

RC1 - This is a fantastic paper. The work that the authors have done has the power to make complex geoengineering simulations more accessible. The development and exploration of the methodology are done quite well. I do have some comments:

AC - We are pleased to hear you think our work matters and is performed quite well.

RC1 - I don't have a sense of pros and cons, i.e., when this simpler method would work versus when you need a more complex model like WACCM. I don't expect anything thorough, but if the authors could provide some opinions on this, it would be helpful.

AC - Yes, that is a good suggestion. We will add our opinion on this in the revised version. Roughly speaking, users more interested in tropospheric or ocean processes who need to limit computational costs may want to use our method, whereas researchers interested in stratospheric circulation or microphysics ought to use WACCM.

RC1 - There are situations where the authors claim to have explored a variety of scenarios and conclude that the method is robust to different scenarios. This is only partially true. You may get different answers if you use a different background scenario, for example one with strong mitigation or changes in tropospheric aerosols, as that will change the spatial patterns of forcing. Some appropriate caveats would be useful.

AC - Indeed, we have only considered changes in the aerosol forcing scenarios, not the background greenhouse gas forcing scenarios. We will ensure that this is evident from our statements and add caveats where appropriate.

RC1 - Lines 152-153: I think providing more details about the simulations here would be useful. I got a little lost. I suggest moving Table C1 into the main body of the text and expanding it so that it has more information about the specifications of each simulation.

AC - For better overview, we will summarise all simulations in an expanded table in the main.

RC1 - Bullets on page 3: This is essentially pattern scaling. There's a lot of literature you can lean on to show that this is a sensible thing to do.

AC - Agreed. We will add a couple of references here to back up our motivation for using this technique.

RC1 - Line 101: What does "similar" mean?

AC - This is quite a general statement indeed. What we roughly meant is a similar response in surface temperature, precipitation, tropospheric winds and temperature to SAI. We will add this to the manuscript.

RC1 - Line 186: I think you mean monotonically?

AC - Yes, we will fix this.

RC1 - Section 2.2: It would be useful if you said somewhere that the approximations you make are good enough for this purpose, as the point of a feedback algorithm is to correct for such uncertainties. MacMartin et al. and Kravitz et al. both say this if you need citations.

AC - Indeed, the original feedforward/feedback algorithm is constructed in such a way as to be very robust w.r.t. suboptimal parameter choices. We found that even for rather suboptimal choices (e.g., a much too strong feedforward due to model differences in sensitivity), errors in the temperature target were limited (order 0.3 degrees), though noticeable. We will stress this robustness in the revised manuscript.

RC1 - Line 241: Your errors seem kind of high. Looking at Figure 2, an error of 0.4°C is a lot. This likely means your controller isn't tuned as well as it could be, which isn't a big deal, but it would be worth saying so.

AC - The relatively high error is due to the fact that the algorithm in its original form (feedback with

integrator and proportional term) is not constructed to deal with sudden changes, such as starting strong SAI in 2080. Rough "dry-testing" experiments with a simple box model emulator for GMST suggest that this can not be resolved by simple re-tuning. Therefore we have made an ad-hoc improvement by resetting the integrated error term, as described in the manuscript, which reduces but does not eliminate the "cooling overshoot" observed in fig. 2b. We will state this more explicitly in the revised manuscript.

RC1 - Figure 2: I found the panels confusing, in that you're mixing and matching units.

AC - We understand the confusion. To improve the figure we will add separate y-axes in blue for the temperature errors. This will avoid the mixing of units.

RC1 - Line 257: Per the above comment, maybe change "well" to "adequately".

AC - Yes, we will change this.

RC1 - Line 292: This is correct but also a strawman argument. You didn't try to restore the climate completely.

AC - Thank you for pointing this out. We will rephrase this more factually such that we still acknowledge local differences without pretending it to be our goal to restore the climate completely.

RC1 - Figure 5: I'm having trouble making sense of how important these results are. I wonder if you could compute z-scores (or something like that) so I would know whether the CAM minus WACCM differences are large compared to the natural variability of WACCM.

AC - Good point. We will add hatching to the revised figure 5 in the manuscript to indicate significance. Additionally, hatching will be applied to other figures that we think will benefit from this.

RC1 - I did not find the paragraph on lines 376, nor Appendix B, terribly convincing. If you heat the stratosphere by 24°C, you are going to have substantial influences on ozone, and we know that ozone has an influence on surface climate. I would be more comfortable if you simply said that this is what you did, its effects of ozone changes on climate are likely smaller than the effects of the stratospheric heating on climate, and this should be explored further. That puts you on much safer ground.

AC - We initially included this section to provide some more background into the matter, but this indeed poses a risk as the processes involved are more complex than shown here, requiring further analysis. As our article is meant to convey a useful protocol for practitioners, it would indeed be safer to state the things you mention. We will do so in a revised version of the manuscript.

RC1 - Lines 400-402: Kravitz et al. (2014) demonstrated that the controller is likely robust to these sorts of differences. That gives some confidence that your controller indeed can handle this.

AC - That is a good suggestion. We will mention this in the revised text.

RC1 - Lines 427ff: See recent work from the Cornell group, specifically led by Farley or Brody. They're doing the initial steps of what you propose.

AC - We will read their relevant publications and refer to these in the outlook.

RC1 - Code availability: Journals tend to want a fixed repository (e.g., Zenodo) rather than a changeable repository (Github).

AC - We will make a Zenodo repository for the revised version.