial, but much remains to be learned. For example, as illustrated by the uncertainty surrounding the climate sensitivity parameter, there still is a very wide range of plausible future paths of global temperatures for any assumed path of GHG emissions. Even less well understood are the regional effects associated with each possible temperature path and the ensuing impacts on local ecosystems and economies. This is not to say that our knowledge is too meager to make informed policy decisions—limited information is not a sufficient condition to prefer the status quo policy. Rather, it is to say that there may be substantial value in gathering additional information in these areas. IAMs can only be as good as the scientific information that is fed into them.

The notion of the value of additional information leads to the second longer-term task. One dimension along which IAMs could be further improved is in their representation of learning and its effects on decision-making over time (e.g., Kelly and Kolstad, 1999; Fisher, 2001; Leach, 2007; Webster et al., 2008). This would require what ecologists and natural resource managers know as an "adaptive management" approach, which is a systematic framework for decision-making in the face of uncertainty that explicitly incorporates the feedbacks between learning and doing (e.g., Holling, 1978; Walters and Hilborn, 1978; Walters, 1986).[12] There are at least three advantages of this approach. First, it would in principle give more accurate estimates of the main quantity of interest: the value of emissions reduction policies in the face of uncertainty, potential learning, irreversibilities, and policy flexibility (or rigidity, as the case may be). Second, rather than a point estimate of the optimal policy it can produce an optimal policy *function,* which can be thought of a set of "contingency plans" covering the full range of possible outcomes. In other words, it can indicate how a policy instrument—such as a target, a tax, or an emissions cap—should be adjusted over time as the carbon stock grows, economic conditions change, and more scientific information accumulates. And third, it allows us to evaluate the trade-offs between the costs of emission reductions per se and the costs of collecting additional information to help reduce the uncertainties, so it can provide a unified framework for adjusting both our policy instruments and research expenditures over time.

In conclusion, our view is that economic methods, including both cost-effectiveness and benefit-cost analysis, can have a useful role in evaluating climate change policies. Nevertheless, researchers and analysts should strive to do a better job of explaining to decision-makers what their models can and cannot do. In particular, economists should better explain the meaning of the social cost of carbon estimates that their models produce so that decision-makers can make proper use of these figures in a policy setting. As an immediate corrective, economic analyses of climate change should clearly characterize the uncertainty in their results to avoid giving

[10] http://www.ipcc.ch/

[11] http://www.climatescience.gov/

[12] The jargon varies among specialties. Economists will recognize this as a "real options" framework (e.g., Dixit and Pindyck 1996, Farrow 2004), and others will know it as a "stochastic dynamic programming" approach (e.g., Ross 1983).

decision-makers and the public a false impression of precision. And moving forward, researchers should continue to improve the existing models—and create new ones as needed—to aid in the development of “contingency plans” for the full range of possible outcomes. If we know anything with certainty, it is this: the probability that the future will unfold along any single deterministic forecast is vanishingly small.

Appendix

This appendix provides a simple numerical example to illustrate the distinctions between the tiers of uncertainty analysis described in the main text.

First, assume that “social welfare” or “utility,” U, depends on aggregate income, Y, and the change in the average global temperature due to greenhouse gas emissions, T. To isolate the effect of risk aversion from time preference, we frame the problem as a static one, so we ignore the crucial dynamic dimension of the climate change problem in this example.

The willingness to pay, WTP, to prevent a change in temperature is the reduction in income with no temperature change that would make society just as well off as with the temperature change but no reduction in income. Formally, WTP is the solution to the following equation: $U(T,Y)=U(0,Y-WTP)$.

Second, assume that the damage from climate change (as a fraction of aggregate income), D, is an S-shaped function of the temperature change, specifically $D=1-\exp\left[-aT^{b}\right]$, where the parameters a and b determine the level and steepness of the damage function. Also assume that utility increases with income at a diminishing rate, specifically $U=Y^{1-\eta}/(1-\eta)$, where η is the elasticity of marginal utility (also referred to as the “coefficient of relative risk aversion”). Therefore, $U(T,Y)=\left[Y\left(1/\left[1+aT^{b}\right]\right)\right]^{1-\eta}/(1-\eta)$. These functional forms are consistent with those used in our previous work (Newbold and Daigneault, 2009).

Third, assume that the best available economic research suggests that the parameter a is between 0 and 0.006 and b is between 1 and 3. Central values are considered more likely than extreme values, so the analyst assumes symmetric triangular distributions for both parameters. (For simplicity, we assume independence between a and b.) Also assume that the best available scientific research suggests that the temperature could change by either 1, 3, or 10 degrees Celsius, with probabilities 0.13, 0.85, and 0.02 respectively.

