nature portfolio

Peer Review File

A framework for estimating the anthropogenic part of Antarctica's sea level contribution in a synthetic setting

Open Access This file is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. In the cases where the authors are anonymous, such as is the case for the reports of anonymous peer reviewers, author attribution should be to 'Anonymous Referee' followed by a clear attribution to the source work. The images or other third party material in this file are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/.

Web links to the author's journal account have been redacted from the decision letters as indicated to maintain confidentiality

24th Aug 23

Dear Dr Bradley,

Your manuscript titled "How do we evaluate the contribution of anthropogenic climate change to sea level rise from Antarctica?" has now been seen by 2 reviewers, and we include their comments at the end of this message. They find your work of interest. One of the reviewers rated your manuscript very high, but some important points are raised by the other reviewer. We are interested in the possibility of publishing your study in Communications Earth & Environment, but would like to consider your responses to these concerns and assess a revised manuscript before we make a final decision on publication.

We find the comments from the second reviewer are important to strengthen the robustness of the study results, as well as to make your paper readable for a broader range of readers. We therefore invite you to revise and resubmit your manuscript, along with a point-by-point response that takes into account the points raised. Please highlight all changes in the manuscript text file.

We are committed to providing a fair and constructive peer-review process. Please don't hesitate to contact us if you wish to discuss the revision in more detail.

Please use the following link to submit your revised manuscript, point-by-point response to the referees' comments (which should be in a separate document to any cover letter), a tracked-changes version of the manuscript (as a PDF file) and the completed checklist:
[Link redacted]
** This url links to your confidential home page and associated information about manuscripts you may have submitted or be reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage first **

We hope to receive your revised paper within six weeks; please let us know if you aren't able to submit it within this time so that we can discuss how best to proceed. If we don't hear from you, and the revision process takes significantly longer, we may close your file. In this event, we will still be happy to reconsider your paper at a later date, as long as nothing similar has been accepted for publication at Communications Earth & Environment or published elsewhere in the meantime.

Please do not hesitate to contact us if you have any questions or would like to discuss these revisions further. We look forward to seeing the revised manuscript and thank you for the opportunity to review your work.

Best regards,

Shin Sugiyama, PhD
Editorial Board Member
Communications Earth & Environment
orcid.org/0000-0001-5323-9558

Heike Langenberg, PhD
Chief Editor
Communications Earth & Environment

EDITORIAL POLICIES AND FORMATTING

We ask that you ensure your manuscript complies with our editorial policies. Please ensure that the following formatting requirements are met, and any checklist relevant to your research is completed and uploaded as a Related Manuscript file type with the revised article.

Editorial Policy: Policy requirements (Download the link to your computer as a PDF.)

Furthermore, please align your manuscript with our format requirements, which are summarized on the following checklist:
Communications Earth & Environment formatting checklist

and also in our style and formatting guide Communications Earth & Environment formatting guide .

*** DATA: Communications Earth & Environment endorses the principles of the Enabling FAIR data project (http://www.copdess.org/enabling-fair-data-project/). We ask authors to make the data that support their conclusions available in permanent, publically accessible data repositories. (Please contact the editor if you are unable to make your data available).

All Communications Earth & Environment manuscripts must include a section titled "Data Availability" at the end of the Methods section or main text (if no Methods). More information on this policy, is available at http://www.nature.com/authors/policies/data/data-availability-statements-data-citations.pdf.

In particular, the Data availability statement should include:
- Unique identifiers (such as DOIs and hyperlinks for datasets in public repositories)
- Accession codes where appropriate
- If applicable, a statement regarding data available with restrictions
- If a dataset has a Digital Object Identifier (DOI) as its unique identifier, we strongly encourage including this in the Reference list and citing the dataset in the Data Availability Statement.

DATA SOURCES: All new data associated with the paper should be placed in a persistent repository where they can be freely and enduringly accessed. We recommend submitting the data to discipline-specific, community-recognized repositories, where possible and a list of recommended repositories is provided at http://www.nature.com/sdata/policies/repositories.

If a community resource is unavailable, data can be submitted to generalist repositories such as figshare or Dryad Digital Repository. Please provide a unique identifier for the data (for example a DOI or a permanent URL) in the data availability statement, if possible. If the repository does not provide identifiers, we encourage authors to supply the search terms that will return the data. For data that have been obtained from publically available sources, please provide a URL and the specific data product name in the data availability statement. Data with a DOI should be further cited in the methods reference section.

Please refer to our data policies at http://www.nature.com/authors/policies/availability.html.

REVIEWER COMMENTS:

Reviewer #1 (Remarks to the Author):

In this paper, the authors explore how anthropogenic climate change and internal variability contribute

to sea-level rise from the West Antarctic ice sheet. By using a synthetic marine ice-sheet example and a Bayesian approach, the study shows how their relative contributions can be assessed, taking into account model uncertainties and climatic variations. Climate change is found to increase median sea-level rise and the weight in the tails of distributions. This research enables attribution of sea-level rise for real-world problems.

I find this study to be a great addition to the current knowledge on attribution of sea-level rise, a problem with significant societal ramifications. With the caveat that I am not an expert in statistical methods, my assessment is that the paper is of high quality and very well written. Therefore, I recommend publication essentially as it is. My only minor point is that, if the style of the journal allows it, DOIs should be added to the references for the readers' convenience.

