Interactive comment on "Constraints on long term warming in a climate mitigation scenario" *by* Benjamin Sanderson

Ben Sanderson

sanderson@cerfacs.fr

Summary: This study explores the long-term warming of climate for heavily mitigated scenarios, discussing the reasons behind changes in the sensitivity of warming to cumulative carbon emitted over time. An efficient model ensemble is generated and integrated with both prescribed concentration pathways (e.g. Figure 1), and an adaptive pathway algorithm to generate warming scenarios that restore towards the desired warming targets (e.g. Figure 2).

The study discusses how the required mitigation efforts are affected by changes in the

sensitivity of surface warming to cumulative carbon emitted over time (the TCRE), and how a constant TCRE framework may not be able to account for such effects.

I found the study, as currently written, difficult to place in the context of existing literature. In particular, this study is missing comparisons to existing literature using observationally-constrained ensembles to explore the future carbon budget, adaptive mitigation pathways, and the time evolutions of effective climate sensitivity and transient climate response to emission.

Before such comparisons are made, it is difficult to say what in this study is new, and how it fits within existing knowledge.

Many thanks to the reviewer,

I have made efforts to clarify the framing in the revised version. The central focus of the study, and novel aspects are:

1 - the adequacy of TCRE-based carbon budgeting for temperature stabilization targets and the potential for hysteresis in the cumulative carbon-temperature relationship.

2 - how including different types of common prior assumptions (which vary across the literature) in the Bayesian model can alter the appropriateness of meta-frameworks for policy such as carbon budgeting.

The use of adaptive pathways and the simple model itself were never intended to by the novel aspects of this study (the adaptive mitigation pathway concept follows previous studies - Sanderson et al 2016, Sanderson et al 2017). That said, I apologise

for missing important comparisons in the methodological aspects of the study. I've endeavoured to place the revised version in the context of the studies raised - which are certainly relevant.

In particular, the Goodwin studies are very relevant and make good efforts to quantify the effects of prior assumptions - but they do not focus on the question of reversibility and hysteresis. However, there are a number of studies and approaches in common usage which make stronger structural assumptions (fixed lambda or RWF) which are used to justify the cumulative emissions budgeting framework. The point of this study is to examine those assumptions, and how they influence model dynamics if imposed.

Major points: Significant areas of existing literature missing from discussion.

1. Observation-constrained ensembles exploring the future carbon budget

The study as currently written is missing comparisons to existing literature on observation-constrained calculations of future carbon budget for this century (Goodwin et al., 2018a) and out to year 2300 (Goodwin et al., 2018b) generated using the WASP model.

Many apologies for these omissions. I now discuss the papers in both results and introduction.

More details are required for the method used here for generating an observationally constrained ensemble. A full methodology needs to be presented containing prior assumptions, observational constraints and how the observational constraints are ap-

plied.

The method has been expanded, and now incorporated into the main text body.

This method used here should then be compared to the Monte Carlo plus history matching method presented in Goodwin et al. (2018a).

This is a good point. I now discuss the Goodwin paper at the start of the optimization chapter - because it represents an important design choice. In FAIR, at least - the Python code requires seconds to run - which made MCMC quite impractical to optimize the model (though a brief inspection suggests that in WASP, as compiled C++ code, might be fast enough).

The solution I proposed was to recode the core elements of the FAIR as a pulse-response model, fast enough that full MCMC is feasible to estimate the posterior (i.e. thousands of simulations per second) - which avoids the design issues of history matching (difficulty in automation of the selection of cutoff values and sampling uniformly from the posterior).

But - I accept that my approach has downsides, the model is an accurate representation of the FAIR core dynamics - but every aspect of the model and forcing must be represented in the pulse/response framework, which might limit further developments (e.g. a more detailed breakdown of non-CO2 forcers

Note that both Goodwin et al. (2018a) and (2018b) studies adopt an efficacy on the ocean heat uptake, which is equivalent to allowing the effective climate sensitivity to change over time (where the ocean heat uptake efficacy is greater than 1 the effective climate sensitivity in the present day is less than the equilibrium climate sensitivity

on multi-century timescales, and where the efficacy is less than 1 the effective climate sensitivity for the present day is greater than the equilibrium value). As such, both studies allow the effective climate sensitivity to vary implicitly over time, and do not assume a relationship between TCR and ECS like the FaIR model studies (an approach which this study critiques).

I've now explicitly noted in the discussion that WASP does not contain this prior assumption.

The fact that no relationship is assumed between TCR and ECS in this study is currently discussed, but to assess whether the method is novel is needs to be compared to the WASP model methodology, for which this is also true.

The novelty in this study is not the omission of the TCR/ECS relationship - but the point that including it or not has large implications on the robustness of the TCRE framework.

2. Comparison to adaptive mitigation pathway algorithms in the literature

The Goodwin et al (2018b) study in Earth's Future presents and uses an adaptive mitigation pathway approach to restore a large ensemble of observation-constrained efficient model simulations to 1.5 and 2.0 âŮę C targets – in a very similar manor to the results presented in this study in Figure 2. Given the similarity of the method, the results in this study should be compared to this existing Goodwin et al (2018b) study in the literature. The 'Adaptive scenario design' used here (Figure 2, Appendix A1.2) should be compared to the 'Adaptive Mitigation Pathway' algorithm presented in Goodwin et al (2018b) and used in Brown et al. (2018) and Nicholls et al (2018). The resulting model output in this study, for compatible carbon emission pathways, should then be compared to the similar output generated in these previous studies in the literature.

