NCEE

NATIONAL CENTER FOR ENVIRONMENTAL ECONOMICS

# Climate Response Uncertainty and the Expected Benefits of Greenhouse Gas Emissions Reductions

**Stephen C. Newbold and Adam Daigneault**

Working Paper Series

Working Paper # 08-06
July, 2008

U.S. Environmental Protection Agency
National Center for Environmental Economics
1200 Pennsylvania Avenue, NW (MC 1809)
Washington, DC 20460
http://www.epa.gov/economics

# Climate Response Uncertainty and the Expected Benefits of Greenhouse Gas Emissions Reductions

**Stephen C. Newbold and Adam Daigneault**

NCEE Working Paper Series

Working Paper # 08-06
July, 2008

# Climate response uncertainty and the benefits of greenhouse gas emissions reductions

Stephen C. Newbold
U.S. EPA, National Center for Environmental Economics
(newbold.steve@epa.gov)

Adam Daigneault
U.S. EPA, Climate Change Division
(daigneault.adam@epa.gov)



ABSTRACT:

Some recent research suggests that uncertainty about the response of the climate system to atmospheric greenhouse gas (GHG) concentrations can have a disproportionately large influence on benefits estimates for climate change policies, potentially even dominating the effect of the discount rate. In this paper we conduct a series of numerical simulation experiments to investigate the quantitative significance of climate response uncertainty for economic assessments of climate change. First we characterize climate uncertainty by constructing two probability density functions—a Bayesian model-averaged and a Bayesian updated version—based on a combination of uncertainty ranges for climate sensitivity reported in the scientific literature. Next we estimate the willingness to pay of a representative agent for a range of emissions reduction policies using two simplified economic models. Our results illustrate the potential for large risk premiums in benefits estimates as suggested by the recent theoretical work on climate response uncertainty, and they show that the size and even the sign of the risk premium may depend crucially on how the posterior distribution describing the overall climate sensitivity uncertainty is constructed and on the specific shape of the damage function.

# 1 Introduction

Recent theoretical research on the influence of uncertainty and potential "catastrophes" associated with global climate change suggests that standard economic assessment models "may give a very misleading picture of the expected utility consequences of alternative GHG-mitigation policies" (Weitzman 2009). Weitzman illustrates this using a constant relative risk aversion (CRRA) utility function, a damage function that rises exponentially with the temperature change, and a fat-tailed probability distribution over the climate sensitivity parameter.[1] With these ingredients, society's willingness to pay to eliminate the risk of climate damages is unbounded. The model can be modified in a variety of ways to give a bounded willingness to pay, but Weitzman argues that the resulting benefit estimates still will be highly sensitive to the necessarily speculative shape of the damage function at very high (and heretofore unobserved) global average temperatures.

In this paper we investigate some of the practical implications of climate response uncertainty for policy analysis. We begin with a highly stylized climate assessment model to explore the influence of several key economic parameters on the willingness to pay to reduce climate change risks when the climate response is uncertain. We then adopt a slightly more realistic model, partly based on the climate dynamics module of the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus and Boyer 2000, Nordhaus 2008), to estimate the benefits of a recent legislative proposal that would impose a declining cap on aggregate U.S. GHG emissions, and a larger emissions reduction policy that assumes the global economy follows the optimal path of emissions as estimated by the most recent version of DICE. Both scenarios are run by first ignoring and then accounting for climate response uncertainty. This exercise provides a further check of the robustness of our stylized model and gives at least a preliminary indication of the quantitative significance of climate response uncertainty for the analysis of potential real-world policies.

[1] A pdf has a fat-tail if the probability of extreme outcomes—in this case the probability of very large temperature changes—approaches zero at a slower than exponential rate. More formally, the integral that defines the moment generating function does not converge.

This paper, motivated largely by the theoretical work of Weitzman (2009), makes five main contributions to the literature on climate change uncertainty. First, we conduct an extensive set of sensitivity analyses based on two simplified climate assessment models. These numerical experiments are intended to put Weitzman's theoretical results and conjectures through their paces—to test their robustness and to assess their quantitative significance for practical policy analyses. Second, we show that a key requirement for large "risk premiums" to emerge from economic climate assessment models is that the posterior distribution describing climate sensitivity uncertainty must be constructed using a Bayesian model averaging approach rather than a Bayesian updating approach.[2] This may provide a partial explanation for some of the divergent views—among economists as well as between economists and natural scientists—on the potential risks of catastrophic climate change. Third, we highlight the dual role of the coefficient of relative risk aversion in models of climate change with uncertainty. This parameter pulls the resulting benefit estimates in opposite directions—towards lower benefits due to the expected increases in consumption and the associated decreasing marginal value of consumption over time, but towards higher benefits due to the stronger aversion to potentially catastrophic outcomes (Heal 2008). We find that the second effect can rapidly dominate the first well within the range of plausible estimates for this parameter found in the economics literature. Fourth, we estimate benefits for two illustrative but realistic emission reduction policies accounting for climate response uncertainty. The models we use are relatively simple, but by isolating the influence of climate response uncertainty and other key economic parameters in a clean and transparent way our results can provide a useful benchmark for interpreting the benefit estimates that emerge from more complex models in

[2] In this paper we use the term "risk premium" to refer to the difference between our estimates of willingness to pay based on an expected utility framework that explicitly accounts for climate response uncertainty and our analogous estimates of willingness to pay based on an analogous deterministic model that effectively ignores the low-probability high-impact risks. 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. Also, by using the term "risk" rather than "uncertainty," we are adopting (the common understanding of) Frank Knight's distinction between these concepts (Runde 1998). According to this convention, "risk" describes a situation where a probability distribution over the potential consequences can be constructed, while "uncertainty" describes a situation where no such probability distribution can be specified. However, this convenient jargon should not disguise the fact that large elements of subjectivity may be embedded in the probability distribution describing climate response uncertainty; we are not necessarily working with purely frequentist pdfs.

future studies. Finally, we elaborate on a feature of the damage function that may be a contributing factor to the negative risk premium found by Nordhaus (2008) in an uncertainty analysis using the DICE model. Our simulations show that a negative risk premium is more likely to emerge the lower is the inflection point of the damage function.

We should clarify at the outset how the approach used in this paper accounts for uncertainty and potential climate catastrophes where standard approaches may not. In most economic climate assessment models, central or "best-guess" point estimates are used for all input parameters. For example, a common assumption about climate sensitivity is that a doubling of the atmospheric carbon concentration will cause a 3$^{o}$C increase in the average global surface temperature. However, the actual value could be smaller or much larger than this commonly cited point estimate (Hegerl et al. 2007, Andronova et al. 2007). Combining the probability distribution over the possible temperature change with a climate change damage function, which translates the actual temperature change into losses in GDP, we can define a "catastrophe" as any very high-impact very low-probability outcome associated with a temperature change above some threshold value. These outcomes are directly incorporated into the expected utility framework and the willingness-to-pay calculations of the models used in this paper, but they may be ignored completely in any standard climate assessment model that uses central parameter estimates alone.

The usual rationale for excluding such high-impact low-probability outcomes seems to be that the associated scientific uncertainty surrounding such possibilities is too large to provide a solid basis for policy evaluations. For example, there still is only limited understanding of the likelihood and timing of potential catastrophic events such as thermohaline circulation collapse or de-glaciation of the Greenland and West Antarctic ice sheets (Keller et al. 2004, Vaughan and Arthern 2007, Ramanathan and Feng 2008). However, this ad hoc rationale focuses exclusively on the "very low-probability" component of the definition. A key point implied by Weitzman (2009) is that the "very high-impact" component can potentially offset or even overwhelm the low-probability component. In general, it is the product of the probability and the value of the impact that is important, rather than one or the other alone.

Parameter uncertainty is typically analyzed in standard models—when it is not excluded altogether based on the "scientific uncertainty" rationale—through sensitivity analysis. The model is used to calculate willingness-to-pay multiple times, where each calculation uses different values for the uncertain parameters, and the range of results is reported. Sensitivity analysis is good modeling practice, but it is not a substitute for a model that explicitly incorporates uncertainty in an expected utility framework. A central goal of this paper is to assess, at least in a preliminary way, the quantitative significance of this distinction for the benefits estimates of GHG emissions reduction policies that emerge from economic climate assessment models.

A few recent studies have used Monte Carlo methods to account for uncertainty in economic climate assessment models, but so far the results have been decidedly mixed. For example, Tol (2003) used the Climate Framework for Uncertainty, Negotiation, and Distribution (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 case 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 by a factor of 20." In contrast, Pizer (1998) 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 a relatively modest 25% compared to its deterministic counterpart. And finally, Nordhaus (2008) used the latest version of DICE and found that "the best-guess policy is a good approximation to the expected-value policy." In this paper we suggest that part of the explanation for these divergent results may lie in the (possibly subtle) differences between the way that each study characterized the climate response uncertainty and the potential economic damages from global average temperature changes.

The remainder of the paper is organized as follows. Section 2 presents some motivating calculations of expected marginal willingness to pay under climate uncertainty. In Section 3 we characterize the current state of scientific uncertainty surrounding the potential response of the climate system to anthropogenic GHG emissions. Section 4 applies our estimated climate sensitivity distributions to a highly stylized economic model. Section 5 applies the same distributions to a more realistic model, based partly on DICE, to investigate two hypothetical but plausible climate change policies. The paper concludes in Section 6 with a summary of our results and a discussion of policy implications. Before proceeding, we note two key features of the climate change economics problem that we do not address in this paper: the effect of learning on the policy response (see the recent special issue of *Climatic Change*, volume 89 issues 1-2, for several articles on this topic), and the costs of emissions reductions. These aspects clearly are vitally important, but, to keep our analysis focused and tractable, we sidestep these issues and focus exclusively on the benefits of emissions reductions, where the baseline rate of economic growth is treated as exogenous.

## 2 Expected marginal willingness to pay

As a preliminary exercise, consider the motivating example by Weitzman (2009). There are two periods. Consumption $c$ in the first period is normalized to one, an uncertain temperature change $\Delta T$ will occur by the second period with probability $p(\Delta T)$, and the growth of consumption between the two periods is $e^{(g-\gamma\Delta T)t}$, where $t$ is the number of years between periods, $g$ is the baseline growth rate, and $\gamma$ is the reduction in the growth rate per unit of temperature change. Utility for the representative agent is of the "constant relative risk aversion" (CRRA) form, $u(c)=c^{1-\eta}/(1-\eta)$, where $\eta$ is the coefficient of relative risk aversion (or marginal utility of consumption). The expected marginal willingness-to-pay—"the amount of present consumption the agent would be willing to give up in the present period to obtain one extra sure unit of consumption in the future period" (Weitzman 2009)—is

$E\left[wtp\right]=e^{-\rho t}\int_{-\infty}^{\infty}e^{-\eta(g-\gamma\Delta T)t}p(\Delta T)d\Delta T$, where $\rho$ is the pure rate of time preference. Weitzman showed that if $p(\Delta T)$ has a fat tail then $E\left[wtp\right]=\infty$. This strong analytical result is all the more striking in that it is completely independent of the values of $\rho$, $t$, $\eta$, $g$, and $\gamma$. However, note that if the appropriate upper limit of integration is less than $\infty$, then $E\left[wtp\right]$ will be bounded and will depend on the values of the other parameters. For example, let $\rho$ = 0.01, $t$ = 100, $\eta$ = 2, $g$ = 0.015, $E\left[\Delta T\right]$ = 3°C, and $\gamma=-\ln(0.975)/(E\left[\Delta T\right]t)=8.44\times10^{-5}$ (i.e., if $\Delta T$ = 3 consumption will be reduced by 2.5% in period 2 relative to what it would be if $\Delta T$ = 0). With these parameter values, the "deterministic" marginal willingness to pay is $wtp=e^{-\left[\eta(g-\gamma E[\Delta T])+\rho\right]t}=0.01956$, around 2% of current consumption. To this deterministic value, we can compare what we will call a "risk adjusted" $E\left[wtp\right]$ based on $p(\Delta T)$ represented by three alternative probability density functions:[3] 1.) a gamma distribution (as in Pindyck 2009), which does not have a fat tail, 2.) a lognormal distribution, which has a fat tail but just barely, and 3.) a distribution based on Roe and Baker (2007), who derived a fat tailed pdf for the climate sensitivity parameter using the framework of feedback analysis. We calibrate all three distributions such that the 50th and 95th percentiles are 3°C and 10°C, respectively. (These are ballpark numbers used for illustration only. In the next section we characterize climate uncertainty in a more systematic way.)

