The paper has improved in this iteration. A few additional comments before I can recommend publication: I'm particually interested in a better explanation for the choice of priors, and its influence on the conclusion.

We thank the reviewer for their careful reading of the manuscript and helpful insight. Below we detail how we have amended the manuscript to address the reviewer's comments.

Major comments
** there are a couple of instances where the authors use variable names that deviate from the standard within the literature. For instance, they call their climate sensitivity 'effective climate sensitivity', a term usually reserved for estimates of climate sensitivity for which lambda is assumed constant (so underselling their work). Furthermore, they abbreviate effective climate sensitivity as ECS, which normally stands for 'equilibrium climate sensitivity'. I urge the authors to replace their use of the term ECS with S, with appropriate subscripts (So S_{20 years}, with S_{140} corresponding to ECS). See https://www.ipcc.ch/sr15/chapter/glossary/, under climate sensitivity.*

We agree that our notation was not in keeping with the literature. We have now amended our notation as suggested, such that 'ECS_{140}' has now become 'S_{140}' and so on for all the timescales.

** the prior on lambda multidecadal seems to have a strong influence on the conclusions. The main text refers to the supplementary information for explanation about the prior given, but I don't really see much information there either. I disagree with the other reviewer (sorry! I know how annoying disagreeing reviewers are..) that a prior symmetrical around zero is physically justified. Evidence from climate models consistently show that changes in lambda on that timescale have an amplifying effect. The authors indicate they want an estimate of S independent of model evidence, but section 4.2 of Sherwood also summarizes observational and process-based argumentation for a decreasing lambda.*
*** I think the previous asymmetry in the prior was fine. If the authors, torn between reviewers, don't want to revert, I'm happy to see the two priors side-by-side.*

We thank the reviewer for their insight. We now present the original prior for multi-decadal climate feedback (symmetrical about zero) alongside an alternative prior that is not symmetrical about zero and derives from Sherwood et al (2020)'s analysis of the pattern effect in section 4.2 therein.

This is discussed in a new sub-section 4.3 (Lines 405 to 435) in a new Figure 8, and in a new column of Table 1.

We also now discuss the implications of these results with an alternative prior in the conclusion (lines 481 to 485 and lines 501 to 505).

*** Please provide a paragraph of explanation of prior choices in the main text*
** the authors use a uniform prior for lambda multidecadal and lambda fast, which is typically not recommended (see the STAN user manual: https://github.com/stan-dev/stan/wiki/Prior-Choice-Recommendations). For consistency, the authors should use the same distribution for lambda Planck, lambda fast and lambda multidecadal (normal distribution seems fine, as the overall feedback is the sum of local feedbacks). There is no physical reason to say that lambda is necessarily between the two given boundaries, and the fact that the posterior does not drop to zero at the boundary indicates the data may indeed by consistent with lambdas < 3.*

We now discuss the choice of priors in the new sub-section 4.3 (Lines 405 to 435). Essentially, we have relatively high confidence in the Planck feedback, and so adopt a normal prior with relatively low uncertainty. However, we adopt a position of ignorance about lambda_fast and lambda_multidecadal and so use a uniform distribution for these in the standard case. The reason lambda_fast and lambda_multidecadal have a lower end of -3 Wm-2K-1 is that we know that the value of each must be greater than minus lambda_{Planck}, explained on lines 411 to 413. This is because the total lambda must always be positive on any timescale.

If we had not imposed this then we could see a situation where lambda_{Planck} + lambda_{fast} was less than zero (and so non-physical) but the model would still work mathematically, provided lambda_{multidecadal} were large enough, due to the finite timestep of the model – event though this would be a non-physical situation in reality because on smaller timesteps than the model resolves total lambda would be negative (and so non-physical in our sign convention). Therefore the hard limit of -3Wm-2K-1 is imposed for both lambda_fast and lambda_multidecadal (where lambda_Planck = +3.3±0.2 Wm-2K-1).

Both lambda_fast and lambda_multidecadal are extended far enough in the positive direction to be past the point where the observation-consistent simulations have already dropped to zero.

Minor comments:
multi-diurnal -> multiday?

Agreed, we have now changed multi-diurnal to multiday (Line 71).

line 380: incomplete sentence

Agreed, we have now completed the sentence which reads (Line 460):

"Note that additional slow feedbacks not considered here, acting from many decades to millennia, may affect how our estimates are comparable to estimates of climate sensitivity from the palaeo-record where any longer feedbacks have been treated as radiative forcing (e.g. Rohling et al., 2012; 2018)."