I have now cited Goodwin 2018b as also using an adaptive scenario design. However, the approach considered here follows our earlier works, which predate the Goodwin studies (Sanderson et al 2016, used in Sanderson et al 2017) - which detail the methodology for adaptive mitigation pathways used here. As such, I didn't consider this aspect of my present study to be novel. I am simply applying an established approach to produce idealized pathways. Apologies if this was unclear in the previous version.

There are, however, some interesting distinctions between the two approaches. Sanderson 2016 is 'forward looking' - i.e. scenario parameters are iteratively adjusted until targets are met, while Goodwin 2018b represents to some degree the state of knowledge of decision makers based on observed TCRE. I've noted this at the start of Section 2.2.

3. Comparisons to existing literature on the reasons behind continued warming after emissions cease and the non-constancy of the TCRE

This study is missing comparisons to existing literature:

on the reasons behind continued surface warming after emissions cease (e.g. Frölicher et al., 2014; Williams et al., 2017a)

These are now cited in the introductory paragraphs on ZEC.

on the reasons behind near-constancy (or otherwise) of the TCRE (e.g. Goodwin et al., 2015; Williams et al., 2017b).

I now cite both papers in the introduction in the paragraph introducing TCRE.

One of the potential reasons that TCRE is non-constant is a change in effective cli-

mate sensitivity. However, other possible reasons are discussed in the studies mentioned here. The reasons behind nonconstancy of the TCRE in the model simulations presented here should be quantified in a way that relates to previous studies in the literature such as these.

This is an interesting point. The pulse-response framing here and the gradient framing of Goodwin 2015 potentially give two different perspectives on the role of the ocean which are perhaps complementary. In the Goodwin framework - the TCRE at a given time is decomposed into the dependence of surface warming on radiative forcing, the fractional dependence of radiative forcing from atmospheric and the dependence of radiative forcing from atmospheric CO_2 on carbon emissions, each of which evolve in time and whose interaction explains why TCRE remains constant (or not).

The Greens function/pulse response method used here could provide a different framing, with discrete responses to emissions which emerge on different exponentially decaying timescales. The model is already defined in terms of the set of exponential decays which describe p_CO2 as a function of emissions, and a second set which define warming as a function of p_CO2. The response parameters do not evolve in time because the time evolution is coded into the pulse response. As such - one can clearly see the effect of a prior on timescales of model response (e.g. imposing the RWF prior suppresses the long timescale sensitivity parameter).

I see these two frameworks as being complementary, the advantage of this approach being that it provides fixed parameters which can then be estimated for the real world (rather than time-evolving diagnostics) - and have added a paragraph to this effect in the discussion. A comprehensive use of the Greens function description as a framework for decomonsing system response to be compared the Goodwin framework is beyond scope in an already long paper, given it's not the focus - but it would make for

an interesting followup which I'd be keen to work on.

4. Constraints on the time-evolution of effective climate sensitivity in the literature

One effect leading to a change in TCRE over different response timescales is that the effective climate sensitivity also evolves over different response timescales. Again, here the discussion is missing sections of the literature.

For example, Goodwin (2018) generates a large ensemble of model simulations with explicitly time-evolving effective climate sensitivity, and then uses historic observations to constrain how effective climate sensitivity evolves on different response timescales. This reveals an increase in effective climate sensitivity over time (Goodwin, 2018) that will, crucially for this study, affect future mitigation requirements for a given warming target. Rohling et al. (2018) presents a similar account of time-evolving effective climate sensitivity from a paleo-perspective. These studies should be discussed and the results of this study compared to these previous findings.

I now discuss this study in some length in the discussion as a proposal for future work. Broadly - the non-stationarity of EffCS and TCRE mean that attempts to quantify these parameters from observations should be qualified - and that an alternative is to explicitly calibrate carbon and thermal feedbacks on different timescales. To do this properly requires future work - specifically considering additional information (other than global mean evolution) which might be used to constrain system response at different timescales.

It should also be noted that unlike the FaIR model methodology that is currently dis-

cussed for comparison (e.g. section 2, Lines 170-185), the WASP model methodology in Goodwin et al. (2018a; 2018b) and Goodwin (2018) does not assume a prior relationship between TCR and ECS or a near-constant TCRE. Therefore, comparisons to the WASP methodology deserve a separate discussion in section 2.

Now noted - I broadly cite the Goodwin 2018 approach as being an example of using a geological prior on ECS, in contrast to FAIR which uses a prior on RWF.

References:

Sanderson BM, O'Neill BC, Tebaldi C. What would it take to achieve the Paris temperature targets?. Geophysical Research Letters. 2016 Jul 16;43(13):7133-42.

Sanderson BM, Xu Y, Tebaldi C, Wehner M, O'Neill BC, Jahn A, Pendergrass AG, Lehner F, Strand WG, Lin L, Knutti R. Community climate simulations to assess avoided impacts in 1.5 and 2 C futures. Earth System Dynamics. 2017 Sep 19;8(3):827-47.