First, we calculate $E\left[wtp\right]$ for all three distributions using an upper limit of integration $\Delta T_{max}$ = 10000°C (to approximate $\infty$). Using the gamma, lognormal, and Roe and Baker distributions gives $E\left[wtp\right]=0.01959$, 8x10$^{40}$, and 5x10$^{64}$, respectively. The latter two estimates confirm Weitzman's analytical result that $E\left[wtp\right]\rightarrow\infty$ under fat-tailed climate uncertainty. However, the first estimate

[3] In using the term "risk adjusted" rather than "uncertainty adjusted" to describe this willingness to pay measure, we are adopting (the common understanding of) Frank Knight's distinction between risk and uncertainty (Runde 1998). According to this convention, "risk" describes a situation where a probability distribution over the potential consequences can be constructed, while "uncertainty" describes a situation where no such probability distribution can be specified. However, this convenient nomenclature should not disguise the fact that large elements of subjectivity may be embedded in the probability distribution describing climate response uncertainty; we are not working with purely frequentist pdfs.

based on the gamma distribution, which is very close to the deterministic *wtp*, illustrates just how crucial the fat tail is to this result. Next, we recalculate $E[wtp]$ for the fat-tailed lognormal and Roe and Baker distributions using an upper limit of integration $\Delta T_{max}$ = 50°C. This is far shy of $\infty$, but this value still is much larger than the range of possible temperature changes typically discussed in the climate science literature, so it would seem like a safe upper limit of integration for the purpose of numerical calculations. However, using this limit of integration gives $E[wtp]$ = 0.01964 and 0.01968 for the lognormal and Roe and Baker distributions. In fact, the upper limit of integration $\Delta T_{max}$ must be increased to nearly 3000°C for the lognormal distribution and to nearly 450°C for the Roe and Baker distribution before the risk premium exceeds the deterministic marginal willingness to pay by more than 20%.

Of course these simple numerical calculations do not overturn Weitzman's analytical results. However, since real-world policy analyses necessarily will have to commit to specific functional forms and parameter values, these examples suggest that a more detailed quantitative exploration of the risk premium may be useful for evaluating the practical implications of climate response uncertainty. This is our task for the remainder of the paper.

## 3 Climate response uncertainty

To investigate the influence of climate response uncertainty on benefits estimates, we first must quantify the uncertainty. To do this, we use results published in the climate science literature to construct a probability distribution over the "climate sensitivity parameter," denoted here as $\lambda$. Specifically, $\lambda$ is defined as the change in the equilibrium atmospheric temperature from a doubling of the equilibrium atmospheric $CO_2$ concentration (Hegerl et al. 2007, Andronova et al. 2007). Climate sensitivity has come to be a key parameter used to summarize the influence of atmospheric GHG concentrations on the average global atmospheric temperature in many scientific studies of climate change, and a wide variety of approaches have been used to estimate it. The overall average of the estimates in the literature tends to hover around a central value near 3°C, but the estimates remain quite variable.[4]

To characterize the overall uncertainty about climate sensitivity, we must make some assumptions about the conditional independence of the information sets and models used in each study. We then can invoke Bayes' rule to construct a posterior distribution summarizing our current knowledge about climate sensitivity. We can write Bayes' rule for using observations (or simulations or other information) on temperature changes in a particular study *i*, $\Delta T_i$, to update a prior distribution for climate sensitivity, $p_0(\lambda)$, which describes the researcher's prior knowledge about $\lambda$ before the study was conducted, as follows:

$$p_i(\lambda|\Delta T_i) = \frac{l(\Delta T_i|\lambda)}{\int l(\Delta T_i|\lambda)p_0(\lambda)d\lambda} p_0(\lambda), \tag{1}$$

where $p_i(\lambda|\Delta T_i)$ is the posterior probability distribution, which describes the researcher's updated knowledge about $\lambda$ in light of the results of the study, and $l(\Delta T_i|\lambda)$ is the likelihood of observing (or simulating or otherwise estimating) $\Delta T_i$ if the true value of climate sensitivity is $\lambda$. To simplify matters, we will assume that each study used an "un-informative" prior for the climate sensitivity parameter (e.g., a uniform distribution with very wide support).[5] This means that we are treating the final probability distributions reported in each study as identical to their likelihood functions, i.e., $p_i(\lambda|\Delta T_i) = l(\Delta T_i|\lambda) \;\; \forall i = 1,2,...,N$. We then can consider two alternative views of the information on climate sensitivity in the climate change science literature. First, if we assume that each study used the

[4] In light of this persistent variability, some recent work attempts to explain why climate sensitivity has been—and may remain—so hard to pin down (Roe and Baker 2007).

[5] "Un-informative" is placed in quotes because, for example, a uniform prior with wide support over the climate sensitivity parameter is *not* equivalent to a uniform prior with wide support over the so-called "climate feedback parameter," $f \propto 1-\lambda^{-1}$ (Frame et al. 2005). Thus, an "un-formative" prior applied to the climate sensitivity parameter would imply some prior information about the value of the feedback parameter, and vice versa.

same dataset, $\Delta T_1$, but used a different model—a different set of assumptions about the "data generating process"— $m_i$, to estimate $\lambda$, then the posterior distribution would be $p_N(\lambda | \Delta T_1) = \sum_{i=1}^{N} l(\Delta T_1 | \lambda, m_i) p(m_i | \Delta T_1)$. (This is based on a standard Bayesian model-averaging approach; e.g., see Hoeting et al. 1999). Under these assumptions the posterior probability distribution for climate sensitivity should be constructed by taking a weighted average of the individual probability distributions, where the weights represent the probability that each model $m_i$ accurately describes the "true" data generating process (and where the weights sum to one, assuming one of the models is the true model). If we have no reason to believe any one model is a better representation of reality than the others *a priori*, then we would use equal weights and take a simple average of the individual pdfs. Note that this represents a relatively pessimistic view of the climate science literature since it effectively assumes that no study uses any new information and each study uses a different model to draw inferences from the single, common information set.

Alternatively, if we assume that each study used essentially the same model but an independent dataset of temperature changes, $\Delta T_i$, then the posterior distribution would be $p_N(\lambda | \Delta T_1, \ldots, \Delta T_N) = \Psi \prod_{i=1}^{N} l(\Delta T_i | \lambda)$, where $\Psi$ is a normalizing constant equal to $1 / \int \prod_{i=1}^{N} l(\Delta T_i | \lambda) d\lambda$. Under these assumptions the posterior distribution for climate sensitivity should be constructed by multiplying (rather than adding) the individual probability distribution functions together and then re-normalizing so the final distribution integrates to one. Note that this Bayesian updating approach represents a much more optimistic view of the climate science literature since it assumes that each study used a completely independent information set to estimate $\lambda$ and that the structure of the climate system is well-understood. Combining the individual pdfs under these assumptions would result in a much tighter posterior pdf for climate sensitivity.

We suspect that the truth lies somewhere between these alternative views—not all studies on climate sensitivity used uninformative priors, not all studies used the exact same or completely independent datasets, and scientific understanding of the nature of climate feedbacks is improving but still is far from complete. For this paper, we construct combined pdfs for climate sensitivity to summarize the uncertainty bounds reported in the available studies, and we use both the Bayesian model-averaging and Bayesian updating assumptions to create two alternative posterior pdfs. Specifically, we gather 28 climate sensitivity ranges from 21 studies that estimated 5th and 95th (or other) percentiles of a probability distribution for the climate sensitivity parameter. Next we solve for the parameters of a climate sensitivity distribution based on Roe and Baker (2007) with the associated percentiles for each study.[6] Then we form a Bayesian model-averaged combined pdf by taking the average of these kernel pdfs, and we form a Bayesian updated combined pdf by taking their product (and in both cases we re-normalize so the combined pdfs integrate to one). This gives two alternative posterior distributions, which we label $p_A(\lambda)$ (averaged) and $p_U(\lambda)$ (updated), that we use to summarize the uncertainty about the climate sensitivity parameter. The studies we use are shown in Table 1, and the constructed kernel density functions from each study and the resulting posterior probability distributions are shown in Figure 1.**[7]**

[6] Roe and Baker (2007) derived a distribution for climate sensitivity based on the definition $\lambda = \lambda_0/(1-f)$, where $\lambda_0$ is the no-feedback climate sensitivity and $f$ is the aggregate feedback parameter, assumed to be distributed normally with mean $\mu$ and standard deviation $\sigma$. This produces the following pdf for $\lambda$: $p(\lambda) = (2\pi\sigma^2)^{-1/2}(\lambda_0/\lambda^2)e^{-0.5[(1-\mu-\lambda_0/\lambda^2)/\sigma]^2}$. Other choices for the kernel pdf also may be reasonable and may produce different results, depending especially on the thickness of the tail of the distribution. See Annan and Hargreaves (2008) for an analysis based on an inverse Gaussian distribution and further discussion of the problem of combining multiple lines of evidence to construct probabilistic estimates of climate sensitivity.

[7] As mentioned above, we do not know where the closest approximation to reality lies relative to these alternatives, but it seems safe to say that a majority of climate scientists would endorse something much closer to the averaged than the updated pdf. The averaged pdf looks broadly similar to the summary provided in the most recent IPCC report (Hegerl et al. 2007), and the closest analog in the literature to the updated version has a 95th percentile around 4.5°C (Annan and Hargreaves 2006). For this reason, and because the *wtp* estimates based on the updated pdf behave virtually the same as the deterministic *wtp* estimates, in the remainder of the paper we place much more emphasis on the model-averaged pdf.

# 4 A stylized model of climate change impacts

We begin with a stylized model, based on Heal and Kristrom (2002) and Weitzman (personal communication), that strips away all but several key features of the problem. While the model is highly simplified, it has the practical advantages of isolating the benefits side of the analysis and of being transparent, conceptually straightforward, and easy to experiment with.

## *4.1 Model description*

As in our motivating examples in Section 2, the model is based on a standard constant relative risk aversion (CRRA) utility function: $u_t = c_t^{1-\eta}/(1-\eta)$. This is a commonly used utility function in economic climate assessment models such as DICE (Nordhaus 2008) and FUND (Tol 2006). Future utilities are discounted at a pure rate of time preference $\rho$. Consumption grows at a constant rate *g* until time *K.* During this initial period the earth's climate does not change, but at time *K* there is a shock to the system caused by an abrupt one-time temperature change such that a fraction of consumption *L* is lost. Consumption then grows again at the same constant rate *g* forever after. This amounts to consumption being reduced by the fraction *L* at every point in time after time *K* relative to what it would have been with no climate change damage.[8]

The climate change damage *L* depends directly on the change in global average surface temperatures $\Delta T$. We will consider two damage functions: the first is what we will refer to as an algebraic damage function, $L = 1 - 1/\left[1+\alpha(\Delta T)^{\beta}\right]$, and the second is an exponential damage function, $L = 1 - \exp\left[-\alpha(\Delta T)^{\beta}\right]$. Both give S-shaped curves, where the steepness of the curve and the inflection point are determined by the parameters $\alpha$ and $\beta$. See Figure 2. Finally, we assume that the equilibrium atmospheric GHG concentration will be exactly doubled at time *K*, so uncertainty about the

[8] Note that this model ignores the gradual temperature changes that are expected over the coming decades and centuries, and therefore might be thought of as focusing exclusively on the possibility of exceeding a "tipping point" in the climate system. The slightly more realistic model used in the following section exhibits gradual temperature changes.

corresponding temperature change under baseline conditions can be described directly by the posterior pdfs constructed in Section 3, $p_A(\lambda)$ or $p_U(\lambda)$.

Now we can ask the following question: What is the maximum fractional reduction in consumption, now and forever, that society would be willing to sacrifice to reduce the probability of future temperature changes—from the baseline $p_0(\Delta T)$ to (some presumably tighter distribution with a lower mean) $p_1(\Delta T)$ [9]—and the associated climate damages? The answer is given implicitly by the following equation:

$$\begin{aligned}
&\int_0^K \left(e^{gt}\right)^{1-\eta} e^{-\rho t}dt + \int_0^{\Delta T_{\max}} p_0(\Delta T)\left[\int_K^\infty \left(\left(1-L(\Delta T)\right)e^{gt}\right)^{1-\eta} e^{-\rho t}dt\right]d\Delta T \\
&\quad = \int_0^K \left(\left(1-wtp\right)e^{gt}\right)^{1-\eta} e^{-\rho t}dt + \int_0^{\Delta T_{\max}} p_1(\Delta T)\left[\int_K^\infty \left(\left(1-wtp\right)\left(1-L(\Delta T)\right)e^{gt}\right)^{1-\eta} e^{-\rho t}dt\right]d\Delta T
\end{aligned} \tag{2}$$

Shortly we will boil this equation down into a very simple expression for *wtp*, but we have written it in the expanded form above to explain each of its terms in turn. The first term on the left hand side is the discounted value of utility from the present to year *K*. We have normalized current consumption to one, and consumption grows at a constant rate *g*, so the term inside the parentheses is consumption in year *t*. Consumption raised to the power of 1-$\eta$ is utility in year *t*, and utility in each year is discounted back to the present using the pure rate of time preference $\rho$.

