Response to Referee Report 2 dp 2011-22
14 December 2011

We thank the referee for finding the time to review our paper and hope our response and revisions will alleviate some of his concerns. In the below, the referee's comments are shown in indented italics; our responses are shown in left-justified, unindented text.

> *Higher-level comments:*
>
> *Unfortunately, in its current form, this paper is extremely poor at communicating. It is opaque, unfocused, not well structured, and insufficiently motivated and justified. As a result, it was impossible to fully understand, appreciate, and review what the authors had done, and why. Overall, I would say the paper wasn't ready for review. Even so, given what I was able to deduce, the paper has a number of problematic elements.*

While the reviewer is the first among the many people with whom we've shared this paper to find it so impenetrable, we will endeavor in the revision process to make it less so.

To state the rationale for this paper:

There are a fair number of estimates of the social cost of carbon in the literature, most relying on a handful of cost-benefit IAMs, particularly DICE, FUND, and PAGE. The damage functions in these IAMs are calibrated against impact studies that are frequently dated and conducted for low levels of warming, generally in the range of a doubling of pre-industrial carbon dioxide levels. This paper does not attempt to address the dated-ness of these calibrations; rather, it for sake of argument concedes these calibrations and focuses on the issue of how to extrapolate damages to higher (sometimes far higher) temperatures -- a question which the authors of these models have not generally focused upon, but is in any case quite uncertain. Conceding the low-temperature calibration used in DICE, we investigate the sensitivity of SCC estimates to different ways of extrapolating from the calibration point to higher temperatures. We also evaluate different levels of risk aversion, since how uncertainty affects the SCC is highly sensitive to this parameter.

> *The authors provide little justification for their decisions. At the moment, they are asking the readers to simply accept what they have done. That's not acceptable. The authors need to provide clear rationale for what they are doing and sound scientific justification for their choices.*

More precision in this comment would have been helpful, since I'm not sure what assertions the reviewer is referring to here. Invited reader #1 did highlight the calibration of the uncertainties in Xc, which we agree are not well explained. As we wrote in response to that commenter:

"Since Xc is intended to be illustrative, and there is really no basis in the literature on which to calibrate the uncertainties, these standard deviations are largely arbitrary. The variance of the log normal distributions for the exponent and the calibration damages were chosen to yield a possibility of extreme poverty of slightly less than 1%, and a probability that our descendants in 2300 would be poorer than us of roughly twice this; other parameters were selected to yield reasonable indicative ranges."

> *The paper provides extraneous information and lacks details needed to understand graphics and facilitate interpretation and appreciate of results. The authors need to*

focus more squarely on what they are trying to show and why (and what they are showing in graphics). For instance, the authors' presentation of damage functions in the literature, as well as their own damage functions, are extremely poor and ineffective at giving the reader the necessary foundation for reviewing and interpreting the results and appreciating the insights. Related to this, the essential initial figures and tables have next to no discussion (figures 1-4) or scattered discussion that needs to be honed to serve the purposes of the paper (tables 1-2).

We agree that some of the early figures and tables could use more discussion and will address this concern in the revision.

The development of the authors' new damage functions is completely inadequate. Here and there across the manuscript, the reader finds a few clues as to what they are, their rationale, and justification for their construction. However, even after good detective work, readers won't be confident that they fully understand. Readers need a single well constructed discussion. Overall, the authors' methodology is unclear, unconvincing, and appears arbitrary. The presentation is far from reader friendly, and throughout the paper the authors take methodological steps (assumptions, formulation, parameterization) but provide insufficient rationale or justification for their decisions.

We will tighten the discussion of the development of Xa, Xb, and Xc in the revision, and will provide more complete mathematical formalism for some of the development in prior subsections to lay the necessary groundwork.

I think the stabilization scenario adds very little to the paper, and even distracts from other points. Therefore, I suggest that the authors drop it. First, the scenario is very ad hoc and not defensible (e.g., only CO2, probably doesn't include lulucf co2, no GDP implications, ignores socioeconomic transformation, no literature reference, hard to believe 50% likelihood with over 500 ppm CO2 only in 2100). Second, what is the motivation for including it? Global society is not on this pathway and marginal damages off of it are pretty much irrelevant to today's decision-making. The authors are better off focusing on clearly explaining what they are doing in the reference case and risk aversion and damage function uncertainty.

