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.

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 ΔT = 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 ΔT 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 ΔT_1 and with probability 1-*p* it will be ΔT_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 ΔT_2 by a small amount $d\Delta T_2$ and simultaneously decreasing ΔT_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 ΔT_2 is greater (less) than the slope at ΔT_1. Therefore, simple intuition based on Jensen's inequality alone does not suffice here since the damage function is not everywhere convex in ΔT —it is convex for low ΔT but concave for high ΔT. Over a wide range, increasing the variance of the distribution—i.e., shifting some of the mass of the pdf towards higher ΔT 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	∞	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:		
ρ , pure rate of time preference	0.01	
g, consumption growth rate	0.015	
η , 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 ΔT = 3 deg C	0.03	
L_{10} , fraction of GDP lost if ΔT = 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

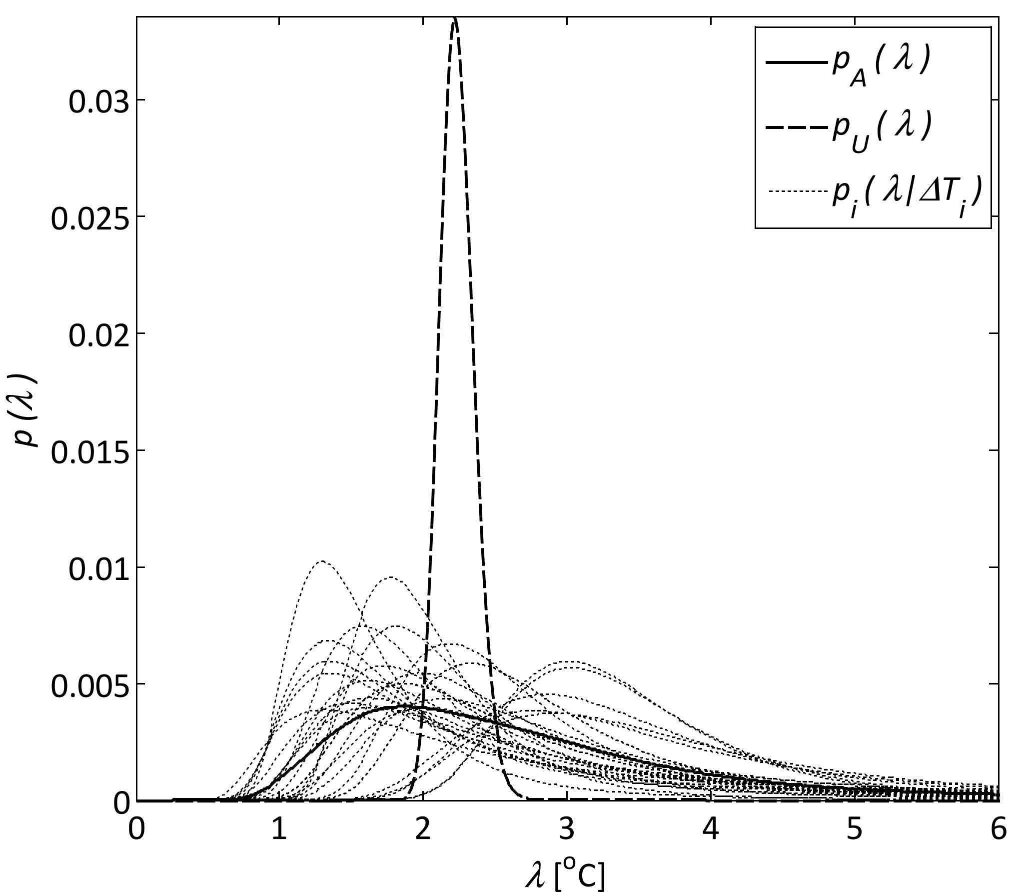

FIGURE 1. Roe and Baker (2007) probability distributions constructed from the 5th and 95th percentiles for the climate sensitivity parameter from each study shown in Table 1 (light dotted lines), the Bayesian model-averaged pdf based on the average of the distributions using equal weights (heavy solid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