The second term on the left hand side is the expected discounted value of utility from time *K* to infinity under the baseline scenario. The term inside the square brackets is the discounted value conditional on a specific temperature change outcome, $\Delta T$. Recall that a fraction $L(\Delta T)$ of consumption is lost at time *K*, so at that time consumption starts at $\left(1-L(\Delta T)\right)e^{gK}$, rather than the level

[9] Our very general notation $p_0(\Delta T)$ and $p_1(\Delta T)$ is intended to accommodate any shift in the probability distribution of temperature changes. In most of the numerical experiments to follow, $p_1(\Delta T)$ will effectively be a spike at $\Delta T = 0$, i.e., the risk is eliminated entirely. In our final exercise in this section we consider cases where $p_1(\Delta T)$ is compressed to varying degrees, starting from the baseline distribution $p_0(\Delta T)$ to the (near) complete elimination of risk, and in the next section we evaluate a "small" and "large" policy that result in partial compressions of the temperature change pdf.

that would prevail if there were no climate shock, $e^{gK}$. To calculate the expected utility from time *K* to infinity, we sum (integrate) the discounted value of utility conditional on each possible value of $\Delta T$, from 0 to the largest possible temperature change $\Delta T_{\max}$, weighted by the probability of occurrence of each $\Delta T$ under the baseline scenario, $p_0(\Delta T)$.

The two terms on the right hand side are directly analogous to those on the left hand side, except for two changes: consumption is reduced by the fraction *wtp* ("willingness-to-pay") in all years, and the probability distribution over the climate damages is changed from $p_0(\Delta T)$ to $p_1(\Delta T)$. Thus, *wtp* is the fractional reduction in consumption under the policy scenario that would make society indifferent between the baseline and policy scenarios. This is the standard "compensating variation" measure of economic benefits, expressed as an annualized consumption equivalent.

Rearranging Equation (2) gives the following formula for the risk-adjusted willingness to pay for a change in the probability distribution over the potential temperature change from $p_0(\Delta T)$ to $p_1(\Delta T)$:

$$wtp = 1 - \left( \frac{E[U_0]}{E[U_1]} \right)^{1/(1-\eta)}, \tag{3}$$

where $E[U_0]$ is the left hand side of (2), the expected present value of utility conditional on $p_0(\Delta T)$ (the baseline scenario), and $E[U_1]$ is the right hand side of (2) with the (1-*wtp*) term removed, the expected present value of utility conditional on $p_1(\Delta T)$ (the policy scenario).

Equation (3) gives the basic recipe we use to calculate the risk-adjusted *wtp* in this paper. However, note that it is valid only if the numerator and denominator, both of which contain expressions of the form $\int_0^{\Delta T_{\max}} p_i(\Delta T)\left(1-L(\Delta T)\right)^{1-\eta} d\Delta T$, are finite. As Weitzman (2009) pointed out, this integral will not converge if $\Delta T_{\max} = \infty$ and $p_i(\Delta T)$ shrinks more slowly than $\left(1-L(\Delta T)\right)^{1-\eta}$ grows with $\Delta T$. To guarantee a bounded *wtp* in our numerical experiments below, we adopt Weitzman's suggestion for closing the model by including a parameter akin to a "subsistence" or "lowest tolerable" level of consumption in the utility function, *D*; specifically, $u = \max\left[c^{1-\eta}/(1-\eta),\ D^{1-\eta}/(1-\eta)\right]$. Therefore, all $\Delta T$ outcomes that would cause consumption to drop below *D* are treated as equally bad worst-case scenarios. By putting such a floor on utility, the integrals implicit in Equation (3) are guaranteed to converge and *wtp* becomes bounded. We also will investigate the effect of truncating the temperature change probability distribution with a sensitivity analysis over $\Delta T_{\max}$ (as in our motivating examples in Section 2).

We want to compare this risk-adjusted estimate of benefits to what might be produced by a standard deterministic climate assessment model that does not account for climate response uncertainty. To do this, we calculate a deterministic willingness to pay by replacing the $E\left[U_i\right]$ terms in with $U\left(E\left[\Delta T_i\right]\right)$ terms, where $E\left[\Delta T_i\right]$ is the expected temperature change under scenario *i* (the baseline or policy scenario). This is analogous to any single run of a typical climate assessment model where the best available central estimate is used for the climate sensitivity parameter.

### *4.2 Results*

The main results from this stylized climate assessment model are shown in Table 2 and Figure 3. Table 2 lists the default parameter values and the associated risk-adjusted and deterministic estimates of willingness to pay. We do not claim that our default parameter values are definitive, but we do intend them to be plausible considering the ranges typically estimated or assumed in the literature. However, the default values we use to calibrate the damage functions deserve special mention. The estimate we use for $L_3$, the economic damages from a 3°C temperature increase, is 2.5% of GDP. This is comfortably within the range of previous estimates from climate assessment models, including DICE. On the other hand, the estimate we use for $L_{10}$, the economic damages from a 10°C temperature increase, is 50% of GDP, which is purely speculative. It strikes us as reasonable in light of the potentially severe impacts discussed in the climate science literature, but unfortunately, since such a large temperature change is well outside of our historical experience and even well outside the range typically considered by climate assessment models, "reasonable speculation" is the best we can claim. Thus, a key task for our numerical exercises is to assess the sensitivity of the results to a wide range of values for $L_{10}$.

Using these default parameter values and the model-averaged pdf, the risk-adjusted *wtp* is substantially larger than the deterministic *wtp* for both damage functions—by a factor of nearly 20 for the algebraic damage function and nearly 15 for the exponential damage function. The updated pdf gives deterministic *wtp* estimates that are substantially lower than those based on the model-averaged pdf because the mean of the updated distribution is lower. Also, the difference between the deterministic and risk-adjusted *wtp* estimates shows up only at the second or third significant digit because the variance of the distribution is much smaller. So our initial impression is that, if the Bayesian model-averaged pdf is more realistic than the Bayesian updated pdf, then explicitly accounting for climate response uncertainty in an expected utility framework can increase benefits estimates substantially. If something closer to the updated pdf is more realistic, on the other hand, then the risk premium that emerges may be negligible and safe to ignore in practice.

How robust is this initial impression to variations in the model parameters? The *wtp* estimates based on the updated pdf respond to changes in the parameters exactly as we would expect based on the standard deterministic formulation. However, the story is not so simple for the Bayesian model-averaged pdf. Figure 3 shows a series of sensitivity analyses where each parameter was varied over a wide range while holding the remaining parameters at their baseline values. In some cases, the response of the model to these *ceteris paribus* changes in the parameters produce the expected directional changes in *wtp*. In particular, the top left and right panels in Figure 3 need virtually no explanation: $\partial wtp/\partial\rho$ and $\partial wtp/\partial g$ are both negative, as we would expect from the standard deterministic case.

The effect of $\eta$ is more striking. In the deterministic case, increasing $\eta$ always decreases *wtp* (when $g > 0$), as shown by the lower curves in the middle left panel of Figure 3. This also is true in the case with uncertainty for low values of $\eta$, but eventually there is a turning point where increasing $\eta$ increases the risk-adjusted *wtp*. This is counter to the standard intuition from the deterministic model, which has figured so prominently in recent discussions of the Stern Review (Stern 2006) in the literature (e.g., Dasgupta 2007, Nordhaus 2007, Stern and Taylor 2007, Sterner and Persson 2007, Weitzman 2007).

The economic intuition behind this result is as follows. Recall that in the deterministic case if $g > 0$ future generations will be wealthier than the present generation and larger $\eta$ leads to a lower net present value of climate change damages. This is because the larger is $\eta$ the more rapidly does marginal utility diminish with consumption and so the lower is the value of incremental losses in consumption for wealthier generations. Conversely, if $g < 0$ future generations will be poorer than the present generation and larger $\eta$ leads to a higher net present value of climate change damages. In the case with uncertainty, however, even if future generations are *expected* to be wealthier than the present generation, there may be some chance—if climate change impacts are much more severe than we expect—that they will in fact be poorer. If the chance of this occurring is large enough, then a larger $\eta$ (greater aversion to risk) again can lead to a higher willingness to sacrifice current consumption to prevent potential future climate change damages.

This feature of the model with uncertainty is highly relevant for the discussion in the literature on the Stern Review. Some who have commented on the Stern Review have implied that the Review's "high" estimates of climate change damages result in part from the "low" value used for $\eta$. Stern used a log utility function, which is equivalent to a CRRA utility function with $\eta = 1$. A number of economists have recommended using $\eta$ values in the range of 2 to 4 instead. Our simulations imply that, insofar as the analysis explicitly accounts for uncertainty, using larger $\eta$ values could increase rather than decrease the estimated benefits of GHG emissions reductions.

Now coming back to the results in Figure 3, note the effect of truncating the temperature change probability distribution by cutting it off at some maximum value $\Delta T_{max}$ in the middle right panel. In this case $\Delta T_{max}$ needs to be very large— greater than 80°C or so for the algebraic damage function and greater than 30°C or so for the exponential model—to produce an accurate estimate of the risk-adjusted *wtp*. The bottom left panel in Figure 3 shows the effect of the subsistence level of consumption on *wtp*. We have used the same vertical axis range for all panels in the figure to allow direct comparison of the influence of each parameter. What cannot be seen in this panel is that using the algebraic damage function *wtp* approaches 0.28 as *D* diminishes, and using the exponential damage function *wtp* approaches 1, reaching 0.99 around $D = 10^{-7}$. The influence of these two parameters taken together— $\Delta T_{max}$ and *D*—illustrate two key points emphasized by Weitzman (2009). First, willingness to pay is unbounded if the pdf has a fat tail and damages rise exponentially with the temperature change. And second, even if *wtp* is bounded, either by a subsistence level of consumption in the utility function or by a (possibly ad hoc) truncation of the climate sensitivity probability distribution, *wtp* still may be highly sensitive to assumptions about damages that are impossible to verify. Weitzman argues that on *a priori* grounds the exponential damage function should be preferred to the algebraic function, but these examples still demonstrate that the choice of the form for the damage function can have a large influence on the results, even when the alternatives diverge only at high (and therefore highly unlikely) temperature changes.

Finally, the bottom right panel in Figure 3 shows the effect on *wtp* of shifting the $\Delta T$ probability distribution. For this sensitivity analysis, rather than calculating the willingness to pay to entirely eliminate the risk of climate damages, the post-policy probability distribution is constructed using

$p_1(\Delta T) = (1-\phi)p_0(\Delta T)$ for $\Delta T > 0$ and $p_1(0) = 1 - \int_{0+}^{\Delta T_{\max}} p_1(\Delta T)\, d\Delta T$. The parameter $\phi$ has the effect of compressing the baseline pdf along the *y*-axis towards zero for all $\Delta T$ > 0 and simultaneously adding probability mass at $\Delta T$ = 0 (thereby allowing us to hold constant the upper limit of integration $\Delta T_{\max}$). If $\phi$ = 0 there is no reduction in the risk; as $\phi$ approaches 1, all of the probability mass is shifted to $\Delta T$ = 0 and the risk is completely eliminated. This experiment is intended to address the question of whether climate response uncertainty and the possibility of "catastrophic" damages associated with the very high but very unlikely temperature changes far out on the tail of the distribution have practical consequences for "marginal" GHG reductions ($\phi$ approaching 0) or only for very large (and likely unrealistic) GHG reductions ($\phi$ approaching 1)? The curves in the bottom right panel of Figure 3 provide an answer. Moving from right to left as $\phi$ shrinks from 1 to 0, the panel shows that the risk premium, in the form of the ratio of the risk-adjusted to the deterministic *wtp*, diminishes but does not approach zero. Even for marginal changes in the risk, where $\phi$ is close to zero, the risk-adjusted *wtp*s are around five times greater than their deterministic counterparts for both damage functions.

