

Interactive comment on “On the Future Role of the most Parsimonious Climate Module in Integrated Assessment” by Mohammad M. Khabbazan and Hermann Held

Anonymous Referee #3

Received and published: 27 November 2017

General comments

This manuscript investigates the performance of a one-box energy balance model (PH99) as an AOGCM emulator for strong mitigation scenarios. The authors find that this simple climate model (SCM) consistently over-predicts future temperatures when the ECS and TCR are transferred directly from AOGCMs. Fitting the PH99 directly to the AOGCM temperature time series eliminates this bias, and reveals that the AOGCMs time series imply a substantially lower ECS and higher TCR than what they had transferred directly. The manuscript briefly discusses the physical interpretation of this discrepancy, and also explore alternative ways of fitting the one-box model that

C1

might be more reasonable for extrapolation in parameter space (of the kind performed when these SCMs are used to investigate the optimal dynamic behaviour of a decision maker under uncertainty).

Before continuing further, I want to briefly highlight two important factors that might reasonably affect how you read this review. First, I have not only read the manuscript, but also the previous reviews and the responses from the authors. My comments primarily address the manuscript itself, but I will also sometimes explicitly agree or disagree with comments that have been made earlier in the process. Second, I have not approached this manuscript as a climate physicist, but rather from the perspective of a researcher who uses the integrated assessment models with SCMs like PH99. My comments will therefore differ in spirit from those of earlier reviewers, and I focus more on issues I believe to be more relevant to those who would use this research.

Specific comments

My overall assessment is that this manuscript offers an interesting contribution and should be published. A previous reviewer expressed concern that the scientific contribution may be inadequate, but I feel that this comment does not adequately consider the policy influence that the one-box model wields (or rather, simple integrated assessment models that use PH99 in one form or another). For instance, two of the three climate-economy models used by the US federal government to calculate the social cost of carbon incorporate one-box energy balance models (notwithstanding recent political developments). The policy analysis in the *Stern Review*, which was commissioned and used by the UK government to formulate climate policy, was also based on a coupled climate-economy model that incorporated a one-box energy balance model. By their simplicity, these SCMs are also have come to serve as tools for translating new climate science for communities that use climate information but generally lack

C2

extensive physics training. Even a relatively small improvement in our understanding and handling of these models would provide a significant contribution.

I do have some concerns about the manuscript, though. First, I think there are parts of the manuscript that will be difficult to decipher for many of the researchers that actually work with SCMs in the context of simple climate-economy models. Second, I think that authors have tended to focus excess attention on concerns related to interpolation and extrapolation of parameter values, at the expense of a fuller and clearer discussion of the physical interpretation of their primary findings. I discuss each point in turn, and offer a few minor comments at the end. Let me state clearly, though, that I expect these concerns can be fully redressed, so I wouldn't consider them reasons for rejecting the manuscript.

1. The heart of this manuscript, as I see it, is the direct transfer of AOGCM characteristics (section 2.1). The central issue is whether or not it is appropriate to use this physical method for deducing the parameter values in one-box model. The subsequent question about whether alternative methods for fitting the parameter values perform better, is also tied to this baseline method. So section 2.1 is really the foundation for all of the analysis in this paper. Yet two crucial pieces seem to be missing from it.

First, at this point in the manuscript the authors should be offering a childishly clear explanation of how (and which) AOGCM outputs can be used to deduce the values of α and μ , which can then be plugged into equations (2) and (3) to retrieve the implicit ECS and TCR, respectively. But I must admit to having some difficulty following their derivations (e.g. not understanding how h is determined in equation (7) where both h and μ appear to be unknowns, and not seeing any expression for α in terms of AOGCM output). A climate physicist will perhaps be so familiar with this material as to be able to perform these calculations with little prompting from the authors, but as a presumptive member of the intended

C3

audience, I would appreciate it if the authors exercised greater pedagogical care in this section.

Second, and just as important, is that the authors have not offered any information to suggest that this is how modellers are currently choosing values for the ECS and TCR. In my experience, many users will not themselves try to deduce these parameters from AOGCM outputs, but rather plug in values of ECS and TCR reported in IPCC chapters or specific academic papers without fully understanding how these values are inferred from AOGCM simulations (and sometimes adding a bit of 'calibration' to make sure the results don't look too dissimilar from MAGICC, say). If those reported ECS and TCR values are derived in this way generally, the authors should state this clearly and cite examples. If not, they should consider whether it would be more appropriate to use a different baseline.

As a suggestion, I think it would be worthwhile to run the model using the actual parameter values assumed in FUND and MIND as a baseline (and maybe PAGE, which the authors seem to have ignored, even though it incorporates a one-box model). If I were to speculate, I would guess that using the default parameter values from these models will give even more discrepant predictions, so in addition to being more relevant to current practice, it might illustrate your point even more clearly.

2. SCMs are used for two distinct purposes: (1) as devices for summarizing and communicating climate science to other modelling communities, and (2) as computationally efficient AOGCM emulators. The analysis performed in this manuscript has important implications for both uses, but the authors are failing to distinguish clearly between them. This creates unnecessary confusion (seen especially clearly in the exchanges with previous reviewers), and has in my opinion led to an unbalanced treatment.