With these assumptions, the *deterministic estimate* of willingness to pay (as a fraction of income) is

$WTP/Y=1-\exp\left[-0.003\times2.88^{2}\right]$ **= 0.0246.**

Next, to show the range of possible estimates of WTP, we conduct a *sensitivity analysis* over each uncertain parameter in turn, holding all other parameters at their expected values:

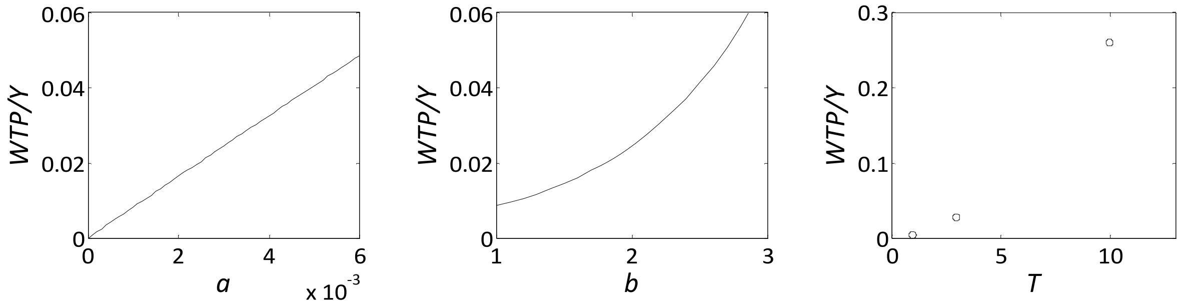

For example, the third graph above reveals that the consumption-equivalent damage from the worst-case scenario in this example is around 27 percent of current consumption.

Next, to construct a *probability distribution for WTP* we perform a Monte Carlo simulation, which involves drawing from the distributions of each uncertain parameter and re-calculating *WTP* for each draw. We then can plot the probability distribution of the results and calculate the expected value of *WTP* based on this distribution:

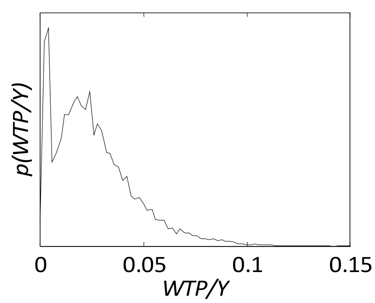

$$E\left[WTP/Y\right]=\int\int\int\left(1-\exp\left[-aT^{b}\right]\right)f\left(a\right)f\left(b\right)f\left(T\right)dadbdT \;= \mathbf{0.0314}.$$

Finally, to calculate willingness to pay using an *expected utility approach*, we find the value of *WTP* that equalizes expected utility with and without the policy, i.e., $U\left(0,Y-WTP\right)=E\left[U\left(T,Y\right)\right]$. Using the above functional forms and assuming $\eta = 2$, this gives

$$WTP/Y=1-\left[\int\int\int\left(\exp\left[-aT^{b}\right]\right)^{1-\eta}f\left(a\right)f\left(b\right)f\left(T\right)dadbdT\right]^{1/(1-\eta)}=1-E\left[U_{0}\right]^{1/(1-\eta)} \;= \mathbf{0.0446}.$$

Notice that the coefficient of relative risk aversion η appears in the calculation of *WTP* only when using the expected utility approach. The following graph shows the effect of η on *WTP*:

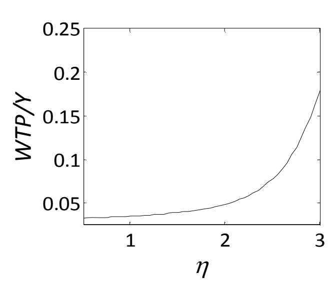

References

Ackerman, F., S. J. DeCanio, R. B. Howarth, and K. Sheeran. 2010. "The Need for a Fresh Approach to Climate Change Economics." In *Assessing the Benefits of Avoided Climate Change: Cost-Benefit Analysis and Beyond*. Gulledge, J., L. J. Richardson, L. Adkins, and S. Seidel (eds.), Proceedings of Workshop on Assessing the Benefits of Avoided Climate Change, March 16–17, 2009. Pew Center on Global Climate Change: Arlington, VA. p. 161–184. Available at: http://www.pewclimate.org/events/2009/benefitsworkshop.

Andronova, N., M. E. Schlesinger, S. Dessai, M. Hulme, and B. Li. 2007. The concept of climate sensitivity: history and development. In: *Human-induced Climate Change: An Interdisciplinary Assessment*. Schlesinger M., Kheshgi H., Smith J., de la Chesnaye F., Reilly J.M., Wilson T., Kolstad C. (eds). Cambridge: Cambridge University Press.

Arrow, K.J. and R. C. Lind. 1970. Uncertainty and the evaluation of public investment decisions. *The American Economic Review* 60(3):364-378.

Ceronsky, M., D. Anthoff, C. Hepburn, and R. S. J. Tol. 2005. Checking the price tag on catastrophe: the social cost of carbon under non-linear climate response. Working paper FNU-87. Available at: http://www.fnu.zmaw.de/fileadmin/fnu-files/publication/workingpapers/catastrophewp.pdf.

Dasgupta, P. S. and G. M. Heal. 1979. *Economic Theory and Exhaustible Resources*. New York, NY: Cambridge University Press.

Dixit, A. K. and R. S. Pindyck. 1996. *Investment Under Uncertainty*. Princeton, NJ: Princeton University Press.