## 5 A slightly more realistic model

The stylized model analyzed above is useful for building intuition and conducting simple numerical experiments, but more realistic models will be needed for quantitative policy analysis. Specifically, one of the more unrealistic simplifications in our stylized model is the assumption of a one-time temperature shock at some known future date that moves the economy immediately from its initial equilibrium growth path to a new lower equilibrium path, but where the consumption growth rates along both paths are the same. In reality, the average global temperature will change gradually, as it is determined by globally interacting bio-geo-chemical systems with various positive and negative feedbacks that play out over long time periods, and we would expect any changes in economic output and consumption to respond gradually as well—though perhaps non-linearly, as certain threshold temperature changes or "tipping points" are crossed and progressively more severe damages occur. In this section we make a first step towards a more realistic model that incorporates uncertainty about the climate response to GHG emissions by using the climate dynamics module from DICE.

### *5.1 Model description*

DICE uses a system of first-order difference equations to represent global carbon dynamics and atmospheric temperatures, calibrated to (approximately) reproduce the aggregate outputs from the simple climate model MAGICC (Wigley 1994). Nordhaus (2008) describes DICE in detail, so we give only an abbreviated account here to indicate how the climate sensitivity parameter $\lambda$ fits into the model. Figure 4 shows a schematic representation of the "three-box" model of global carbon flows used in DICE, which takes human GHG emissions (*E*) as an input and tracks the resulting concentrations of carbon in the atmosphere ($C_A$), the upper ocean ($C_U$), and the lower ocean ($C_L$). The atmospheric temperature ($T_A$) depends on the radiative forcing ($F$), which gives the net effect of heat input from the sun and the trapping of heat partly due to atmospheric greenhouse gasses. Therefore, *F* depends in part on the atmospheric carbon concentration. This is modeled in DICE as $F(t) = F_{EX}(t) + \theta \ln\left[C_A(t)/C_A(0)\right]/\ln 2$. The forcing increases with the log of the atmospheric carbon concentration, and $\theta$ is the increase in forcing from a doubling of $C_A$. ($F_{EX}$ is "exogenous forcing," not related to human emissions of greenhouse gasses.) The atmospheric temperature evolves in response to changes in the radiative forcing and the difference between the atmospheric and oceanic temperatures: $T_A(t) = T_A(t-1) + \upsilon\left[F(t) - (\theta/\lambda)T_A(t-1) - \tau\left(T_A(t-1) - T_L(t-1)\right)\right]$. Combining these two equations and setting all differences equal to zero gives the equilibrium relationship between atmospheric temperature and carbon concentration: $T_A^{eq} = (\lambda/\theta)F_{EX}^{eq} + \lambda \ln\left[C_A^{eq}/C_A(0)\right]/\ln 2$. Thus, $\lambda$ is the change in the equilibrium atmospheric temperature from a doubling of the atmospheric carbon concentration, the same “climate sensitivity parameter” whose uncertainty we summarized in Section 3 above. (Of course this is only one of many structural parameters in DICE, so the uncertainty analysis we conduct here is not comprehensive.)

We retain all other features of the stylized model from above, namely the CRRA utility function, exogenous consumption growth (net of climate damages), subsistence consumption *D*, and the same two alternative forms for the damage function. We also use the same default parameter values. In this way, we isolate the effect of switching from the one-time-climate-shock to the more realistic climate dynamics of the DICE model. As before, we begin by calculating results for the default parameters and then conduct sensitivity analyses by varying several key parameters in turn while holding all other parameters at their default values.

Now the risk-adjusted willingness to pay is defined implicitly by

$$\int_0^{\lambda_{\max}} p(\lambda)\left[\int_0^{\infty}\left(c_t\left(\Delta T_t\left(\mathbf{E}_{0t},\lambda\right)\right)\right)^{1-\eta} e^{-\rho t}dt\right]d\lambda = \int_0^{\lambda_{\max}} p(\lambda)\left[\int_0^{\infty}\left(\left(1-wtp\right)c_t\left(\Delta T_t\left(\mathbf{E}_{1t},\lambda\right)\right)\right)^{1-\eta} e^{-\rho t}dt\right]d\lambda, \tag{4}$$

where $c_t\left(\Delta T_t\left(\mathbf{E}_{it},\lambda\right)\right)$ is consumption in year *t* under emissions scenario *i*. Consumption in each year is a function of the temperature difference, $\Delta T_t$, which in turn is a function of the path of emissions up to time *t* under that scenario, $\mathbf{E}_{it}$, and the value of $\lambda$. As before, the risk-adjusted *wtp* is calculated using Equation (3), and the deterministic *wtp* is calculated by replacing the expected net present values of all future utilities integrated over the uncertainty in $\lambda$ with the net present values conditional on the expected value of $\lambda$, under both the baseline and policy scenarios.

Note that the function $c_t(\bullet)$ captures only the climate change damages, it does not include the costs of changing the emissions path from $\mathbf{E}_0$ to $\mathbf{E}_1$. So as before we are focusing solely on the benefits side of the equation, which means that the interpretation of *wtp* here is directly analogous to that in our stylized model above: it is the willingness to pay (express as an annualized consumption equivalent) for the reduction in climate damages from a change in the emissions path from $\mathbf{E}_0$ to $\mathbf{E}_1$. These benefits would then be compared to the expected costs of making such a change in the emissions path to determine whether the policy would pass a benefit-cost test.

We analyze two hypothetical changes from the baseline emissions path, i.e., two hypothetical $\mathbf{E}_1$'s. The first represents a "large" policy, based directly on the optimal emissions path as estimated by the latest version of DICE when run in its deterministic mode (Nordhaus 2008). (Note that this policy would not necessarily be optimal in a model that accounts for uncertainty.) The second represents a "small" policy, based loosely on the Climate Stewardship and Innovation Act (S.280) introduced by Senators Lieberman and McCain. This proposal involves a cap on U.S. GHG emissions that is tightened in a step-wise fashion over time (USEPA 2007). To simulate this policy, we start with the baseline global emissions path from DICE. We then use baseline projections from the MiniCAM model (Smith and Wigley 2006) for 1990-2095 to calculate the U.S. proportion of global emissions (in total $CO_2$-equivalents). The proportion of total emissions contributed by the U.S. declines approximately exponentially in the MiniCAM simulation, so we fit an exponential decay to the U.S. proportion of total emissions to extrapolate beyond 2095 to the end of the DICE time horizon in 2405. This defines the baseline U.S. GHG emissions path. To construct the policy emissions path, we make proportional reductions in the baseline U.S. emissions according to the Lieberman-McCain caps over time, which are as follows: 2015 emissions = 2005 emissions, 2025 emissions = 1990 emissions, 2035 emissions = 18% lower than 1990 emissions, 2045 emissions = 18% lower than 1990 emissions, and 2055 [and beyond] emissions = 60% lower than 1990 emissions. We then add the reduced U.S. emissions and the rest of world's baseline emissions to get total policy emissions.[10] Thus, we are analyzing a scenario where the U.S. follows the proposed Lieberman-McCain path and the rest of the world adopts no additional measures to reduce emissions, which amounts to a relatively small change in global GHG emissions.

[10] Also, around the year 2250 fossil fuels are exhausted—or become uneconomical—in DICE and the backstop technology is adopted. When this occurs baseline emissions drop dramatically, so after this date we set the policy emissions equal to the baseline emissions.

Figure 5 shows the predicted paths of atmospheric temperature differences conditional on $\lambda = E[\lambda]$ (top panel), and the probability distributions for the temperature difference 100 years hence for the baseline scenario and both hypothetical policy scenarios (bottom panel).

Before discussing the results, it should be clear that the main purpose of using the DICE optimal path and the simulated Lieberman-McCain path, as described above, is to compare hypothetical but plausible "large" and "small" policies, where "large" and "small" refer to the magnitudes of the emissions reductions relative to the baseline emissions path at the global scale. Due to the many simplifying assumptions of our model, we have much more confidence in the relative sizes of the risk-adjusted and deterministic *wtp*s for these two hypothetical policies than in the absolute estimates of benefits for either of them alone. In particular, a rigorous and detailed analysis of the Lieberman-McCain proposal would require a more realistic integrated assessment model and is beyond the scope of this paper.

### *5.2 Results*

Since we use an exogenous growth function, we cannot solve for the optimal path of emissions reductions—the path that maximizes the net present value of consumption net of abatement costs—as in a complete optimal growth model such as DICE. Nevertheless, the deterministic version of our model using the DICE damage function and setting $\lambda$ = 3 produces an estimate of the net present value of benefits for the DICE optimal path that is very close to the estimate produced by DICE itself. Specifically, the net benefit for the optimal emissions path from DICE is 0.17% of the net present value of total future income (Nordhaus 2008 p 84), and the benefit-cost ratio for the optimal policy is 2.36 (p 82). Together these give benefits equal to 0.276% of the net present value of total future income. Our simplified model gives *wtp* equal to 0.280% of consumption for the DICE optimal emissions path. This close correspondence to the DICE results is reassuring, and suggests that the differences we see in the risk-adjusted and deterministic *wtp* estimates from this model are produced mainly by the same influences that led to the differences in our stylized model above—namely, the risk premium that emerges from the nonlinear damage and utility functions and the uncertainty about the climate response—rather than our simplifying assumptions about economic growth.

The *wtp* estimates using our default parameter values are shown in Table 3. As in the previous section, the risk premium using the exponential damage function is large. Using the Bayesian model-averaged pdf, the risk-adjusted *wtp*s are approximately 5 and 7.5 times the deterministic *wtp*s for the DICE optimal path and the simulated Lieberman-McCain path, respectively. In stark contrast, however, the risk premiums using the algebraic damage function are negligible for all practical purposes. Furthermore, the risk "premium" is actually *negative* for the simulated Lieberman-McCain path using the algebraic damage function. (We will explore this result in more detail below, when we look at the sensitivity of the outputs to the shape of the damage function.)

As before, the updated pdf produces results that behave virtually the same as the deterministic model, so next we check the sensitivity of the risk-adjusted *wtp* using the model-averaged pdf to $\eta$, *D*, and $L_{10}$. These are the key parameters that determine the severity of the worst-case scenarios (i.e., how "catastrophic" the climate change outcomes can possibly be) and how risk-averse the representative agent is to these outcomes. The results of these sensitivity analyses are shown in Figure 6. The graphs on the left apply to the DICE optimal emissions path, and those on the right apply to the simulated Lieberman-McCain emissions path.[11]

The risk-adjusted *wtp* estimates decline monotonically with $\eta$ using the algebraic damage function, but using the exponential damage function they initially decrease then increase rapidly with $\eta$

[11] We can make a very rough comparison between these benefits estimates and estimates of costs for the Lieberman-McCain proposal as follows. An analysis by the U.S. EPA (http://www.epa.gov/climatechange/downloads/ s280fullbrief.pdf) estimated that the Lieberman-McCain proposal would reduce the average annual growth rate of consumption in the U.S. by approximately 0.04 percentage points. We can convert this to an annualized willingness to pay figure analogous to our benefits estimates using $\int_0^\infty e^{((1-\eta)(g-\Delta)-\rho)t} dt = \int_0^\infty (1-wtp)^{1-\eta} e^{((1-\eta)g-\rho)t} dt$, which gives $wtp = 1 - \left[ \left(g(1-\eta)-\rho\right) / \ \left((g-\Delta)(1-\eta)-\rho\right) \right]^{1/(1-\eta)}$. Using $\Delta = 0.0004$ and our default parameter values from Table 2 in this formula gives a cost estimate of 0.016, or 1.6% of consumption, which is larger than any of the benefits estimates for the simulated Lieberman-McCain emissions path reported in Table 3. However, as seen in Figure 6, risk-adjusted benefits estimates substantially larger than this could emerge using other reasonable values of $\eta$, *D*, or $L_{10}$.

as in the stylized model above. This gives another illustration of the sensitivity of the risk-adjusted *wtp* to the specific functional form of the damage function, even when the alternative forms have the same general shape and are "calibrated" to give the same damages at low to intermediate temperature changes. The next two graphs show the influence of *D* on the *wtp* estimates. Using the exponential damage function, as *D* approaches zero the risk-adjusted *wtp*s approach approximately 0.15 and 0.12 for the DICE optimal and simulated Lieberman-McCain emissions paths, respectively. Thus, the risk adjusted *wtp* is seen to increase substantially as *D* approaches zero, but *D* = 0 is not a sufficient condition for *wtp* to approach 1 (in particular, $\lambda_{\max}$ may need to reach very high values as well). The final two graphs show the effect on the *wtp* estimates of changing $L_{10}$, the loss in per capita consumption if $\Delta T$ = 10°C. Adjusting $L_{10}$ shifts the entire damage function, but its main effect is to change the rate at which the damage function approaches 1 as $\Delta T$ increases to very high levels. Recall that our default value for $L_{10}$ was 0.5. The risk-adjusted *wtp* rises rapidly after this point for the optimal DICE emissions path using the exponential damage function. This suggests that the precise rate at which the severity of the damages increases with the equilibrium temperature difference in the "globally catastrophic" range of outcomes can have a disproportionate influence on the benefits estimates of climate change policies. But of course these are consequences about which we can only speculate, precisely because they are so far outside our past experience and even most climate change simulation models.