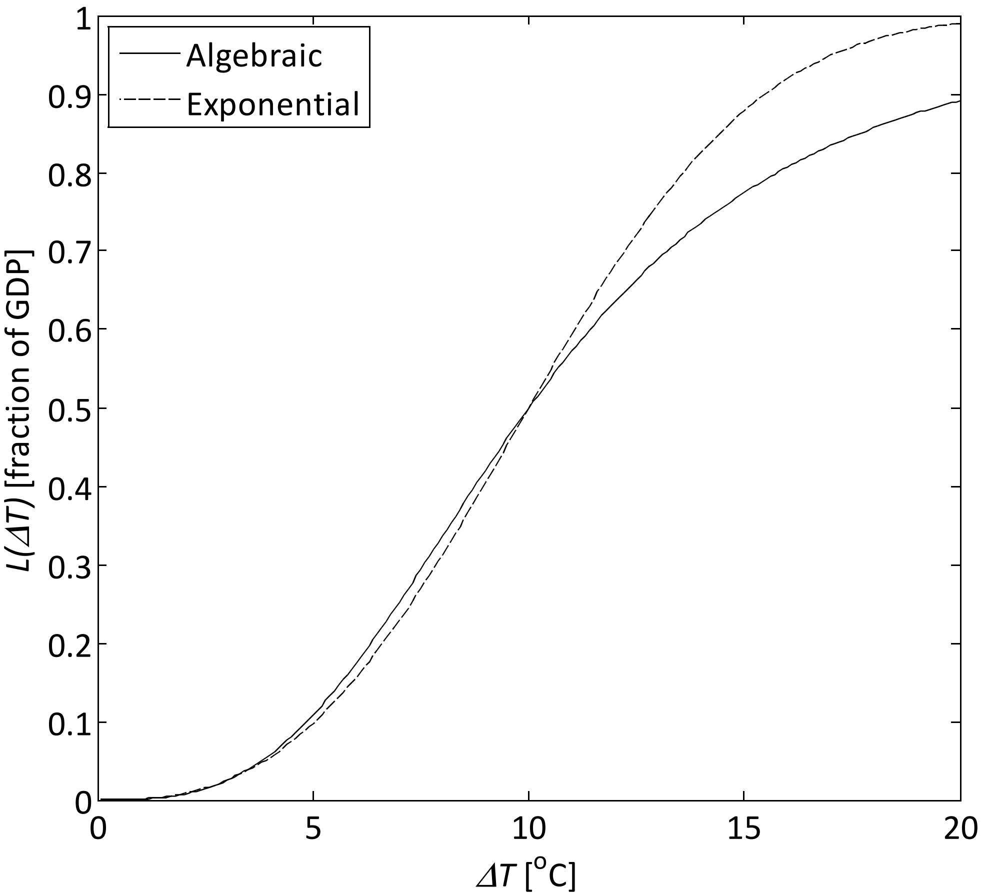

FIGURE 2. Two damage functions used for the stylized climate assessment model in Section 3.