The authors appear to recognise the role of PH99 as a communication device when they, in their Discussion (section 5), briefly mention the idea that the 'tran-

C4

sient climate sensitivity' might be lower than the ECS, as a potential physical explanation for the lowering of the ECS when the parameters are calculated by fitting the one-box model to AOGCM temperatures instead of derived from AOGCM forcings. But what is the chief physical mechanism behind this? And does this mean that PH99 users should interpret $\frac{\Delta T}{\Delta F} \ln(2)$ as the 'transient climate sensitivity' rather than the ECS? Can we do this without undermining the physical basis for the one-box model? And what about the higher TCR value that you get when fitting the one-box parameters instead of transferring them directly? What is the physical interpretation of this second important change? You also acknowledge that measurement error in AOGCM outputs could lead to biased values of the ECS and TCR in the one-box model, but can you do anything to show that the biases would actually go in the direction of inflating the ECS and deflating the TCR? Perhaps you could just add a random sample of Gaussian deviations to your input data and feed them through the non-linear PH99 mapping to see what the resulting distribution of ECS and TCR would look like?

I realise I have given you a lot of questions to answer, but I really do feel that this part of the discussion has been unduly neglected, and the paper would benefit greatly from extending it. It seems a very interesting fact that, for a given TCR, the ECS value transferred directly from an AOGCM is systematically higher than the value that would yield the best fit to that same AOGCM (and vice versa for TCR). Anything the authors are able to do to help the reader understand the causes of the differences between fitted and transferred parameter values, and how this might affect the physical interpretation of the one-box model parameters, would be very welcome.

I would, compensatingly, recommend shortening the discussion of the second use of PH99, as an AOGCM emulator, which currently takes up the majority of sections 4 and 5. I think it is interesting to consider the advantages and disadvantages of alternative methods for choosing ECS and TCR for emulation purposes,

C5

but I often felt lost in this discussion and think it can be done more concisely. The fitting method, perhaps unsurprisingly, does a pretty good job of fitting the AOGCM temperature time series. But the key drawback of the fitting-method is that it's inappropriate for obtaining probability distributions for the ECS and TCR that can be used to simulate PH99 under uncertainty. These kinds of simulations are now standard practice for economic assessments of climate policy based on coupled climate-economy models, so this is indeed an important issue to wrestle with.

The quadratic and cubic fitting in Figure 5 seems useful mostly as a cautionary example of 'what not to do.' The authors already explain that it's likely to lead to unphysical parameter values, and as a previous reviewer pointed out, the curvature seems largely a consequence of a single AOGCM run with a low ECS. Overall, the authors can probably devote less space on this particular exercise and be even clearer that it is ill-advised. Instead, they should focus on the more physical interpolation/extrapolation methods considered in section 4, and try to offer users more concrete advice about when they might prefer the Lorenz curve method, or the ECS-to-ECS and TCR-to-TCR fit, or the ECS/TCR-to-ECS fit, or when all three are likely to perform poorly.

I think a slight reorganization of sections 4 and 5 would probably be the most effective way of accomplishing all of this. The new section 4 would take the first two paragraphs of the current section 5 as its starting point, but elaborate along the lines I have discussed above in order to offer a discussion of the physical interpretation of fitted ECS and TCR relative to the directly transferred ones. The new section 5 would merge the current section 4 with the remainder of the current section 5, in order to offer a discussion of the appropriate and inappropriate ways to interpolate and extrapolate ECS and TCR values in PH99, in light of their physical reinterpretation in the new section 4. This separation would also make it much clearer how the choice of parameter values for PH99 depends on whether

C6

one is using it as a communications device or as an emulator.

Technical corrections

1. p. 1, line 9 (and throughout): The manuscript refers to FUND and MIND as two coupled climate-economy model that employ a one-box model. PAGE does as well (see discussion in Calel & Stainforth, 2017, BAMS, already cited).
2. p. 2, line 30: Typically, these models are used to study optimal climate policy, so it would be good if you could cite a few studies where these models are used specifically to study 2 degree stabilisation scenarios.
3. p. 5, line 12: Typo. "APGCM" should be "AOGCM."
4. p. 5, lines 17-22: While RCP4.5 is certainly out-of-sample, it's less obvious to me that it serves the purpose of validating the method for 2 degree stabilisation scenarios. As a validation exercise, wouldn't it be preferable to fit α and μ using RCP4.5 and then drive the one-box model using the RCP2.6 forcings?
5. p. 8, line 11: Typo. "againsta" should be "against a."
6. p. 8, line 15: Typo. "radiative active" should be "radiative activity" or "radiative forcing."
7. p. 8, line 19: I think it's inaccurate to say that "studies based on PH99 implicitly worked with ECS values that were larger than announced." I think you've made the point that they might be using a higher ECS than would be appropriate for emulating AOGCMs, but this is quite different. They declare their ECS values, and the question raised in this manuscript is whether they shouldn't be using ECS but rather some 'effective ECS' or 'transient climate sensitivity' instead. Please rephrase this.

C7

8. p. 8, lines 20-28: The discussion of log-Normal distributions seems to come out of nowhere. I think the reorganization I have suggested above may resolve this, but please make an effort to link this more strongly to the rest of the discussion.
9. p. 8, line 29: Typo. "boefore" should be "before."
10. p. 10, line 7: Typo. "generally" should be "it generally."

Interactive comment on Earth Syst. Dynam. Discuss., <https://doi.org/10.5194/esd-2017-40>, 2017.

C8