Now we come back to the negative risk premium seen in Table 3 for the simulated Lieberman-McCain path using the model-averaged pdf and the algebraic damage function. Note that the low end of the $L_{10}$ range in the bottom left panel of Figure 6 is where we find the damage function used in the DICE model by Nordhaus (2008), which is of the algebraic form with $L_3 = 0.024913$ and $L_{10} = 0.22111$.

Using these damage function parameters, the deterministic *wtp* is around 15 percent *larger* than the risk-adjusted *wtp* for both damage functional forms and for both hypothetical emissions paths.

This is consistent with Nordhaus' result based on sensitivity analysis using a Monte Carlo approach (Nordhaus 2008 Ch 7). Nordhaus ran the DICE model 100 times with different randomly-drawn values for a number of key input parameters, including the climate sensitivity parameter. In each case the "social cost of carbon" (SCC) was calculated and then averaged and compared to the SCC from the baseline version where all parameters were fixed at their expected values. The surprising result was that the average SCC from the Monte Carlo analysis was $26.85 per ton of carbon, approximately 5% lower than the deterministic estimate of $28.10. Nordhaus explained this result by calculating the correlation between the predicted temperature difference and the level of consumption 100 years in the future. These turn out to be positively correlated, which means that states of the world with high climate damages tend also to be states with high consumption. This naturally leads to a negative risk premium.

However, because multiple parameters were varied in Nordhaus' Monte Carlo analysis, it is difficult to trace the specific source(s) of the negative risk premium. For example, one of the uncertain parameters was the growth rate of total factor productivity. It is easy to see how this parameter contributes to the negative risk premium. When total factor productivity is varied, a positive correlation between consumption and temperature naturally will emerge because high total factor productivity will lead to high production (which allows for higher consumption) and therefore high emissions (which causes higher temperatures). However, this particular feature is absent from our model. We assume that the consumption growth rate net of climate damages is exogenous and fixed, which means that high temperature outcomes are not positively correlated with high consumption outcomes. We must look elsewhere for the explanation of our negative risk premium.

In our case we are able to trace the source of the negative risk premium directly to the less severe damage function used in DICE. To see this, consider a simplified scenario with only two possible climate change outcomes: with probability *p* the temperature difference will be $\Delta T_1$ and with probability 1-*p* it will be $\Delta T_2$ $(>\Delta T_1)$. The expected damage is $E\left[D(\Delta T)\right]= pD(\Delta T_1)+ (1-p)D(\Delta T_2)$. Now consider a small mean-preserving spread of this two point distribution, i.e., increasing $\Delta T_2$ by a small amount $d\Delta T_2$ and simultaneously decreasing $\Delta T_1$ by a corresponding amount $d\Delta T_1 = -p/(1-p)d\Delta T_2$. Taking the total derivative of the expected damage gives $dE[D]/d\Delta T_2 =$ $(1-p)[\partial D_2/\partial\Delta T_2 - \partial D_1/\partial\Delta T_1]$, which shows that the expected damages will increase (decrease) if the slope of the damage function at $\Delta T_2$ is greater (less) than the slope at $\Delta T_1$. Therefore, simple intuition based on Jensen's inequality alone does not suffice here since the damage function is not everywhere convex in $\Delta T$—it is convex for low $\Delta T$ but concave for high $\Delta T$. Over a wide range, increasing the variance of the distribution—i.e., shifting some of the mass of the pdf towards higher $\Delta T$ values where marginal damages are decreasing—makes the risk premium more negative. And the negative risk premium is more likely to emerge the lower is the inflection point of the damage function, all else equal. This may help to explain both the negative risk premium seen in Table 3 for the simulated Lieberman-McCain emissions path and the negative risk premiums estimated for both damage functional forms and both hypothetical emissions paths using the DICE damage function parameters. It also may be part of explanation for the negative risk premium found by Nordhaus (2008).

# 6 Discussion and Conclusions

The main common thread running through most of our results is that, due to the non-linearities in the utility function and the damage function, the greater the level of uncertainty the more the representative agent is willing to pay to reduce the risks of climate change. However, the magnitude of this risk premium appears to be very sensitive to several key parameters. In particular, seemingly subtle changes in the upper bound on the possible temperature change and the shape of the damage function in this extreme territory can produce a very wide range of benefits estimates, from a negative risk premium all the way to a willingness to pay approaching the full value of global economic output. These results are especially striking considering that these crucial assumptions about the damage function and temperature differences are confined to the extreme tail of the (Bayesian model-averaged) climate sensitivity probability distribution. However, as we saw in Section 5, even the qualitative conclusion that "increasing climate response uncertainty generates larger *wtp*" is not universal. If the inflection point of the damage function is low enough relative to the pdf of climate sensitivity, then increasing the uncertainty can decrease the benefits.

Our results also suggest that developing a more realistic posterior probability distribution for climate sensitivity should be a high priority for further research. The approaches we used in Section 2 to construct our posterior pdfs for the climate sensitivity parameter are admittedly simplistic and intended mainly to illustrate the relevance of this issue for benefits estimation. As noted earlier, while the updated pdf is clearly overly optimistic, some nontrivial tightening of the posterior distribution may be possible by combining multiple (at least partially) independent lines of evidence that have thus far been analyzed separately. Developing a more accurate characterization of climate response uncertainty in the short run will require a careful sorting-out of the common features and differences among the various models and datasets that have been used to date to estimate climate sensitivity. (To our knowledge, the most detailed attempts to do this so far are those of Annan and Hargreaves 2006, 2008.) In the longer run, more climate model testing and additional lines of empirical evidence relevant for estimating climate sensitivity will be required to further narrow the posterior distribution.

However, there may be a limit to the rate at which we can narrow this uncertainty over time (Weitzman 2009, Roe and Baker 2007). Therefore, another key direction for further research is to investigate the potential for learning about the true value of climate sensitivity, how rapidly such learning can occur, and how the policy response may be influenced by learning (e.g., Hammitt et al. 1992, Kolstad 1996, Webster 2008). On this score at least, the fundamental dynamics of the system may work in our favor. If climate sensitivity is high, learning should occur rapidly since in this case it would be possible to distinguish the long-run temperature rise from the background natural climate variability relatively quickly. If climate sensitivity is low, learning will occur more slowly since it would take longer to identify the signal through the noise. In this case, however, the social costs of this slow pace of learning also would be lower since the eventual climate damages would be less severe.

A third task for future work is to conduct additional uncertainty analyses using more sophisticated climate assessment models. In the meantime, by isolating the influence of climate response uncertainty and other key economic parameters in a clean and transparent way our results can provide a useful benchmark for interpreting the benefit estimates that emerge from more complex models. However, an important consideration when conducting such analyses will be to properly represent the current information and expectations of the agent(s) in the model with respect to future climate changes. We skirted this issue in our numerical experiments by treating the consumption growth rate as exogenous. When using a fully-specified dynamic optimization model where savings (investment) and abatement are control variables, it will be important to avoid making the (implicit) assumption that all agents in the model know the true climate sensitivity. For example, this would be the effect of simply nesting a deterministic dynamic programming-based climate assessment model in a Monte Carlo analysis. Each iteration of the model would run a deterministic scenario, each with different values for the (assumed known) model parameters. A proper accounting of uncertainty in the context of a complete climate assessment model would require endogenizing the learning process. Markets would clear every period based on the current expectations about future climate changes, but those expectations may evolve over time as new information about climate sensitivity accumulates.

In the meantime, where does this leave the economist who wishes to use benefit-cost analysis to evaluate alternative climate change policies? There is an unavoidable tension between the usefulness of benefit-cost analysis as a means to dispassionately weigh the advantages and disadvantages of proposed environmental policies (Arrow et al. 1996) and the danger of overconfidence and undue reliance on such exercises (e.g., DeCanio 2003 Ch 5). This inherent tension strikes us as even more salient in light of our numerical experiments. Our initial intuition was that the risk-adjusted *wtp* estimates should be fairly robust and "well-behaved" in the relevant space of the key parameters. We expected that the risk premium generally would be positive and possibly non-negligible, but would reach extremely large values only as *D* diminished and $\Delta T_{max}$ increased to values possibly far outside the plausible ranges for these parameters. Some of our results were consistent with these expectations—e.g., the risk-adjusted marginal willingness to pay calculated in Section 2 is very large only if $\Delta T_{max}$ is orders of magnitude larger than the highest values typically discussed in the climate science literature. But overall we found the risk-adjusted *wtp* to be more fragile than we initially expected well within the range of typical values discussed in the economics literature for several key parameters, including the coefficient of relative risk aversion. We do not believe this invalidates the use of climate assessment models for benefit-cost analysis, but we are persuaded that analyses intended to inform real-world climate policy decisions should account for climate response uncertainty in a rigorous way. Simplified analyses that ignore the effects of uncertainty may be useful for illustrative or screening purposes, but strictly deterministic models should be viewed as provisional only and given proportionately less weight in policy deliberations. Thus, our view is that these results increase *both* the importance of dispassionately weighing the advantages and disadvantages of alternative climate policies *and* the potential dangers of misplaced confidence in the economic models currently available for conducting this task.

To sum up, in this paper we have conducted a series of numerical simulation experiments to assess the influence of climate response uncertainty on economic benefits estimates of GHG emissions reductions. The risk-adjusted estimates of *wtp* in Sections 4 and 5 using the model-averaged pdf constructed in Section 3 are often many times larger than their deterministic counterparts, and they are highly sensitive to seemingly subtle differences in assumptions about the damage function at very high (and necessarily speculative) damage levels. The results of our sensitivity analyses easily span the range of results found in previous economic studies of climate change uncertainty, from the very large risk premiums seen in some results by Tol (2003) to the small negative risk premium seen in the results by Nordhaus (2008). In this respect, our results based on the model-averaged pdf reinforce and further illustrate the theoretical model of Weitzman (2009). However, our results also indicate that the very large risk premiums seen in some of our results could be reduced if at least some of the existing estimates of climate sensitivity are in fact based on substantially independent lines of evidence, in which case a tighter posterior distribution for climate sensitivity may be appropriate.

## References

Andronova NG, Schlesinger ME (2001) Objective estimation of the probability density function for climate sensitivity. Journal of Geophysical Research 106(D19):22605–22611

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

Annan JD, Hargreaves JC (2006) Multiple observationally based constraints to estimate climate sensitivity. Geophysical Research Letters 33:L06704 http://www.jamstec.go.jp/frcgc/research/d5/jdannan/GRL_sensitivity.pdf

Annan JD, Hargreaves JC (2008) On the generation and interpretation of probabilistic estimates of climate sensitivity.