Reviewer #2 (Remarks to the Author):

Dear Authors,

the manuscript aims to provide a framework for determining the anthropogenic component in sea level rise contributions from Antarctica. I do think that the importance of this kind of attribution studies can hardly be overstated. With recent advances in the simulation of ice sheets, in particular regarding more and larger ensemble simulations, this manuscript seems very timely and has the potential to shape future high profile studies. It can further contribute to the crucial discourse of how to address uncertainties in ice sheet predictions. I agree with the authors that there is no established framework to do so and that this creates another layer of uncertainty (which is sometimes refereed to as ‘deep uncertainty’). The main novelty lies in the combination of a three dimensional glacial flow model with stochastic forcing, a Bayesian calibration and the subsequent analysis of the anthropogenic enhancement ratio. This is a big step towards answering one of the most crucial question of the field.

However, I believe the authors could do a better job in justifying and testing the proposed framework and the steps it involves. In my option the advantages of the proposed framework have to be made more clear. Where possible, this should be done analytically, for example by a ‘perfect model test’ and sensitivity tests (see below). In other cases additional discussions would help others to follow, apply and develop the proposed framework.

I hope my comments help the authors to revise the manuscript in order to make it easier for other scientists to follow or amend the proposed framework/methodological choices. This will require to add more focus on the justification and discussion of methodological decisions. This applies in particular to the topics of the calibration parameter, the prior probability distribution and the model parameters perturbed. In this context, the motivation of each step of the analysis should be made more clear. For example: You propose a Bayesian calibration while Christian at al. (2022) does not. Maybe you can expand on the reasons for this decision.

For these reasons I would encourage the authors to revise their manuscript before a final decision is reached.

Main comments:
Calibration: The calibration is performed on the average ocean melt at several time slices. The reference (‘observation’) is a coupled ice-ocean model which is run on the ice geometry of the specific geometry from the ice-only model at that time. I expect this calibration to be capable of judging the quality of the melt rate parametrization (including the melt factor) but I do not expect it to be able to calibrate other components of the ice flow model. Many other model parameters, when varied in future

studies, might well have no (or only marginal) impact on the difference between coupled and parameterized melt rates. Similarly, I expect this calibration to be insensitive to past fluctuations of the forcing. In my understand this procedure looks like a hybrid between a parameterized and coupled ocean forcing. If correct it should be described as such. The use of reference data from a coupled model including future time slices (e.g. 2100) illustrates that this dataset cannot be replaced by real word observations, as would be desirable for a classic probabilistic calibration.

At the same time the prior distribution is centered at M=1.25, which is justified by artificial melt observations of 23m/year at the end of the second initialization phase. The prior is therefore (in my view) taking the form of a typical Bayesian update (i.e. the calibration) instead of describing the expectations prior to including observational targets. The description of the origin of the prior is in parts hidden in the caption of supplementary Figure 7.

The width of the prior (sigma_P) is 0.2 in M (unit-less). This is described as wide (line 349), however it is not explained how this value has been chosen, nor how one might find reasonable values for this in general. Looking at figure 7b, a change of M of 0.2 seems to correspond to less than one m/year in melt rate. This seems in fact to me to be a quite narrow definition of the prior since I would assume that observation of the mean basal melt rate of an ice shelf would have larger uncertainties than 1 m/year (but I am not an expert on this kind of observations). This prior effectively rules out values of M<0.75 (compare Figure 8b and 8f) which should discussed.

The framework should include some tests on how big the influence of methodological decisions is on the final results. This could include a perfect model test and sensitivities to selected values (such as sigma_P) and ideally which/how many parameters are selected. How would you expect the anthropogenic enhancement ratio to change if more model parameters would be included in the parametric uncertainty?

Generally, I think this manuscript would benefit from a sort overview of further challenges the authors see for the attribution of the anthropogenic component of Antarctic sea level rise.

Minor comments:

The forcing is solely based on oceanic temperature variations. In the coming decades an increase in surface accumulation is discussed for parts of Antarctica. I would have liked a mention of this topic in the manuscript.

Line 15: you highlight the importance of structural uncertainty but neglect it in the framework. Can you expand in the manuscript why previous approaches to structural uncertainty cannot be recommended?

Can you provide advice on how to choose model parameters to be included in parametric uncertainty estimates? How did you decide on choosing M?

Chapter titles are sometimes misleading:
Line 86: this chapter is less about ‘distributions of sea level rise’ than about the impact and interactions of the two types of uncertainty.
Line 364: What end members?

Line 93 “no study has yet combined parametric and aleatory uncertainty”: technically I think studies using Monte-Carlo methods, like Wernecke et al. (2022), do this. But I am not aware of any studies

combining parametric and aleatory forcing uncertainties (which is what I believe you mean). I am not trying to promote my paper (no need to add the citation), just an example I know of.

107-108: As mentioned, using model results for 2100 can hardly be considered equivalent to "observations".

Looking at the ensemble means in Figure 2g: The counterfactual seems to have a small positive trend, which I assume is by chance and is absolutely fine. However, please check whether the dashed lines show the actual trend of the ensemble mean (which I would expect to see here, and is the way I understand the caption) or the theoretical (by-design) trends of 0 and 100 m/century are plotted.

Line 122 'trends': do you mean scenarios?

I am a bit concerned, that the model parameter and forcing are that closely connected (equation 10), since, as you say, the model parameter is basically a scaling parameter for the melt forcing. Would you expect the results be different if using other parameters and forcing which act on the model system in ways less related to each other?

