Interactive comment on "Complementing CO₂ emission reduction by Geoengineering might strongly enhance future welfare" *by* Koen G. Helwegen et al.

Anonymous Referee #2

Received and published: 23 February 2019

This article makes a very nice start towards seriously integrating solar radiation management (SRM) into an economic model of climate change. A policymaker can choose to reduce CO2 emissions or to inject aerosols into the stratosphere. The former option reduces future CO2 and temperature, whereas the latter only reduces temperature. The key extensions to the benchmark economic model of climate change are a model of the relation between SRM, temperature, and precipitation, a decomposition of climate change damages that distinguishes SRM from CO2 emission reductions, and a dynamic programming formulation (borrowed from other literature) that allows for optimal policy to account for uncertainty about SRM failure and about a type of permanent economic loss that the authors refer to as a tipping point.

It is kind of obvious that the policymaker will choose to use both SRM and emission reductions. The potentially interesting conclusion is the relative mix, which the authors report includes modest SRM. However, this conclusion is likely to be sensitive to hard-to-pin-down assumptions about the cost and damages from SRM (as well as the damages from CO2 and the costs of eliminating it). It is unclear how big of an advance this result is relative to previous literature (Heutel et al). I would prefer to see a more insightful conclusion and will offer suggestions to that effect.

First, a more interesting question about SRM is, to me, not the precise level at which we should be using it now but whether we should start using it at all. The current model has no force that would lead a policymaker to delay introducing SRM. I can think of two relevant forces. First, there may be a fixed cost to beginning SRM, whether political, economic, or ecological. Second, the authors assume (I think) that the risk of SRM failure is constant over time. A more realistic model would have a small chance of new research finding a fatal flaw in SRM as long as it is not used and a much larger chance of actual consequences revealing a fatal flaw once it is used. I am suggesting that the probability of failure should be low as long as SRM is not used and then be higher once it is used. A policymaker might then choose to delay beginning SRM until it is really needed, since SRM may not last for long once it is begun. I strongly encourage the authors to explore these (or other) ways of making the paper about when to begin using SRM.

Second, the interesting aspect of SRM value is, to me, its insurance value. The authors analyze a dynamic programming problem, so they have the machinery to answer that kind of question, but they fail to get there. Instead, they analyze perhaps the least interesting form of "tipping" that they could have. Their tipping point permanently reduces economic output, but the interesting aspect of SRM in a world with tipping points is the potential for SRM to manage a tipping process that is underway. Had the authors chosen to model tipping points in the way that Lemoine and Traeger (2014) did (where a tipping point is a sudden change in a parameter of the physical system,

which manifests itself in temperature and economic output only over time and in a fashion that depends on subsequent policy), SRM could play a key role by allowing the policymaker to intervene after triggering a tipping point so as to mitigate the consequences of that tipping point. Alternately, had the authors studied uncertainty about future warming, SRM could have been used to control the consequences of ending up in a high-warming world. Finally, if the model allowed for the more interesting type of tipping point (or for warming uncertainty) along with a reason to delay beginning SRM, then we may see the policymaker delay SRM until a bad state of the world warrants using it. Learning whether/when that plan is optimal would be interesting.

Smaller comments:

-The authors should spend more space clearly laying out the methodological, calibrational, and conceptual contrasts with Heutel et al. They also should elaborate on the differences in results and speculate as to their origin,

-Technical notes: 1) The authors describe how they mitigate the problem being non-smooth, but it seems far more straightforward and more realistic to simply assume that the chance of tipping is smooth. Assuming no risk of tipping below some temperature is arbitrary. 2) The authors impose an upper bound on injection rates. This is a curious choice. I would prefer the authors to model the source of the constraint. As it is, either the constraint binds or it does not. If it binds, then it is important, the authors should highlight it in their results discussion, and the authors should consider eliminating it. If it does not bind, then it is unimportant and the authors should consider eliminating it (or at worst justify it in terms of numerical convenience but say it is irrelevant in practice).

-Sensitivity checks: I want to see more sensitivity analysis with respect to the fairly arbitrary damage parameters, and especially to the ψ_S that I'm not sure has been explored. This should be the core of the results. I am far more interested in learning about what we need to believe to get near-term SRM to be high or low than in learning about the spread of future policy.

-I wasn't happy with the discussion of precipitation damages in 2.1.3. There are plenty of ways precipitation can matter, independently of temperature. Think of crops. One could have just as easily said that many temperature channels are mediated by precipitation. Plus geoengineering changes precipitation in ways that are not mrely determined by temperature. For instance, might geoengineering change patterns of rainfall (such as monsoons)? Perhaps these kinds of effects could be captured by ψ_S, about which very little is said other than that its level is arbitrary.

-The objective in (10) must be wrong. I think the policymaker must be maximizing expected welfare, not just welfare. And in light of that comment, the motivation for the final paragraph in 2.2 doesn't make a lot of sense. A risk-averse policymaker (as modeled here) has already chosen policy in light of the range of possible cases. It is not internally consistent to postulate caring about percentiles of the subsequent policy performance because of risk aversion.

-In s 3.2, the authors report that abatement decreases after triggering a tipping point. This is purely an artifact of the type of tipping point modeled here. It would probably not arise if abatement had any role in controlling the consequences of a tipping point, as in Lemoine and Traeger 2014.

-Why is the scc unaffected by the possibility of tipping in the abatement+SRM world?

-It seems to me that the final sentence on page 16 would be stronger with a comparison between maximized welfare in the "realistic" world and welfare in the same world if the policymaker used the policy from the abatement-only world. How much could today's policymaker gain by anticipating the future possibility of SRM?

-The discounting result in 3.4 was one of the more interesting results of the paper: a higher discount rate favors SRM because it causes damages, may fail, and doesn't curb future CO2. I would suggest highlighting it more and perhaps doing more with it.

2019.

C5