Arrow KJ, Cropper ML, Eads GC, Hahn RW, Lave LB, Noll RG, Portney PR, Russel M, Schmalensee R, Smith VK, Stavins RN (1996) Is there a role for benefit-cost analysis in environmental, health, and safety regulation? Science 272: 211-222

Ceronsky M, Anthoff D, Hepburn C, Tol RSJ (2005) Checking the price tag on catastrophe: the social cost of carbon under non-linear climate response. Working paper FNU-87 http://www.fnu.zmaw.de/fileadmin/fnu-files/publication/working-papers/catastrophewp.pdf

Dasgupta P (2007) Comments on the Stern Review's Economics of Climate Change http://www.econ.cam.ac.uk/faculty/dasgupta/STERN.pdf

DeCanio SJ (2003) Economic Models of Climate Change: A Critique. Palgrave MacMillan: New York

Forest CE, Stone PH, Sokolov AP, Allen MR, Webster MD (2002) Quantifying uncertainties in climate system properties with the use of recent observations. Science 295:113-117

Forest C, Webster M, Reilly J (2004) Narrowing uncertainty in global climate change. *The Industrial Physicist* AugSep 7/23/04:22-25. http://www.aip.org/tip/INPHFA/vol-10/iss-4/p20.pdf

Forest DJ, Stone PH, Sokolov AP (2006) Estimated PDFs of climate system properties including natural and anthropogenic forcings. Geophysical Research Letters 33:L01705

Forster PMDF, Gregory JM (2006) The climate sensitivity and its components diagnosed from Earth radiation budget data. Journal of Climate 19:39–52

Frame DJ, Booth BBB, Kettleborough JA, Stainforth DA, Gregory JM, Collins M, Allen MR (2005) Constraining climate forecasts: The role of prior assumptions. Geophysical Research Letters 32:L09702

Gregory JM, Stouffer RJ, Raper SCB, Stott PA, Rayner NA (2002) An observationally based estimate of the climate sensitivity. *Journal of Climate* 15(22):3117–3121

Hammitt JK, Lempert RJ, Schlesinger ME (1992) A sequential-decision strategy for abating climate change. Nature 357:315-318

Heal G, Kristrom B (2002) Uncertainty and climate change. Environmental and Resource Economics 22:3-39

Heal G (2008) Climate economics: a meta-review and some suggestions. NBER working paper. http://www.nber.org/papers/w13927

Harvey LDD, Kaufmann RK (2002) Simultaneously constraining climate sensitivity and aerosol radiative forcing. Journal of Climate 15(20):2837-2861

Hegerl GC, Crowley TJ, Hyde WT, Frame DJ (2006) Climate sensitivity constrained by temperature reconstructions over the past seven centuries. Nature 440:1029-1032

Hegerl GC, Zwiers FW, Braconnot P, Gillett NP, Luo Y, Marengo Orsini JA, Nicholls N, Penner JE Stott PA (2007) Understanding and Attributing Climate Change. In: Solomon S, Qin D, Manning M, Chen Z, Marquis M, Averyt KB, Tignor M, HM, eds. Climate Change 2007: The Physical Science Basis.

Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press

Hoeting JA, Madigan D, Rafterty AE, Volinsky CT (1999) Bayesian model averaging: a tutorial. Statistical Science 14(4):382-417

Hoffert MI, Covey C (1992) Deriving global climate sensitivity from paleoclimate reconstructions. Nature 360: 573–576

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(1):19-56. http://journals.sfu.ca/int_assess/index.php/iaj/article/view/227/190

Keller K, Bolker BM, Bradford DF (2004) Uncertain climate thresholds and optimal economic growth. Journal of Environmental Economics and Management 48:723-741.

Kerr RA (2004) Three degrees of consensus. Science 305:932-934

Knutti R, Stocker TF, Joos F, Plattner G-K (2002) Constraints on radiative forcing and future climate change from observations and climate model ensembles. Nature 416: 719–723

Kolstad CD (1996) Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. Journal of Environmental Economics and Management 31:1-18

Morgan MG, Keith D (1995) Subjective judgments by climate experts. Environmental Science and Technology 29:468A-476A

Murphy JM, Sexton DMH, Barnett DN, Jones GS, Webb MJ, Collins M, Stainforth DA (2004) Quantification of modeling uncertainties in a large ensemble of climate change simulations. Nature 430:768-772

National Academy of Sciences (NAS) (1979) Carbon Dioxide and Climate: A Scientific Assessment. Washington, DC: US National Academy of Sciences

Nordhaus W (2007) Critical assumptions in the Stern Review on Climate Change. Science 317:201-202

Nordhaus W (2008) A Question of Balance: Weighing the Options on Global Warming Policies. http://nordhaus.econ.yale.edu/Balance_2nd_proofs.pdf

Nordhaus WD, Boyer J (2000) Warming the World: Economic Models of Global Warming. Cambridge, MA: MIT Press

Piani C, Frame DJ, Stainforth A, Allen MR (2005) Constraints on climate change from a multi-thousand member ensemble of simulations. Geophysical Research Letters 32(23):L23825

Pindyck RS (2009) Uncertainty, extreme outcomes, and climate change policy. Working paper

Pizer WA (1998) Optimal choice of policy instrument and stringency under uncertainty: the case of climate change. RFF discussion paper 98-XX. http://www.rff.org/~pizer/re.pdf

Pizer WA (1999) The optimal choice of climate change policy in the presence of uncertainty. Resource and Energy Economics 21:255-287

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

Roe GH, Baker MB (2007) Why is climate sensitivity so unpredictable? Science 318:629-632

Runde J (1998) Clarifying Frank Knight's discussion of the meaning of risk and uncertainty. Cambridge Journal of Economics 22:539-546

Schneider von Deimling T, Held H, Ganopolski A, Rahmstorf S (2006) Climate sensitivity estimated from ensemble simulations of glacial climate. Climate Dynamics 27:149–163

Smith SJ, Wigley TML (2006) Multi-gas forcing stabilization with MiniCAM (Mini Climate Assessment Model). Energy Journal Special Issue #3: 373-392

Stainforth DA, Aina T, Christensen C, Collins M, Faull N, Frame DJ, Kettleborough JA, Knight S, Martin A, Murphy JM, Piani C, Sexton D, Smith LA, Spicer RA, Thorpe AJ, Allen MR (2005) Uncertainty in predictions of the climate response to rising levels of greenhouse gases. Nature 433:403-406.

Stern N (2006) The Stern Review on the Economics of Climate Change. Cambridge, UK: Cambridge University Press

Stern N, Taylor C (2007) Climate change: risk, ethics and the Stern Review. Science 317:203-204

Sterner P, Persson UM (2007) An even Sterner review: introducing relative prices into the discounting debate. Resources for the Future working paper http://www.rff.org/Documents/RFF-DP-07-37.pdf

Tol RSJ, de Vos AF (1998) A Bayesian statistical analysis of the enhanced greenhouse effect. Climatic Change 38:87–112

Tol RSJ (2003) Is the uncertainty about climate change too large for expected cost-benefit analysis? Climatic Change 56:265-289

Tol RSJ (2006) The climate framework for uncertainty, negotiation and distribution (FUND), technical description, version 2.8

USEPA (2007) EPA analysis of the Climate Stewardship and Innovation Act of 2007: *S.*280 in the 110th Congress. http://www.epa.gov/climatechange/downloads/s280fullbrief.pdf

Vaughan DG, Arthern R (2007) Why is it hard to predict the future of ice sheets? Science 315:1503-1504

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

Weitzman ML (2007) The Stern Review of the Economics of Climate Change. http://www.economics.harvard.edu/faculty/weitzman/files/JELSternReport.pdf

Weitzman ML (2009) On modeling and interpreting the economics of catastrophic climate change. Review of Economics and Statistics 91(1):1-19

Wigley TML (1994) MAGICC (Model for the Assessment of Greenhouse-gas Induced Climate Change): User's Guide and Scientific Reference Manual. National Center for Atmospheric Research: Boulder, CO. http://sedac.ciesin.org/mva/MAGICCUM/MAGICCUM.html

Wigley TML, Ammann CM, Santer BD, Raper SCB (2005) Effect of climate sensitivity on the response to volcanic forcing. Journal of Geophysical Research 110:D09107

## Tables and Figures

**TABLE 1.** Studies used to construct the combined pdfs for the climate sensitivity parameter.

| **Authors (year)** | **$5^{th}$** | **$95^{th}$** | **Notes** |
|---|---|---|---|
| NAS (1979) | -0.6 | 6.5 | Established the original range of 1.5-4.5$^{O}$C for the $25^{th}$ and $75^{th}$ percentile later cited in all IPCC assessment reports. (Our analysis uses the $25^{th}$ and $75^{th}$ percentiles in constructing the kernel pdf for this study.) |
| Hoffert and Covey (1992) | 1.4 | 3.2 | Derived global climate sensitivity from paleoclimate reconstructions using last glacial maximum (LGM) approach. |
| Morgan and Keith (1995) | -0.8 | 5.8 | Expert elicitation of 16 climate researchers. (To construct a log normal distribution for this study, we substitute the lowest positive $5^{th}$ percentile from the remaining studies in this table for authors' own negative estimate.) |
| Tol and de Vos (1998) | 1.6 | 8.9 | Time series analysis of historical estimates of atmospheric temperatures and $CO_2$ concentrations. Bayesian approach. |
| Andronova and Schlesinger (2001) | 1.0 | 9.3 | Used observed global mean and hemispheric difference in surface air temperature 1856-1997. Monte Carlo analysis using an energy balance model (EBM). |
| Forest et al. (2002) | 1.4 | 7.7 | Used non-uniform expert prior distribution of equilibrium climate sensitivities (ECS). |
| Knutti et al. (2002) | 2.2 | 9.1 | Used global mean ocean heat uptake and global mean surface air temperature increase. Directly include the indirect forcing effect. |
| Gregory et al. (2002) | 1.1 | $\infty$ | Used surface air temperature space and time patterns, and an atmospheric-ocean GCM. (To construct a probability |

| | | | |
|---|---|---|---|
| | | | distribution for this study, we substitute the highest 95th percentile from the remaining studies in this table for the authors' own estimate of ∞ .) |
| Harvey and Kaufmann (2002) | 1.0 | 3.0 | Separated forcing from aerosols and GHGs to remove some uncertainties using inverse method. |
| Forest et al. (2004) | 1.4 | 7.8 | Monte Carlo analysis of economic and earth system parameters using an integrated global systems model. |
| Kerr (2004) | 2.4 | 5.4 | Applied "perturbed physics" approach to global climate model. |
| Murphy et al. (2004) | 1.8 | 5.2 | Used a "perturbed physics ensemble method" (PPEM) with all model versions assumed equally likely. |
| | 2.4 | 5.2 | Used PPEM with reliability-based weighting of model versions according a climate Prediction Index. |
| Frame et al. (2005) | 1.2 | 11.8 | Applied global change in surface temperature to EBM. Used a uniform prior distribution that extends beyond 10°C sensitivity. |
| Stainforth et al. (2005) | 1.5 | 11.5 | Used PPEM method for six model parameters. |
| Wigley et al. (2005) | 1.3 | 6.3 | Estimated effect of climate sensitivity on the response to volcanic forcings using individual volcanoes. Agung volcano. |
| | 0.3 | 7.7 | El Chichon volcano. |
| | 1.8 | 5.2 | Mt. Pinatubo volcano. |
| Piani et al. (2005) | 2.2 | 6.8 | Used PPEM via distributed computing project. |
| Annan and Hargreaves (2006) | 1.7 | 4.5 | Used Bayesian method to sharpen the posterior distribution of ECS. |
| Forest et al. (2006) | 2.1 | 8.9 | Used approach similar to Forest et al. (2002), including both |

| | | | |
|---|---|---|---|
| | | | natural and anthropogenic forcings. |
| | 1.4 | 7.7 | Without natural forcing. |
| | 1.4 | 4.1 | Using expert priors. |
| | 1.4 | 7.7 | Using uniform priors. |
| Forster and Gregory (2006) | 1.2 | 14.2 | Used Earth Radiation Budget Experiment (ERBE) combined with surface temperature observations based on a regression approach. |
| Hegerl et al. (2006) | 1.2 | 8.6 | Used multiple palaeoclimatic reconstructions of Northern Hemisphere mean temperatures over the last 700 years. |
| | 1.5 | 6.2 | With non-uniform prior distributions. |
| Schneider von Deimling et al. (2006) | 1.2 | 4.3 | Used PPEM with varied atmospheric and ocean parameters in simulation of the LGM. |

**TABLE 2.** Default parameter values and the resulting deterministic and risk-adjusted estimates of willingness to pay to prevent climate change damages as a fraction of GDP using the climate sensitivity probability distributions constructed in Section 2 and the stylized IAM described in Section 3.

| | | |
|---|---|---|
| *Default parameter values:* | | |
| $\rho$, pure rate of time preference | 0.01 | |
| *g*, consumption growth rate | 0.015 | |
| $\eta$, elasticity of marginal utility | 2 | |
| *D*, "subsistence" consumption | 0.01 | |
| *K*, year when climate shock hits | 100 | |
| $\Delta T_{\max}$, maximum possible temperature change | 100 | |
| $L_3$, fraction of GDP lost if $\Delta T$ = 3 deg C | 0.03 | |
| $L_{10}$, fraction of GDP lost if $\Delta T$ = 10 deg C | 0.5 | |
| *Baseline results* | | |
| | Algebraic | Exponential |
| Bayesian model-averaged pdf: | | |
| Deterministic *wtp* | 0.00315 | 0.00304 |
| Risk-adjusted *wtp* | 0.06099 | 0.04469 |
| Bayesian updated pdf: | | |
| Deterministic *wtp* | 0.00086 | 0.00093 |
| Risk-adjusted *wtp* | 0.00087 | 0.00094 |