Line 166-167 "As a result, a small change in the target calibration value would result in a small change in the selected value of M" I disagree. A reduction of <10% in calibration value to 21m/year would reduce mu_M from 1,25 to 0.5 i.e. more than halve the melt at a given depth. So a small change in calibration value (remember we are talking about the prior, so 'calibration value' is a bit odd choice of words) would result in a much larger change in M, correct?

Line 205-207 "(more specifically, with a linear anthropogenic trend, the melt anomaly scales with the square of time, which rapidly outweighs any time-integrated negative internal component)"
If I understand correctly, this is an interesting point. Since the accumulated melt seems to be important (until it reaches some threshold/tipping point) the system more vulnerable against long lasting trends then against short terms variability in melt. Maybe make your point more prominently.
Line 245-246 "(there are periods when relatively more retreats are initiated than the background trend, see figure 5)". Unclear to me.
Line 275 "'real-world' simulations" rephrase (there is no such thing)
line 304 " natural" rephrase? Maybe 'internal' (Humans are natural)
line 361 "sigma_m" do you mean sigma_L?
Line 363 sigma_L=10: What is this value based on? Where could it be based on in general?
Caption Figure 7: (a) and (b) "calibration phase" do you mean initialization phase?
Figure 8/general: In which stage of the calibration do interpolate the probabilities and/or SLR values (if at any stage) and in which way? If there is no interpolation each forcing realization would result in a frequency distribution with only 5 weighted members, correct?

Thank you for your work and the consideration of my comments,
Andreas Wernecke

References:

Wernecke, Andreas, et al. "Quantifying the impact of bedrock topography uncertainty in Pine Island Glacier projections for this century." Geophysical Research Letters 49.6 (2022): e2021GL096589.

We thank the two referees for their time in this review process and the comments they have provided, which we feel have improved the quality and readability of our manuscript. Below, we respond to these comments in detail.

Reviewer #1 (Remarks to the Author):

In this paper, the authors explore how anthropogenic climate change and internal variability contribute to sea-level rise from the West Antarctic ice sheet. By using a synthetic marine ice-sheet example and a Bayesian approach, the study shows how their relative contributions can be assessed, taking into account model uncertainties and climatic variations. Climate change is found to increase median sea-level rise and the weight in the tails of distributions. This research enables attribution of sea-level rise for real-world problems.

I find this study to be a great addition to the current knowledge on attribution of sea-level rise, a problem with significant societal ramifications. With the caveat that I am not an expert in statistical methods, my assessment is that the paper is of high quality and very well written. Therefore, I recommend publication essentially as it is. My only minor point is that, if the style of the journal allows it, DOIs should be added to the references for the readers' convenience.

We are extremely pleased to read the referee's positive assessment. We have added DOIs to the references in the updated manuscript.

Reviewer #2 (Remarks to the Author):

Dear Authors,

the manuscript aims to provide a framework for determining the anthropogenic component in sea level rise contributions from Antarctica. I do think that the importance of this kind of attribution studies can hardly be overstated. With recent advances in the simulation of ice sheets, in particular regarding more and larger ensemble simulations, this manuscript seems very timely and has the potential to shape future high profile studies. It can further contribute to the crucial discourse of how to address uncertainties in ice sheet predictions. I agree with the authors that there is no established framework to do so and that this creates another layer of uncertainty (which is sometimes refereed to as ‘deep uncertainty’). The main novelty lies in the combination of a three dimensional glacial flow model with stochastic forcing, a Bayesian calibration and the subsequent analysis of the anthropogenic enhancement ratio. This is a big step towards answering one of the most crucial question of the field.

However, I believe the authors could do a better job in justifying and testing the proposed framework and the steps it involves. In my option the advantages of the proposed framework have to be made more clear. Where possible, this should be done analytically, for example by a 'perfect model test' and sensitivity tests (see below). In other cases additional discussions would help others to follow, apply and develop the proposed framework.

I hope my comments help the authors to revise the manuscript in order to make it easier for other scientists to follow or amend the proposed framework/methodological choices. This will require to add more focus on the justification and discussion of methodological decisions. This applies in particular to the topics of the calibration parameter, the prior probability distribution and the model parameters perturbed. In this context, the motivation of each step of the analysis should be made more clear. For example: You propose a Bayesian calibration while Christian at al. (2022) does not. Maybe you can expand on the reasons for this decision.

For these reasons I would encourage the authors to revise their manuscript before a final decision is reached.

We are grateful to the referee for their time in reading our manuscript and providing detailed comments. We are pleased to read that the referee shares our enthusiasm for the work and the crucial questions which motivate it. Below, we respond to the referee's comments in detail. In revising our manuscript in response to these comments, we have improved the justification and discussion of methodological choices, particularly regarding the Bayesian calibration; we believe that these changes improve the manuscript, not only in strengthening the robustness of the results, but also to make it more readable for a broad range of readers. (Note that we have adopted a different latex template for the second draft to allow a track changes version and thus line numbers are sometimes quite different in the updated ms, and we mention line numbers where we have made changes.)