We agree with the reviewer that the basis for the stabilization scenario was not clearly expressed in the text. We believe some of the reviewers concerns relate more to DICE as a model than to the particulars of our calculation of a stabilization scenario with it. The reviewer is correct that, in DICE, LULUCF emissions and non-CO_2 forcings are exogenously specified and not subject to mitigation. Unstated in text (with our apologies for the omission) is that we have alleviated this limitation of DICE in our application and now allow for control of land use emissions and non-CO_2 emissions, albeit with the same single control parameter and cost function as applied to fossil fuel CO_2 emissions.

I am unclear as to what socioeconomic transformations the author would like us to represent in our stabilization scenario; most IAMs -- including all cost-benefit IAMs as well as many more disaggregated IAMs -- use an exogenously specified population for both reference and policy scenarios. As to the 50% likelihood of keeping peak temperature below 2.5ºC, I am unclear whether the reviewer thinks the model allows too little or too many emissions for this target; in any case, his quarrel is with the simple climate and carbon cycle box model used in DICE (which is slightly more sophisticated than the exponentially decays used for carbon dioxide concentration and temperature in PAGE and FUND), not with our paper. Van Vuuren et al. (2011) compared the transient climate response of DICE 2007 to a range of climate models and found that it performed reasonably toward the center of the distribution, though had a bit slower response than MAGICC.

Our stabilization scenario peaks at 593 ppm CO_2e in 2065, then decays, falling below 550 ppm in 2120 and reaching 520 ppm in 2300. With an equilibrium climate sensitivity of 3 C/CO_2 doubling, surface temperature hits 2.5 degrees warming in 2095 and thenceforth stabilizes. Note that the planet has not reached equilibrium by 2300; at that time point, deep ocean temperature is still warming and has reached 1.9 C. There is no need for a literature reference for this scenario; it is simple to compute from DICE, as we have done. There is no particular reason to believe these climatic details, though the Van Vuuren et al. comparison indicate that they are not unacceptable bad; but they are, in any case, besides the point.

The reason for including it is illustrated in the discussion is that it highlights the point that the SCC is not always lower in lower CO_2 scenarios than in higher CO_2 scenarios. This is a point made in the companion paper by Kopp and Mignone, building on a point made by Baumol (1972). If damages are non-convex, as the comparison of the reference and stabilization scenarios illustrate for Wa, and later in the century for D and Xc in the zero risk aversion case, then applying the SCC calculated off the reference path as a Pigouvian tax will not suffice to drive emissions toward an optimal level. This is an important point, but one which we agree could be made more directly; we will do so in the revision.

> *Related to this last point, the authors start their conclusion with the following: "Our analysis highlights the importance of jointly considering risk aversion and uncertainty in damages when estimating the SCC." This is a concise statement of their goal. It should have appeared at the beginning of the paper, and the paper needs to be more focused on building a story that delivers on this goal. As to whether the paper achieved this goal, I have to say no. The paper isn't clear, has too much going on, and the reader cannot understand what was done to generate the results, why it was done, and that what was done was legitimate.*

We will endeavor to make the paper clearer in the revision.

> *More specific comments:*
> *Calibration to 2.5 DegC – If I understand this correctly, this step doesn't seem scientifically legitimate. First, these are not identical damage functions (in terms of what is included alone), so it is problematic to squeeze or stretch them to fit the calibration point.*

I am not sure what the reviewer is objecting to here. The DICE damage function claims to be a comprehensive summary of the economic impacts of climate change, so all damages are in principle included. Likewise for each of the substitute specifications we present.

> *Second, this is not simply a normalization that shifts damage functions up or down. The damage functions are non-linear, therefore the calibration affects the curvature. The authors do not justify the calibration or explain the implications of the calibration (how different are the calibrated functions from the originals?).*

We will add more discussion of Figure 1 in order to address this point.