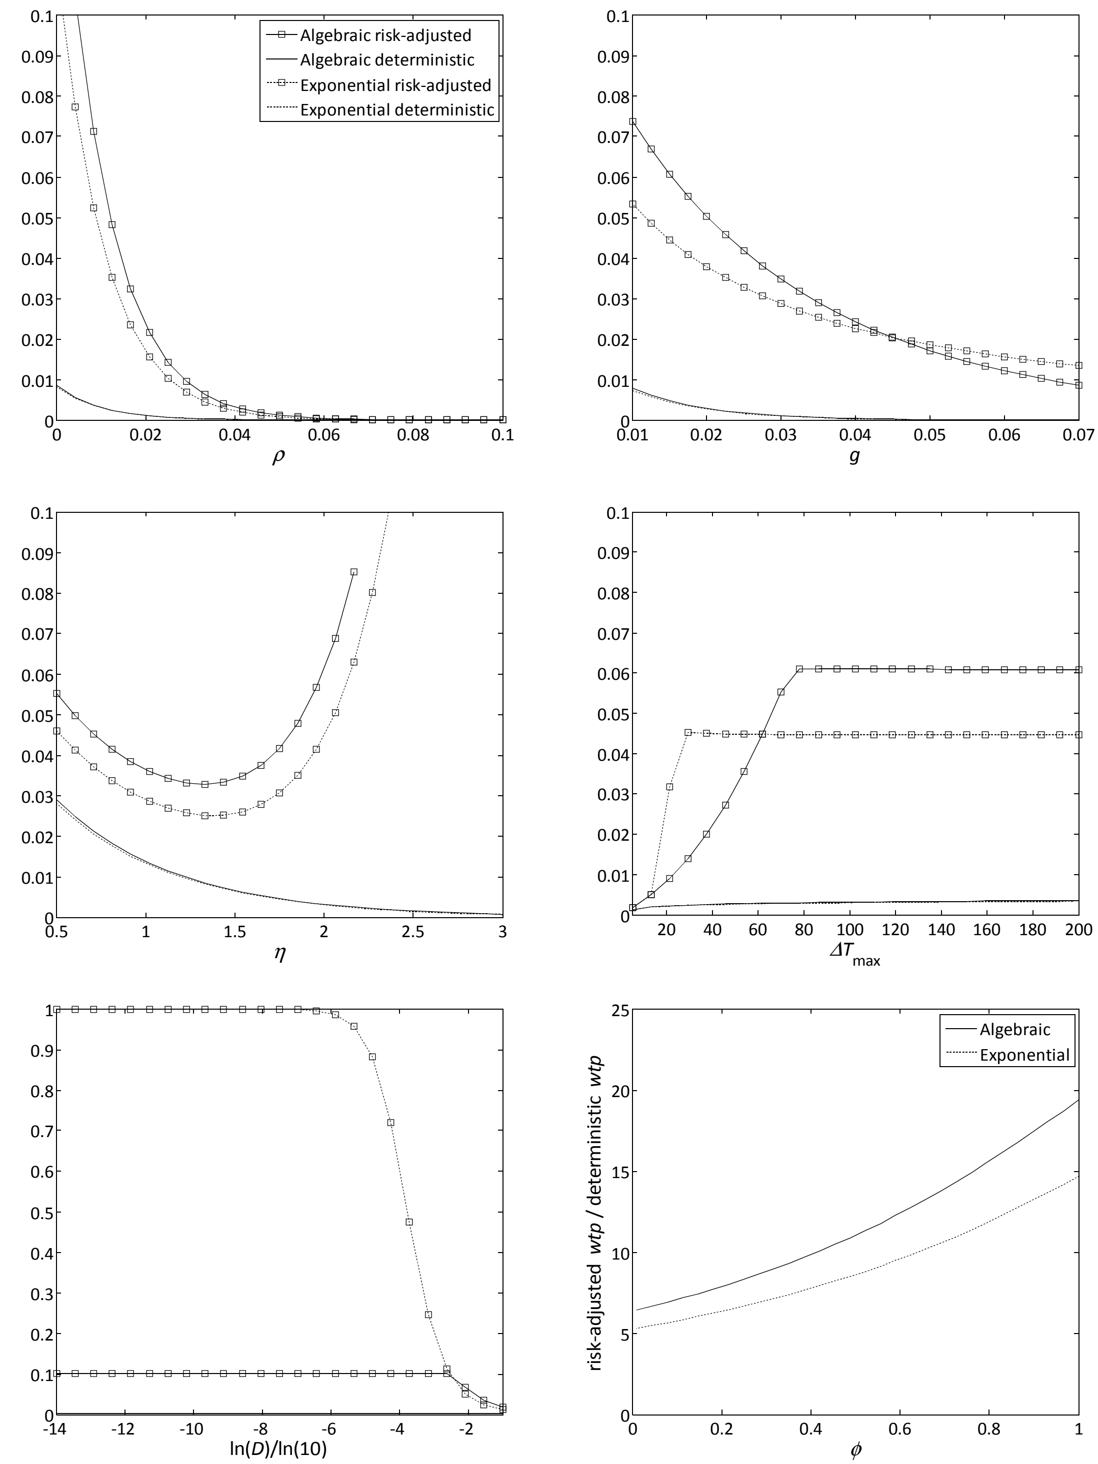

FIGURE 3. Sensitivity analyses using the stylized climate assessment model and the Bayesian model-averaged pdf, varying each parameter in turn while holding all other parameters at their default values, which are ρ = 0.01, g = 0.015, η = 2, ΔT_{max} = 100, D = 0.01, and ϕ = 1.