**TABLE 3.** Deterministic and risk-adjusted estimates of willingness to pay to prevent climate change damages as a fraction of GDP using the default parameter values from Table 2 in the slightly more realistic climate assessment model described in Section 4.

| | Algebraic | Exponential |
|---|---|---|
| Bayesian model-averaged pdf: | | |
| *DICE optimal path:* | | |
| Deterministic *wtp* | 0.00610 | 0.00542 |
| Risk-adjusted *wtp* | 0.00610 | 0.02701 |
| *Simulated Lieberman-McCain path:* | | |
| Deterministic *wtp* | 0.00112 | 0.00100 |
| Risk-adjusted *wtp* | 0.00109 | 0.00770 |
| Bayesian updated pdf: | | |
| *DICE optimal path:* | | |
| Deterministic *wtp* | 0.00256 | 0.00241 |
| Risk-adjusted *wtp* | 0.00256 | 0.00242 |
| *Simulated Lieberman-McCain path:* | | |
| Deterministic *wtp* | 0.00048 | 0.00045 |
| Risk-adjusted *wtp* | 0.00048 | 0.00045 |

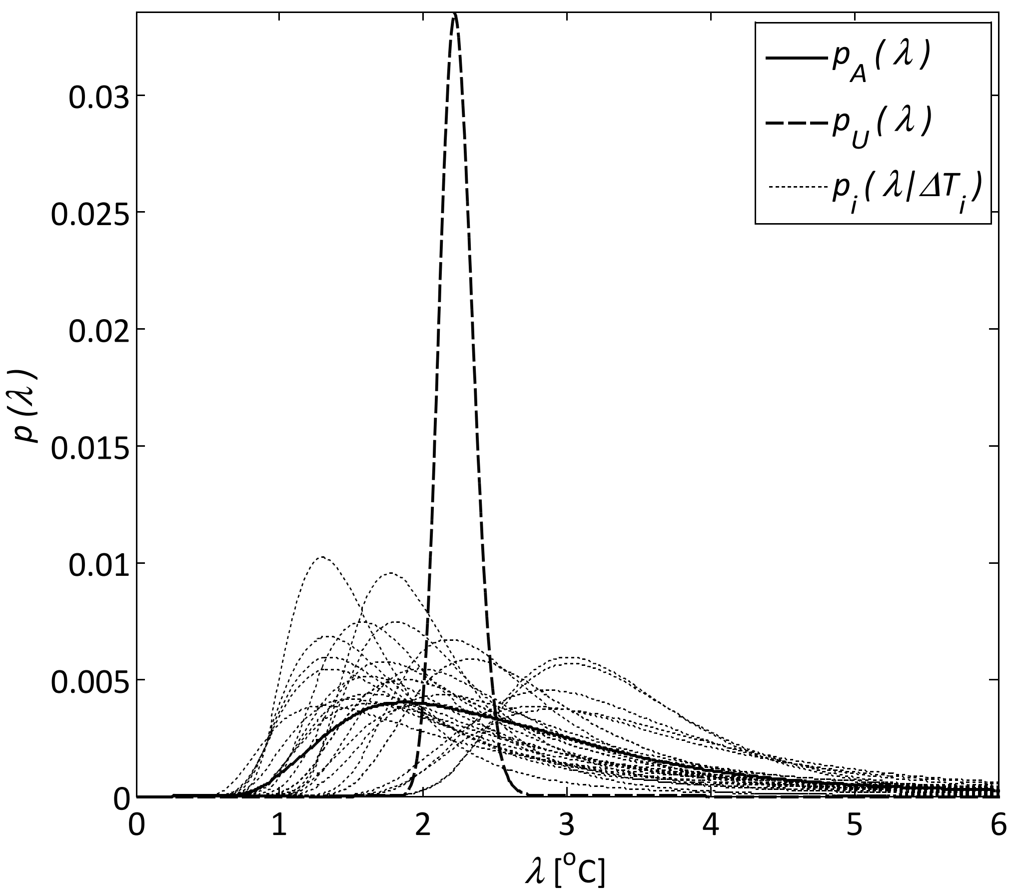


**FIGURE 1.** Roe and Baker (2007) probability distributions constructed from the 5$^{th}$ and 95$^{th}$ percentiles for the climate sensitivity parameter from each study shown in Table 1 (light dotted lines), the Bayesian model-averaged pdf based on the average of the distributions using equal weights (heavy solid line), and the Bayesian updated pdf based on the product of the distributions (heavy dotted line). The mean of the model-averaged pdf is 3.42$^O$C and the 5$^{th}$, 50$^{th}$, and 95$^{th}$ percentiles are 1.23$^O$C, 2.49$^O$C, and 7.45$^O$C. The mean of the updated pdf is 2.24$^O$C and the 5$^{th}$, 50$^{th}$, and 95$^{th}$ percentiles are 2.04$^O$C, 2.22$^O$C, and 2.43$^O$C.

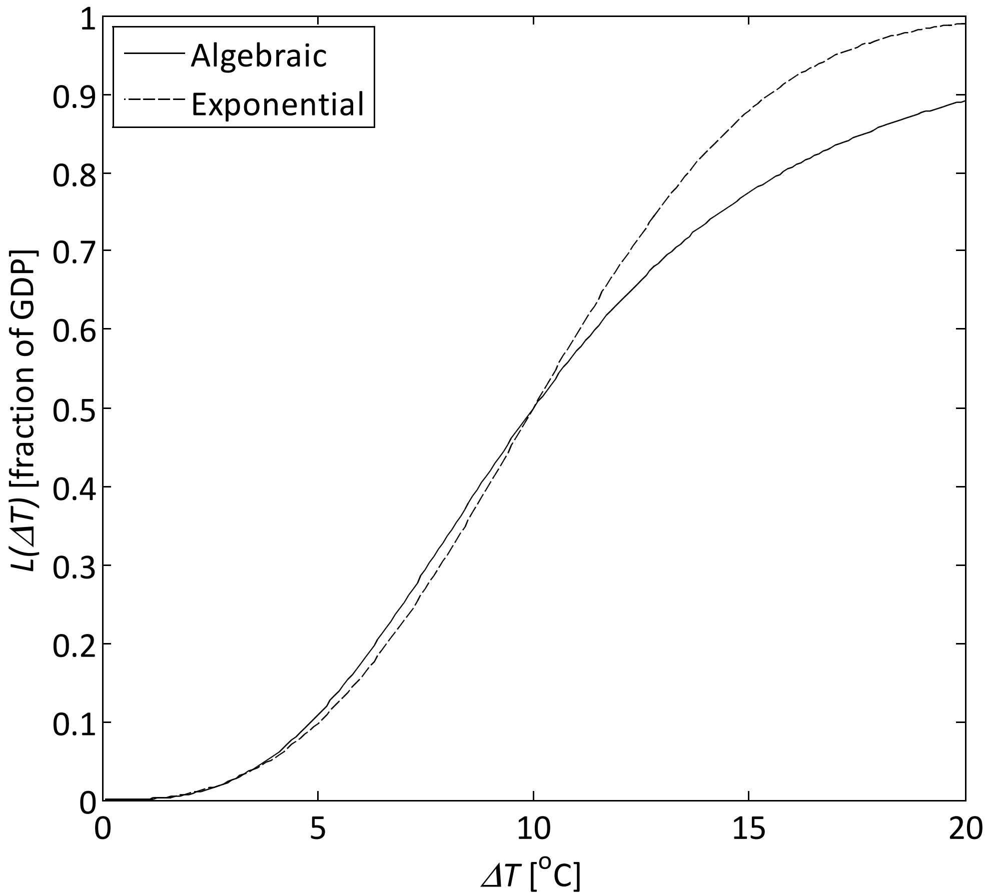


**FIGURE 2.** Two damage functions used for the stylized climate assessment model in Section 3.

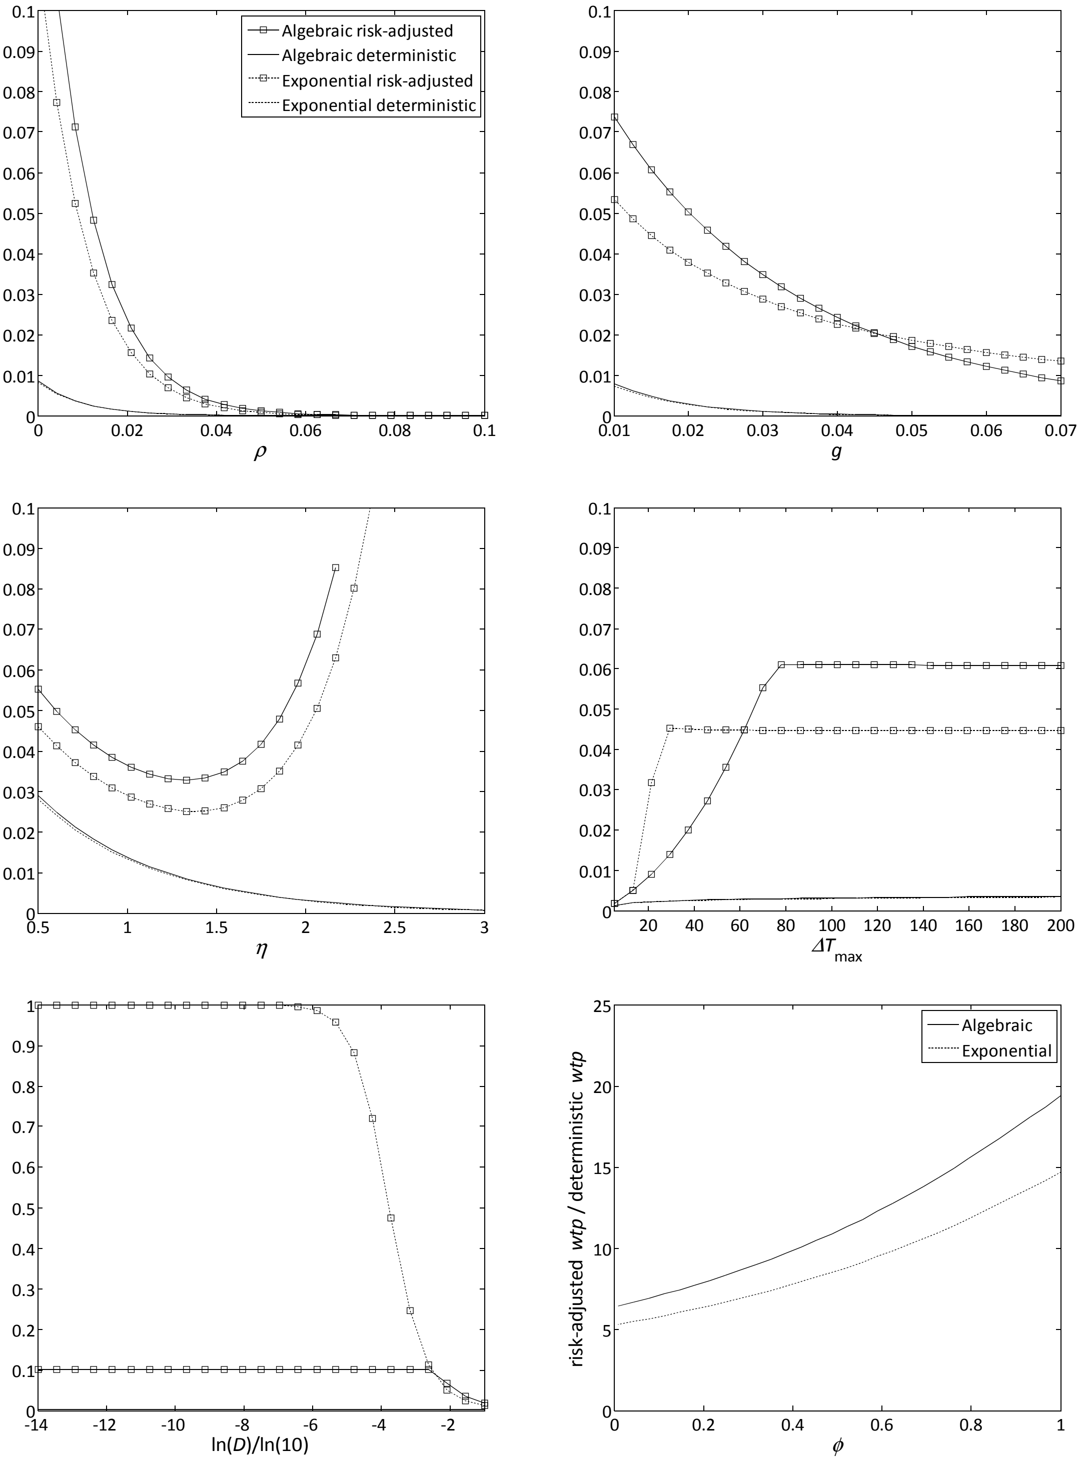


**FIGURE 3.** Sensitivity analyses using the stylized climate assessment model and the Bayesian model-averaged pdf, varying each parameter in turn while holding all other parameters at their default values, which are $\rho$ = 0.01, $g$ = 0.015, $\eta$ = 2, $\Delta T_{max}$ = 100, $D$ = 0.01, and $\phi$ = 1.

