Referee report

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.

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.

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). 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. 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.

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.

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. 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?). 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.

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.

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.

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.

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.

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.

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).