We agree that the benefit of our Bayesian approach, over the frequentist approach of Christian et al. was not made sufficiently clear. In the updated manuscript, we have elaborated on this point in the introduction, writing [lines 121-133]:

"Here, we present a framework for attributing SLR contributions from the WAIS which uses a Bayesian approach integrating multiple realizations of forcing; we build upon Christian et al. in two main ways: firstly, we consider SLR contributions, rather than retreat, as the metric to be attributed. By using SLR as the attribution metric, we are able to quantify the role of anthropogenic climate change for observed SLR within any interval, rather than only exceedance of a single, predefined retreat threshold. This alleviates the common 'event definition' problem which commonly impacts

attribution studies. Secondly, we explicitly account for the role of variable model parameters in the attribution assessment. Bayesian approaches naturally permit the joint probability density of multiple model parameters, which may be poorly constrained in general, to be represented within a projection of SLR. This avoids the need to specify the precise values of model parameters at the outset, which yield very different attribution results depending on the particular choice of parameters in the framework of Christian et al."

We have also added a note of this in the discussion/conclusions sections, where we state [lines 497-498] that *"Our use of a Bayesian framework allows us to treat parametric uncertainty within attribution assessments and avoids the need to specify a single event to be detected."*

Main comments:

Calibration: The calibration is performed on the average ocean melt at several time slices. The reference ('observation') is a coupled ice-ocean model which is run on the ice geometry of the specific geometry from the ice-only model at that time. I expect this calibration to be capable of judging the quality of the melt rate parametrization (including the melt factor) but I do not expect it to be able to calibrate other components of the ice flow model. Many other model parameters, when varied in future studies, might well have no (or only marginal) impact on the difference between coupled and parameterized melt rates. Similarly, I expect this calibration to be insensitive to past fluctuations of the forcing. In my understand this procedure looks like a hybrid between a parameterized and coupled ocean forcing. If correct it should be described as such. The use of reference data from a coupled model including future time slices (e.g. 2100) illustrates that this dataset cannot be replaced by real word observations, as would be desirable for a classic probabilistic calibration.

The referee has painted a clear picture of the procedure and is absolutely correct that the calibration is only really capable of judging the quality of the melt rate parameterization (there is a flow calibration component, since the flow affects the ice shelf cavity geometry, which in turn affects the melt rates, but this expected to be small and we therefore don't mention it). In the updated ms, we have made it clearer that we only calibrate the melt aspect of the system and focus on this part as it primarily controls the interactions between aleatory and parametric uncertainties (discussed further below). In addition, calibrating other parts of the ice sheet model already has a well established procedure; calibrating with an ocean model for precise melt rate calibration is the novel aspect. In the updated ms, we write [lines 187-193]:

"The melt rate calibration procedure is only capable of calibrating the melting aspects of the flow model; other parameters, such as those related to basal sliding and ice viscosity, which are important

in determining ice flow (and thus SLR) are not calibrated. Other studies (e.g Nias et al. 2019, Bevan et al. 2023) have established procedures for calibrating many such aspects of ice-sheet models using observational data; the novelty of our calibration methods is that it permits precise calibration of basal melt rates, which have previously only been indirectly calibrated via the effect of melting on ice flow."

The referee is correct to describe the procedure as a 'hybrid' between parametrizations and coupled ice-ocean models. We this was only mentioned implicitly in the original ms; in the updated ms, have stated this more clearly [lines 174-181]:

"Our approach can be considered a `hybrid' between a parametrisation of melting and a coupled ice-ocean model: we use a parameterisation of basal melting for computational efficiency and adopt a Bayesian approach to the model parameters within: simulations are weighted by comparing their predictions of basal melt rates with those from an offline ocean model at different snapshot times through a simulation (methods); the ocean model thus plays a role analogous to a ground-truth in a traditional Bayesian update, i.e. it is the information assimilated into the model."

At the same time the prior distribution is centered at M=1.25, which is justified by artificial melt observations of 23m/year at the end of the second initialization phase. The prior is therefore (in my view) taking the form of a typical Bayesian update (i.e. the calibration) instead of describing the expectations prior to including observational targets. The description of the origin of the prior is in parts hidden in the caption of supplementary Figure 7.

We agree that this is a key unconventional aspect of our procedure: in our framework, we assume that the mean melt rate of 23m/yr is not treated as an observation in the conventional Bayesian sense, but rather an arbitrary piece of prior information, which informs the prior only. In the updated ms, we refer to the mean melt rate of 23m/yr as 'prior information' to make this distinction clear. The value of 23m/yr is somewhat arbitrary and was simply selected to illustrate the procedure. In our procedure, new information is assimilated via the mitgcm simulations, which are akin to 'observations' in the traditional Bayesian sense, and are treated as independent of the prior information.

This is clearly a slight philosophical shift from the traditional Bayesian update technique, in which satellite observations, rather than model outputs, are used to inform the Bayesian update. In the updated ms, we write [line 181-184]:

"It should be noted that this is a slightly different philosophy to a typical Bayesian update in ice sheet modelling, in which agreement with satellite observations, rather than with results of more detailed models, are typically used to update probabilities."

As we outline in more detail below, we expect that, in practical applications, satellite observations must also be assimilated and say this in the updated ms. We have also moved the description of the origin of the prior from the caption of supplementary figure 7 to the methods section, where we write [line 662-665]:

"In the results contained herein, we use \mu = 1.25, based on agreement in the mean melt rate after the initialization stage (in this case, a mean melt rate of 23m/yr, which can be thought of as an arbitrary piece of prior information)"

The width of the prior (sigma_P) is 0.2 in M (unit-less). This is described as wide (line 349), however it is not explained how this value has been chosen, nor how one might find reasonable values for this in general. Looking at figure 7b, a change of M of 0.2 seems to correspond to less than one m/year in melt rate. This seems in fact to me to be a quite narrow definition of the prior since I would assume that observation of the mean basal melt rate of an ice shelf would have larger uncertainties than 1 m/year (but I am not an expert on this kind of observations). This prior effectively rules out values of M<0.75 (compare Figure 8b and 8f) which should discussed.