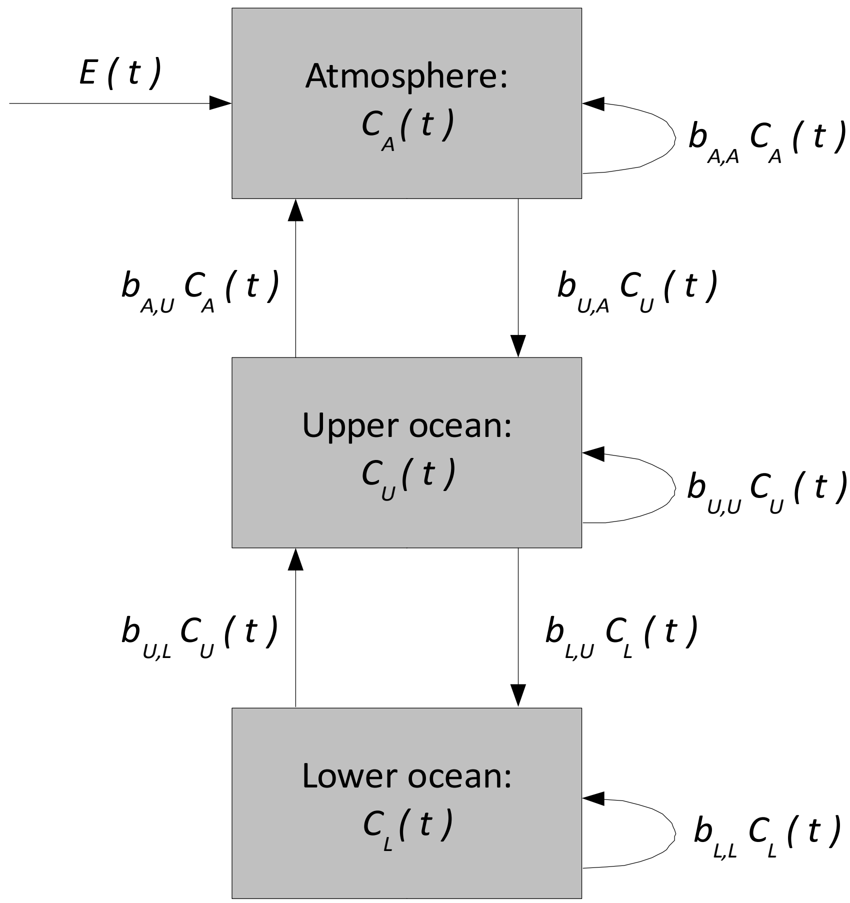


**FIGURE 4.** Schematic representation of the "three-box" model of global carbon flows in DICE. The $b_{i,j}$ parameters indicate the fraction of the carbon in box *i* that flows to box *j* during each time step.

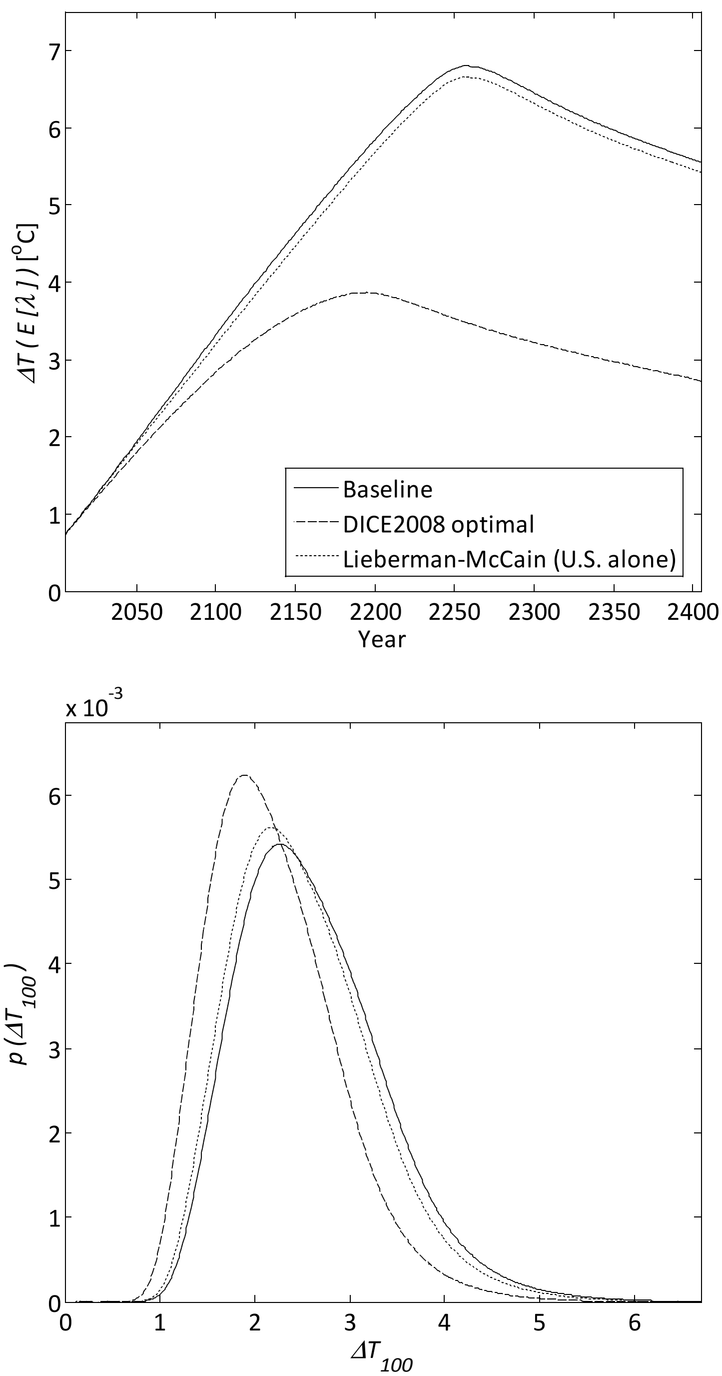


**FIGURE 5**. Predicted paths of atmospheric temperature differences using the DICE climate module, conditional on the climate sensitivity parameter being equal to its expected value (top graph), and the probability distributions over the temperature difference 100 years from now based on the Bayesian model-averaged pdf (bottom graph), for the DICE baseline emissions path, the DICE optimal emissions path, and the simulated Lieberman-McCain emissions path.

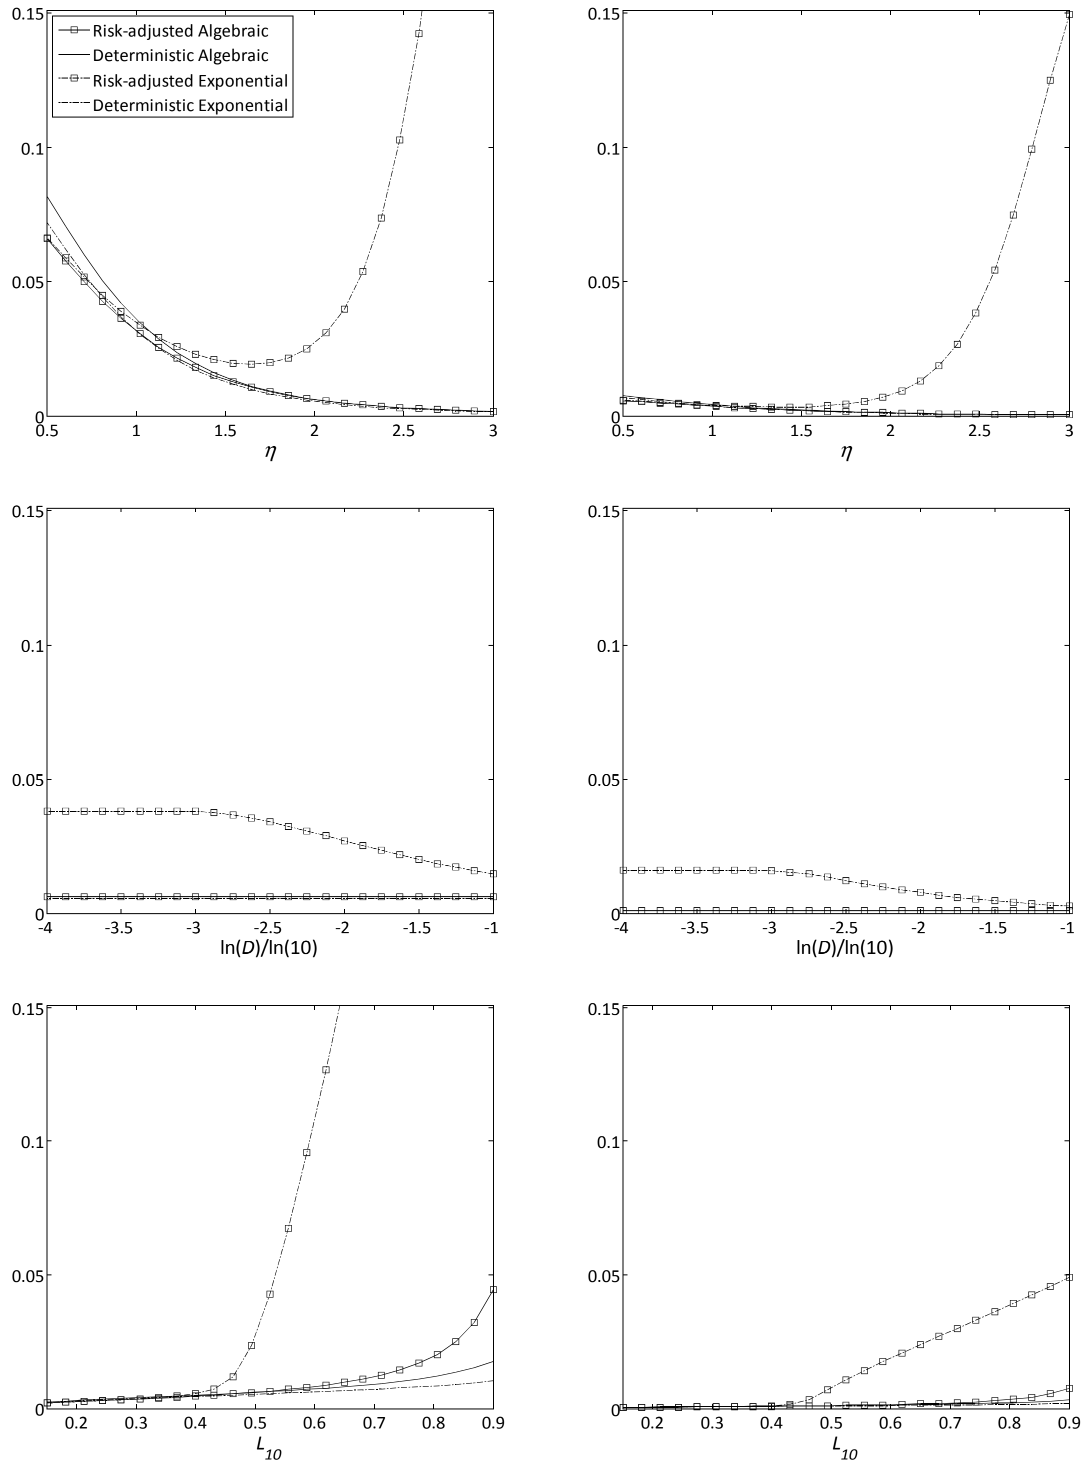


**FIGURE 6.** Sensitivity analyses using the slightly more realistic climate assessment model and the Bayesian model-averaged pdf, varying $\eta$ , $D$, and $L_{10}$ in turn while holding all other parameters at their default values, for the DICE optimal path (left) and the simulated Lieberman-McCain path (right).