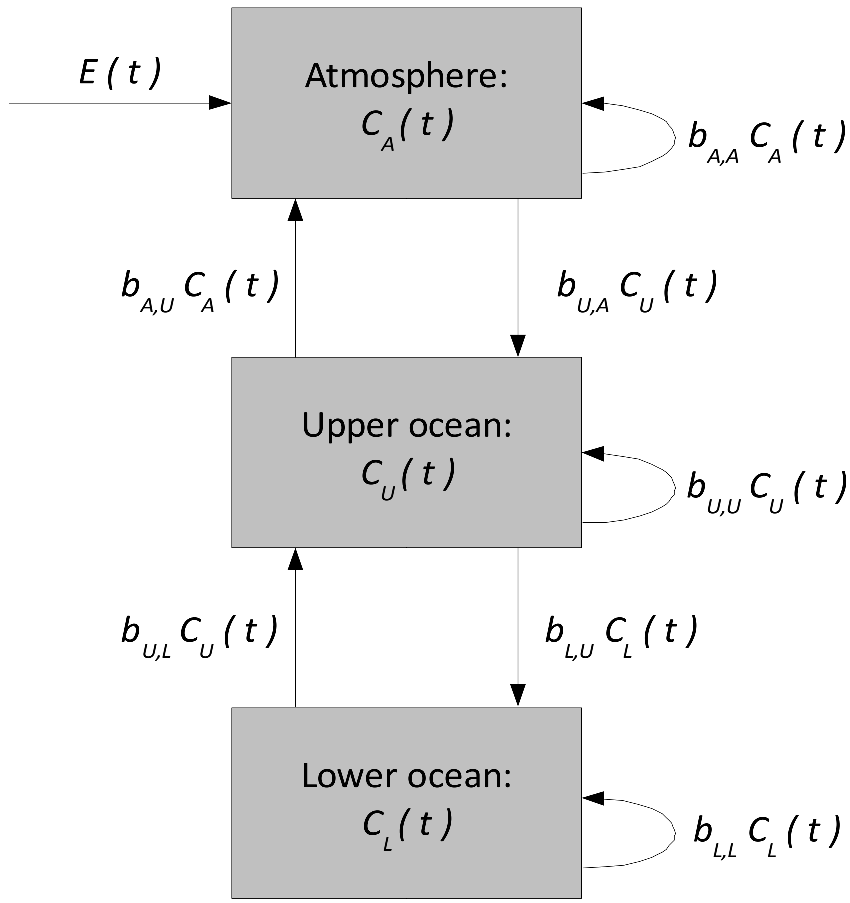

FIGURE 4. Schematic representation of the "three-box" model of global carbon flows in DICE. The $b_{i,j}$ parameters indicate the fraction of the carbon in box *i* that flows to box *j* during each time step.