The width of the prior was described as 'wide' to reflect the fact that the range of values of M is restricted to 0.5-1.5, with a standard deviation of 0.2 reflecting a reasonable proportion of this, but the referee is correct that this does indeed effectively rule out M < 0.75, and thus describing it as wide is inaccurate. In the updated ms, we have removed the description of 'wide' [line 665], preferring to stay agnostic in this description.

To help to constrain the effect of the prior standard deviation, and demonstrate that our framework is robust to this choice, we have added a new figure, in which we plot the anthropogenic enhancement ratio (AER) as a function of SLR and time (i.e. as in figure 4) for different values of the parameters sigma_P and sigma_L. At the largest value of sigma_P shown here (sigma_P = 0.5, for which the prior is essentially agnostic to the value of M), the same picture essentially emerges (although it is complicated by the presence of an area where the AER < 1 just inside the tail, but this is a finite size effect). Essentially, with a wider prior, low M values are given higher weight than with a narrower prior. Lower M values are associated with a less distinguishable anthropogenic signal, because a larger forcing anomaly is required to achieve the same integrated melt anomaly and thus the proportion of simulations which have been triggered to retreat shows less difference between the natural and anthropogenic ensembles.

In the updated ms, we have noted that choices of the parameters in the error models can influence the attribution assessment, stating that [lines 540-543]

"...the effect of parameters which control the strength of Bayesian updates must be explored; although we find that varying these parameters within reasonable ranges does not qualitatively change the results (methods), they may influence the precise values of anthropogenic enhancement. "

We also expound upon this in the methods section, where we describe the Bayesian update procedure (we choose not to go into this level of detail in the main text to avoid confusing the reader and obfuscating the message) [lines 702-726]:

"Supplementary figure 10c shows the AER as a function of SLR and time (i.e. as in figure 4 of the main text) for different values of the prior parameter sigma_P and melt error covariance \sigma_L within reasonable ranges. We see that, while varying these parameters adjusts the precise value of the AER, the overall picture -- that higher observed SLR are concomitant with stronger anthropogenic influence -- remains. The small exception to this is for large \sigma_P and small \sigma_L, for which the anthropogenic signal is most obscured (see below) and a band of AER < 1 emerges close to the tail. This is a finite size effect, and would disappear in the limit of a large number of simulations, emphasizing the need for large ensembles of simulations.

For smaller \sigma_L, errors in melting are penalized more harshly; in this study, smaller \sigma_L tends to shift weight towards smaller M, which typically display smaller errors in melting (see supplementary figure 8a--b for an example from one realization of forcing). Simulations using a smaller value of M require a larger time-integrated forcing anomaly to achieve the same integrated melt anomaly required to initiate retreat. Simulations in which this is achieved in the anthropogenic case and not in the counterfactual case, tend, therefore, to be observed later on average, when the ensemble mean difference in forcing is greater. Thus, for a given time, the ratio of ensemble members which have retreated in the anthropogenic ensemble to those which have retreated in the counterfactual ensemble is closer to unity for smaller M, leading to a dampened anthropogenic effect. Conversely, smaller \sigma_P shifts weight towards M = \mu = 1.25 (in this case), which is at the higher end of the M range considered here, enhancing the anthropogenic effect."

The framework should include some tests on how big the influence of methodological decisions is on the final results. This could include a perfect model test and sensitivities to selected values (such as sigma_P) and ideally which/how many parameters are selected. How would you expect the

anthropogenic enhancement ratio to change if more model parameters would be included in the parametric uncertainty?

As mentioned above, we have included a new figure, which shows the AER as a function of time and SLR (i.e. as in figure 4 of the ms) for different values of sigma_P and sigma_L. The salient point is that although the absolute magnitude of the AER may change, the picture of increasing AER with SLR remains for values in a feasible range.

We restrict our sensitivity analysis to these two parameters, and leave the assumption that model parameters which are in practice uncertain (e.g. basal sliding and ice viscosity) are known, to ensure that the space of things we vary in our sensitivity tests remains reasonable. Varying the basal sliding and/or viscosity coefficients requires the running of a new ensemble of simulations and is computationally difficult, particularly if many different parameters are to be tested. We hope the referee will agree that this is beyond the scope of our work; as we proceed with real world attribution studies, this will certainly be a key factor, however. In addition, our use of a detailed ocean model to calibrate melt rates is the novel factor of our study, while there is already established literature on ice sheet model parameters; while these parameters are clearly important, we have chosen to focus on the novel aspect of our work.

That said, we have added several paragraphs in the main text providing an overview of further challenges and questions of note in the WAIS attribution question (see next referee comment). There, we describe how other parameters, including the basal sliding coefficient and ice viscosity are in practice poorly constrained, and recommend the use of Bayesian methods to constrain them. Of course, this has computational challenges in terms of the sheer number of computations required, and that is one thing we hope to highlight with this paper. It is not clear to us how the AER would change in general if more model parameters were included in the parametric uncertainty, and we therefore refrain from commenting on this in the updated ms. For example, if Bayesian calibration meant that simulations with a higher basal friction coefficient were preferentially selected, this would be expected to decrease the AER, since the configuration would be 'more stable' – a larger time-integrated melt anomaly would be required to destabilize it – as discussed for lower M values above.