Farrow, S. 2004. Using risk-assessment, benefit-cost analysis, and real options to implement a precautionary principle. *Risk Analysis* 24(3):727-735.

Fisher, A. C. 2001. Uncertainty, irreversibility, and the timing of climate policy. Pew Center on Global Climate Change working paper, Arlington, VA.

Hansen, J., Sato M, Ruedy R, Kharecha P, Lacis A, Miller RL, Nazarenko L, Lo K, Schmidt GA , Russell G, Aleinov I, Bauer S, Baum E, Cairns B, Canuto V, Chandler M, Cheng Y, Cohen A, Del Genio A, Faluvegi G, Fleming E, Friend A, Hall T, Jackman C, Jonas J, Kelley M, Kiang NY, Koch D, Labow G, Lerner J, Menon S, Novakov T, Oinas V, Perlwitz Ja, Perlwitz Ju, Rind D, Romanou A, Schmunk R, Shindell D, Stone P, Sun S, Streets D, Tausnev N, Thresher D, Unger N, Yao M, S Zhang S. 2007. Dangerous human-made interference with climate: A GISS model study. *Atmospheric Chemistry and Physics* 7:2287-2312.

Hegerl, G. C., F. W. Zwiers, P. Braconnot, N. P. Gillett, Y. Luo, J. A. Marengo Orsini, N. Nicholls, J. E. Penner, P. A. Stott. 2007. Understanding and Attributing Climate Change. In: *Climate Change 2007: The Physical Science Basis.* S. Solomon, D. Qin, M. Manning, M. Marquis, K. Averyt, M. M. B. Tignor, H. L. Miller Jr, and Z. Chen, eds. Geneva, Switzerland, IPCC.

Holling, C.S. (ed.) 1978. *Adaptive Environmental Assessment and Management*. Caldwell, NJ: Blackburn Press.

Hope C. 2006. The marginal impact of CO_2 from PAGE2002: an integrated assessment model incorporating the IPCC's five reasons for concern. *The Integrated Assessment Journal* 6:19-56.

Keller, K., B. M. Bolker, and D. F. Bradford. 2004. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management* 48:723-741.

Kelly, D.L. and C. D. Kolstad. 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23(4):491-518.

Leach, A. J. 2007. The climate change learning curve. *Journal of Economic Dynamics and Control* 33:1728-1752.

Newbold. S and A. Daigneaut. 2009. Climate response uncertainty and the benefits of greenhouse gas emissions reductions. *Environmental and Resource Economics*. http://www.springerlink.com/content/n522w44187535035.

Nordhaus, W. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.

OMB (Office of Management and Budget). 2003. Circular A-4: Regulatory analysis. Executive Office of the President, Washington, DC. http://www.whitehouse.gov/omb/circulars/.

Pindyck, R. S. 2000. Irreversibilities and the timing of environmental policy. *Resource and Energy Economics* 22:233-259.

Pindyck, R. S. 2009. Uncertainty, extreme outcomes, and climate change policy. Working paper.

Pizer, W. A. 1999. The optimal choice of climate change policy in the presence of uncertainty. *Resource and Energy Economics* 21:255-287.

Ramanathan, V. and Y. Feng. 2008. On avoiding dangerous anthropogenic interference with the climate system: formidable challenges ahead. *Proceedings of the National Academy of Sciences* 105:14245-14250.

Roe, G. H. and M. B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318:629-632.

Ross, S. M. 1983. *Introduction to Stochastic Dynamic Programming*. San Diego, CA: Academic Press Inc.

Roughgarden, T. and S. Schneider. 1999. Climate change policy: quantifying uncertainties for damages and optimal carbon taxes. *Energy Policy* 27:415-429.

Ruddiman, W. F. 2005. *Plows, Plagues, and Petroleum: How Humans Took Control of Climate*. Princeton, NJ: Princeton University Press.

Stern, N. 2006. *Stern Review: The Economics of Climate Change.* Cambridge, UK: Cambridge University Press.

Sunstein, C. R. 2007. *Worst-Case Scenarios*. Cambridge, MA: Harvard University Press.

Tol, R. S. J. 2003. Is the uncertainty about climate change too large for expected cost-benefit analysis? *Climatic Change* 56:265-289.

Tol, R. S. J. 2005. The marginal damage costs of carbon dioxide emissions: an assessment of the uncertainties. *Energy Policy* 33(16):2064-2074.

Tol, R. S. J. 2008. The social cost of carbon: trends, outliers, and catastrophes. *Economics—the Open-Access Open-Assessment E-Journal* 2(25):1-24.

Tol, R. S. J and G. W. Yohe. 2009. *The Stern Review*: a deconstruction. *Energy Policy* 37:1032-1040.

Walters, C. J. 1986. *Adaptive Management of Renewable Resources.* New York: Macmillan.

Walters, C.J. and R. Hilborn. 1978. Ecological optimization and adaptive management. *Annual Reviews of Ecology and Systematics* 9:157-188.

Webster, M., L. Jakobivits, and J. Norton. 2008. Learning about climate change and implications for near-term policy. *Climatic Change* 89:67-85.

Weitzman, M. L. 2009. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics* 91:1-19.