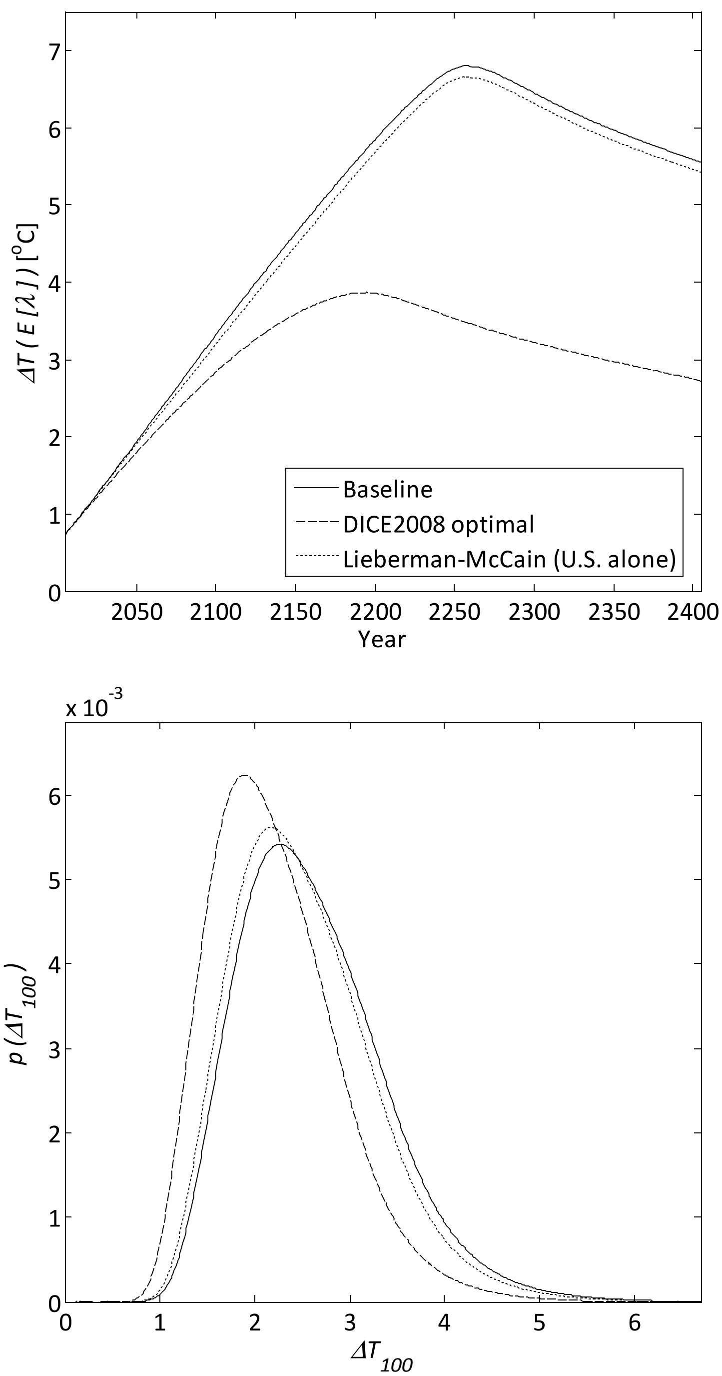

FIGURE 5. Predicted paths of atmospheric temperature differences using the DICE climate module, conditional on the climate sensitivity parameter being equal to its expected value (top graph), and the probability distributions over the temperature difference 100 years from now based on the Bayesian model-averaged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