Regarding a perfect model experiment: for a study in which only variations in model parameters are considered, a perfect model experiment would take the form of running the ice sheet model with the known parameters then using the results of this experiment as data, which constrains the ensemble. The distributions of SLR from the resulting ensemble would then be compared to the single SLR from the perfect model. However, a direct comparison between data (i.e. melt rates in our case) from the perfect model and other simulations is not possible because the cavity

geometry is variable. Cavity geometry will have a big influence: for example, a configuration grounded at the ridge will naturally have a much lower melt rate than one which has retreated off the ridge, because the latter is exposed to much warmer ocean water. In lieu of a sensible perfect model experiment, therefore opt to present a sensitivity analysis, as outlined above, demonstrating our results are robust under different parameter values.

Generally, I think this manuscript would benefit from a sort overview of further challenges the authors see for the attribution of the anthropogenic component of Antarctic sea level rise.

In the updated ms, we have expanded substantially upon the brief description of further challenges that was contained in the original ms [lines 503–568]. We discuss what we see as the three key challenges, devoting a paragraph in the discussion section to each:

- Computational challenges, including assimilation of large numbers of parameters to be varied and calibrated (i.e. a full treatment requires the inclusion of uncertainties in ice-internal processes, such as the sliding law, ice rheology, damage evolution, calving etc), as well as determining reasonable values of parameters used in the Bayesian updates
- State uncertainties, including knowledge of initial states, particularly where the ice is grounded and where the ice front is located
- Determining trends in climatic forcing, particularly including trends in both accumulation and ocean forcing simultaneously, and the compound attribution question of what is driving trends in these quantities.

It is worth noting that this is just a brief overview of what really requires a dedicated review paper of its own!

<u>Minor comments:</u>

The forcing is solely based on oceanic temperature variations. In the coming decades an increase in surface accumulation is discussed for parts of Antarctica. I would have liked a mention of this topic in the manuscript.

The referee is correct that variability in accumulation should have been mentioned. In the updated ms, we mention this when we describe the forcing. We also mention that we focus on the ocean variability for several reasons: firstly, for simplicity (the attribution problem becomes more difficult when different forcings have different trends and attributions of their own, and we note this in the new paragraph discussing challenges), secondly, because this is thought to be the key driver of WAIS retreat, and thirdly, variability in accumulation is much smaller than variability in ocean forcing.

In the updated ms, we write [lines 216-221]:
"Although accumulation rates also feature significant interdecadal internal variability, and are projected to display an anthropogenic trend in the future, this variability is smaller than in the ocean forcing for West Antarctica. In addition, changes in melting, rather than accumulation, are understood to have been the dominant driver of recent WAIS retreat, and having multiple forcings, each with a unique anthropogenic trend complicates the attribution task."

We also note in the updated discussion section that [lines 565-568]
"Additionally, anthropogenic trends in accumulation, not considered in this study, must be considered simultaneously with trends in ocean warming; in the future, trends in accumulation are expected to partly offset ice loss from WAIS, potentially obscuring the anthropogenic signal."

Line 15: you highlight the importance of structural uncertainty but neglect it in the framework. Can you expand in the manuscript why previous approaches to structural uncertainty cannot be recommended?

We have removed the word "structural" from the updated ms, which was incorrectly used here. We simply intended to say that there are uncertainties in ice sheet models. We are sorry for this confusion. In the updated ms, we now mention structural uncertainty – and the challenges surrounding it – in the new paragraph on challenges in future attribution studies. The most prominent challenge is that assessing such uncertainties requires multiple different ice sheet models to be used, running the same experiments, which is a significant community effort (as well as computationally expensive, given the large ensembles required). In the updated ms, we state [line 543-546]:

"…multiple different ice sheet models should be used in order to assess structural uncertainties arising from those processes not represented in some ice sheet models, further adding to the computational challenge."

Can you provide advice on how to choose model parameters to be included in parametric uncertainty estimates? How did you decide on choosing M?

We are not sure whether the referee means "which values of the parameter to sample" or "which parameters within the model to vary", so we respond to both.

For the first: we note in the methods that our constraints that the ice sheet remain grounded at the ridge crest under cold forcing, and retreat under warm forcing, means that we have to restrict to

0.5 < M < 1.5 (this is essentially the same as choosing finite support for the prior), stating that [lines 870-878]:

"Note that for a consistent estimate of SLR contributions from simulations with different values of M, we require similar initial conditions, chosen to be a grounding line at or near the seabed ridge crest. For M >~ 1.5, the ice retreats irreversibly down the ridge during the second initialization stage. We therefore consider only M values smaller than this. In addition, we should impose that a constant warm forcing applied to the shelf should initiate retreat (WAIS retreat was, in practice, hypothesised to be initiated with forcing oscillating between warm and cold); we found that for M <~ 0.5, no ice sheet retreat was initiated under warm forcing. Therefore, we restrict ourselves to the range 0.5 <= M <= 1.5."