Leaving out Table 2 damage functions – There is no clear explanation for why the damage functions in Table 1 were chosen and those in Table 2 were not. In particular, Table 2 includes FUND and PAGE, which are two of the three models used by the USG, and well represented in the literature. It would argue that it is essential to include those.

In this paper, we focus on damage specifications that can be implemented in a one-region, DICE-type model, in order to highlight the sensitivity of SCC results to choices about functional form. Table 2 provides parallel descriptors for damage specifications that appear in cost-benefit IAMs other than DICE. In the context of multi-region models like PAGE and FUND, other authors have dealt with questions related to how to combine social costs experienced in different regions with different levels of vulnerability; addressing these issues in the context of this paper would expand the scope too broadly. Like DICE and PAGE and unlike FUND, we do not model different economic sectors using different functional specifications; such a choice would similarly obscure our focus, which is on how different choices about functional extrapolation affect the SCC. Nonetheless, most of the core elements of the functional forms used in all these sectors do appear – albeit within in a single sector – in the specifications we examine.

Damage function K has uncertain threshold damages, like PAGE; damage function ASB has an uncertain exponent, like PAGE; damage function Xc has both. Like FUND, damage function L has rate dependency, and damage functions SP, Wa and Xc have damages that effectively depend upon wealth.

Casual use of economic terms – consumption, output, wealth, income, welfare, utility are distinct economic concepts. For instance, C does not equal GDP. In addition, consumption is not the same as wealth and welfare, yet it is referred to as both. The paper interchanges some of these as if they are synonyms. These differences should be respected in the discussion, treatment of damage functions, and results.

We agree that we do use the term 'wealth' in a fuzzy sense, which we will strive to clarify in the revision; by 'wealthy,' we really mean 'high consumption' or 'high income.' (In the Solow growth model, consumption and income are directly proportional.)

We will also clarify some of our definitions by formally describing the economic component of DICE to clarify our discussion of utility/output/capital damages.

We are unclear as to whether the reviewer's objection is broader than the use of the term 'wealth.' We can find no occurrences in the text where we say 'consumption' and mean anything other than 'consumption.' We can find no occurrences in the text where we say 'welfare' and mean anything other than integral of discounted utility; we see no points where we use the terms 'output' and 'GDP' (which we do, admittedly, use as synonyms) and mean anything except the same. The term 'income' does not appear in the paper.

Section 2.3 – strange how adaptation only receives a modest mention at the end of this section. I expected it to be the first thing discussed. Adaptation is a very important aspect of net damages, yet it is only mentioned for a model that the authors didn't include in their analysis. More discussion is needed regarding the implications of accounting for adaptation.

We will note when we introduce DICE that adaptation is incorporated implicitly and indirectly, via the assumption underlying the sectoral studies used for calibration. This criticism seems aimed more at DICE than at the experiment we are trying to conduct.

> *Decomposition of utility, output, & capital – first, I couldn't make sense of Section 2.5 (and Fig. 5). Utility, output, and capital are not additive. Effects on capital will affect output, which will affect utility. They are distinct end-points, but they are not independent. Also, the size of the effects in one is not indicative of the size of the effects in the other. For instance, modest GDP effects (and even total consumption effects) can have large welfare effects (e.g., agriculture damages can have huge welfare implications, especially in developing countries, because food is a large share of consumption). These facts calls into question the authors Xb and Xc damage functions, which attempt to decompose damages into these three things. This decomposition, as far as I can tell, is theoretically unsound and, and given the way I think I understand it was implemented, completely arbitrary. The authors would be on more solid ground if they abandon Xb and Xc and focused on Xa and damage function uncertainty.*

We agree with the reviewer that utility, output, and capital are not additive. As we write in the manuscript:

"Damages to utility impact well-being but not the growth of the material economy. Damages to output leave current capital untouched but reduce investment and therefore future capital and output. Damages to capital will produce the same investment reduction as damages to output, while also impacting current capital. Consider the effect of a short but very severe shock to each of these factors: a severe shock to utility will temporarily make people very unhappy; a severe shock to output will cause a recoverable depression; and a severe shock to capital will require a protracted period of rebuilding. For the same damage function, the SCC will therefore be higher the greater the proportion of damages accruing to capital, and lower the greater the proportion accruing to utility (Figure 5)."