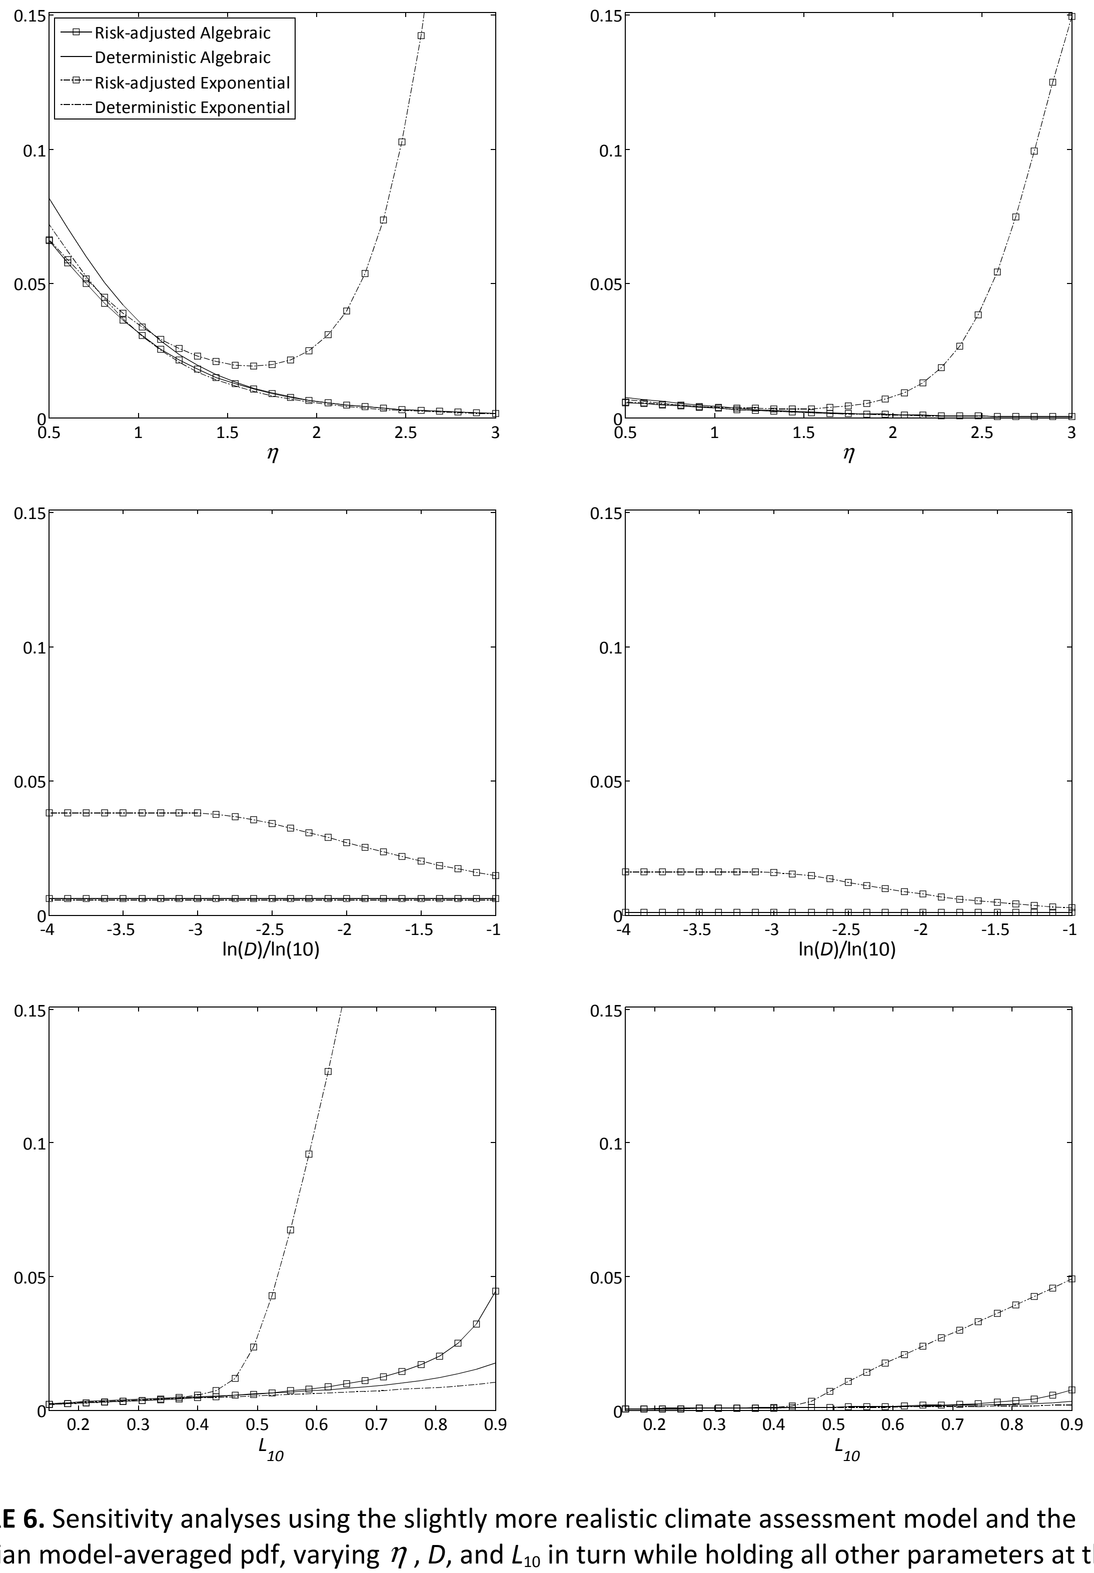

FIGURE 6. Sensitivity analyses using the slightly more realistic climate assessment model and the Bayesian model-averaged pdf, varying η, D, and L_{10} in turn while holding all other parameters at their default values, for the DICE optimal path (left) and the simulated Lieberman-McCain path (right).