For the second: as we note above, we have chosen to focus on varying the parameter in the melt rate parametrization in order to elucidate the relationship between aleatory and parametric uncertainty and to focus on the novel aspect of our work: directly calibrating melt rates via the use of a detailed ocean model. In the updated ms, we have stressed the need to vary other parameters in 'real-world' assessments, stating that [lines 530–537]

"In considering a generic marine ice sheet, we are able to neglect uncertainty arising from model parameters governing basal sliding and ice viscosity, as well as processes such as damage, calving, sliding law, and ice rheology, which might obscure (or amplify) long-term climatic trends in the forcing, but should be included in assessments of SLR and thus its attribution to anthropogenic climate change. Additional parametric uncertainties can be succinctly integrated into the Bayesian approach taken here, and should be calibrated with observations."

Chapter titles are sometimes misleading:
Line 86: this chapter is less about 'distributions of sea level rise' than about the impact and interactions of the two types of uncertainty.

We have renamed this section to [line 149-150]: *"Interactions between aleatory and parametric uncertainties in sea level rise projections"* to reflect the referee's correct comment, and thank the referee for this helpful suggestion.

Line 364: What end members?

This title is a relic of a previous draft and should have been removed. We apologize for this and thank the referee for spotting. In the updated ms, this section is titled [line 729] *"Details of ice sheet configuration"*

Line 93 "no study has yet combined parametric and aleatory uncertainty": technically I think studies using Monte-Carlo methods, like Wernecke et al. (2022), do this. But I am not aware of any studies combining parametric and aleatory forcing uncertainties (which is what I believe you mean). I am not trying to promote my paper (no need to add the citation), just an example I know of.

This is a fair point, and we have adjusted our wording to make this statement clearer. In the updated ms, we write [line 159-161]:

"Although accounting for parametric uncertainty in this way is now standard, no study has yet probed the interaction between parametric and aleatory forcing uncertainties, primarily because of the computational expense of doing so"

107-108: As mentioned, using model results for 2100 can hardly be considered equivalent to "observations".

We hope that our response to the previous comment has clarified this, but we have also updated this statement to reflect, adding that the ocean model information is analogous to observations in a typical Bayesian update (i.e. it is the new information that is assimilated). We write [line 179-181]: *"the ocean model thus plays a role analogous to a ground-truth in a traditional Bayesian update, i.e. it is the information assimilated into the model."*

Looking at the ensemble means in Figure 2g: The counterfactual seems to have a small positive trend, which I assume is by chance and is absolutely fine. However, please check whether the dashed lines show the actual trend of the ensemble mean (which I would expect to see here, and is the way I understand the caption) or the theoretical (by-design) trends of 0 and 100 m/century are plotted.

This is a good point, and we thank the referee for spotting this error: what is plotted is rather the externally imposed trend in individual simulations, i.e. the T(F) in equation 13, which is 100m/century shallowing of the pycnocline in the anthropogenic case and 0m/century shallowing of the pycnocline in the counterfactual case. This was incorrectly labelled in the caption of the original ms. We feel it's more instructive to leave the externally imposed trend plotted (i.e. as was in the original ms), and have updated the caption to reflect this *("dashed lines indicate the externally imposed trend (the T(F) term in equation 13), i.e 100m/century and 0m/century shallowing of the pycnocline in the anthropogenic and counterfactual cases, respectively.").* This is to allow the reader to clearly see the trends in forcing in both cases. (Of course, in the limit of large number of ensemble members, these would be identical.)

Line 122 ‘trends’: do you mean scenarios?

Here we refer to the trend itself (i.e. 0m/century or 100m/century), which is indicative of the scenario (i.e. counterfactual or anthropogenic). We agree that it’s perhaps clearer to refer to the scenario here and have used that in the updated ms.

I am a bit concerned, that the model parameter and forcing are that closely connected (equation 10), since, as you say, the model parameter is basically a scaling parameter for the melt forcing. Would you expect the results be different if using other parameters and forcing which act on the model system in ways less related to each other?

This is a good point. It’s precisely because the model parameter and forcing are closely related that we have strong interactions between parametric and aleatory uncertainty. As we note above, we expect that variations in other parameters wouldn't change the key findings and give an example of basal sliding. However, we wouldn’t expect the interactions between aleatory and parametric uncertainty to be so strong in this case.

Line 166-167 “As a result, a small change in the target calibration value would result in a small change in the selected value of M” I disagree. A reduction of <10% in calibration value to 21m/year would reduce mu_M from 1,25 to 0.5 i.e. more than halve the melt at a given depth. So a small change in calibration value (remember we are talking about the prior, so ‘calibration value’ is a bit odd choice of words) would result in a much larger change in M, correct?

This is a good point. This comment made us realize that there are really two cases: (1) the target prior information (here, the 23m/yr mean melt rate) is not sensitive to parameters and (2) the target prior information is sensitive to parameters. However, because of the extreme sensitivity of SLR to M, these two are basically the same – both a small and a large change in M (resulting from a small change in the observation) may result in a large change in SLR. The situation we described in the original ms actually refers to case (1), with mean melt rate not sensitive to M (as the referee correctly points out, albeit phrased the other way round), and SLR sensitive to M, so the final SLR is sensitive to the observation value, but we had described the system as if it were case (2). We feel that this makes our statement somewhat stronger, since it is not just restricted to cases (like that we have here) where the parameter value is sensitive to the observation. We have clarified this in the updated ms, stating [lines 309-324]:

“ In the example presented here, the mean melt rate at the start of the simulation (at the end of the initialization stage, which is performed separately for different values of M) is only weakly dependent

on the melt rate parameter M (supplementary figure 7d), owing to a feedback between melting and ice geometry (methods). As a result, a small change in the single target metric to be matched would result in a large change in the selected value of M (supplementary figure 7d), which would ultimately result in a large change in the simulated SLR at the end of the simulation (figure 2h). In other cases where the target observation is more sensitive to parameters, a small change in the observation would be expected to result in a small change in the selected parameter, but this may also ultimately result in a large change in the simulated SLR at the end of the simulation, owing to the non-linearity of SLR in M."