We will formally show how we divide damages among capital, output and effective consumption so that their total damages to utility remains at a specified level. (Essentially damages to capital go as $(1 - (1-D)^{f_{\text{capital}}})$; damages to output go as $(1 - (1-D)^{1-f_{\text{capital}}-f_{\text{utility}}})$, and damages to effective consumption go as $(1 - (1-D)^{f_{\text{utility}}})$, where f_{capital} and f_{utility} are the fractions of total damages to utility caused by the respective sources.)

> *3% discount rate (flat or average) – both of these are problematic due to inconsistency with changing annual growth over time, which implies changing consumption discount rates over time. Matching up to the USG constant 3% is not a good argument, since the constant rate was a poor choice in the first place due to this inconsistency issue. Also problematic is simultaneously adjusting the pure rate of time preference. You don't have a true risk aversion experiment this way.*

As noted in our response to invited reader comment 1 (and as we will make clear in the revised manuscript):

"We agree that our choice of values for η and ϱ are partially influenced by the desire for consistency with the central discount rate used in assessing climate change impacts in US regulatory impact assessments. This is a deliberate choice, as our intent is for this paper to be relevant to US policy discussions. But there is also a good theoretical basis for our approach of calibrating ϱ and η to maintain (in the absence of climate change) a fixed average interest rate. Namely: we take as given the statement (implied by the US government scenarios) that risk-free interest rates will average 3% between 2015 and 2115, and we explore the different consumer preferences (in terms of elasticity of the marginal utility of consumption and rate of pure time preference) that might explain this "observed" interest rate. This interest rate might result entirely because consumers value the future less than the present (i.e., η could equal 0); or it might result entirely from consumer expectations about future wealth (i.e., ϱ could equal 0). Varying ϱ and η independently would necessitate changing the constraint provided by this observable (scenario-specified) "fact.""

> *The paper redefines the SCC (initially, p10, p14) – this is confusing. If different concepts are being defined (which appears to be the case in at least one case), it is far from clear. They need to be more clearly defined and distinguished.*

We will consolidate our definitions of the SCC into the introduction; the definition given in the introduction ("The social cost of carbon (SCC) is a monetized estimate of the change in social welfare that results from a marginal change in carbon dioxide (CO_2) emissions.") and the more technical definitions on page 10 ("the ratio of the change in expected welfare from a unit of emissions to the change in expected welfare from a unit of material consumption in the period of emissions") and 14 ("the change in expected welfare from a unit emission of carbon dioxide in a given year, normalized to change in expected welfare from a unit of consumption in the same year") refer to the same concepts. On page 14, we reintroduce the definition for the purpose of distinguishing the SCC from the deterministic SCC, which is the SCC conditioned upon a specific state of the world.

> *Cost-benefit IAMs: this is not a good term. It is misleading. First, some of these highly aggregated IAMs do not compare cost and benefits. Second, and more importantly, the SCC estimates being generated in this paper are not estimates for economically optimal mitigation pathways. A more accurate label for these models would be "highly aggregated" IAMs (to distinguish them from the IAMs that do stabilization cost-effectiveness analyses that identifies potential energy macroeconomic transformation pathways).*

While "highly aggregated" IAMs is certainly a valid descriptor of the class of models we are referring to, we mildly prefer the term "cost-benefit IAMs," because the key characteristic for the purpose of this paper is that they can estimate the benefits of climate change mitigation as well as the costs. That is, they include estimates of the whole-economy impacts of climate change. The more disaggregated models, which have much greater detail and perform better for estimating the costs of transformation pathways, could be considered "cost IAMs;" they do not generally attempt to include mitigation benefits/climate change damages. In principle, the economic benefits of climate mitigation could be incorporated into disaggregated IAM, in which case the two dichotomies (cost-benefit/cost and highly aggregated/more disaggregated) would no longer be equivalent. We will explain our terminology in a footnote.