The referee is correct that 'calibration value' is confusing, given that we are discussing the prior. We have changed this to 'single target observation to be matched' to emphasize that we are comparing to studies that use only a single 'observation' to calibrate parameters.

Line 205-207 "(more specifically, with a linear anthropogenic trend, the melt anomaly scales with the square of time, which rapidly outweighs any time-integrated negative internal component)" If I understand correctly, this is an interesting point. Since the accumulated melt seems to be important (until it reaches some threshold/tipping point) the system more vulnerable against long lasting trends then against short terms variability in melt. Maybe make your point more prominently.

This interpretation is absolutely correct, and we appreciate the referee's succinct wording, which explains the situation nicely. In the updated ms, we have removed this sentence from parenthesis to make it more prominent, and added a few words to make the point clearer, stating [lines 376-379]:

"More specifically, with a linear anthropogenic trend, the melt anomaly scales with the square of time, which rapidly outweighs any time-integrated negative internal component: the system is more vulnerable to long-lasting trends in melting than to short term variability."

Line 245-246 "(there are periods when relatively more retreats are initiated than the background trend, see figure 5)". Unclear to me.

In the updated ms, we have clarified this, explaining that [lines 453-457]:

"in the limit of infinite ensemble members, the proportion of retreats initiated would be smooth, whereas in our finite size ensemble, the proportion of retreats initiated oscillates around the trend in

this quantity, with periods when relatively more, and periods when relatively few, retreats are initiated compared to the background trend"

In our finite size ensemble, that is not the case, which manifests as bands of relatively high and low values in the AER corresponding to period of relatively more and less retreats, respectively, being initiated than the background trend. We don't wish to dwell on this point, as it is really just a finite size effect, and the comment only serves to head off a question about the patterning in figure 4.

Line 275 "'real-world' simulations" rephrase (there is no such thing)

We have replaced this with [lines 504-505] *"using geometries and parameters representing real world conditions"* in the updated ms.

line 304 "natural" rephrase? Maybe 'internal' (Humans are natural)

We have rephrased this as [line 586] *"with and without an anthropogenic trend"*

line 361 "sigma_m" do you mean sigma_L?

Yes, thank-you for spotting this typo. It has been corrected in the updated ms [line 694].

Line 363 sigma_L=10: What is this value based on? Where could it be based on in general?

We have used sigma_L = 10 based on the fact that errors in the melt model are on the order of 10s of meters per year. In general, the error covariance should be on the same order as the errors; in the updated ms, we write [line 697-700]:

"In the results presented here, we use sigma_L = 10m/year. In general, the sigma_L should be on the same order of magnitude as errors in melting; in our simulations, melt errors are typically on the order of 10s of meters per year"

In addition, our new figure demonstrating how the anthropogenic enhancement ratio varies with sigma_P and sigma_L indicates that, with sigma_L on the order of 10s of metres, variations so not make a qualitative difference to the results.

Caption Figure 7: (a) and (b) "calibration phase" do you mean initialization phase?

Yes, thanks for spotting. We've corrected this in the updated ms.

Figure 8/general: In which stage of the calibration do interpolate the probabilities and/or SLR values (if at any stage) and in which way? If there is no interpolation each forcing realization would result in a frequency distribution with only 5 weighted members, correct?

This is an excellent point, which we should have mentioned. We interpolate both the melt errors (we need D as a continuous function of M, but only have five discrete points) and the sea level rise as a function of M (so that we can invert and find M as a function of SLR). In the updated ms, we have added a description of the interpolation procedure applied to both:

[lines 700-704]:
"To assimilate timeslice errors into the Bayesian update, we require P(\mathcal{O}|M, \mu) as a function of M, but the simulation only provides sparse points (figure 8b). To overcome this, we interpolate between the data points using a smoothing spline fit, via the fit function in MATLAB."

[line 627-628]
"Here, we obtain SLR as a function of M by linearly interpolating between individual M (see figure 8e)"

Thank you for your work and the consideration of my comments,
Andreas Wernecke

5th Dec 23

Dear Dr Bradley,

Your revised manuscript titled "How do we evaluate the contribution of anthropogenic climate change to sea level rise from Antarctica?" has now been seen by our reviewers, whose comments appear below. In light of their advice we are delighted to say that we are happy, in principle, to publish a suitably revised version in Communications Earth & Environment under the open access CC BY license (Creative Commons Attribution v4.0 International License).

We agree with the comments from the reviewer that readers would like to understand the focus and novelty of your study (e.g. ocean model, melt rate) before reading the Results section. For example, please briefly explain your novel methods in the final paragraph of the Introduction section, where you describe the overview of the study. Adjustment of the title, abstract and introduction (see our editorial requests table) should also help send your message to the readers.

We therefore invite you to revise your paper one last time to address the remaining concerns of our reviewers. At the same time we ask that you edit your manuscript to comply with our format requirements and to maximise the accessibility and therefore the impact of your work.

EDITORIAL REQUESTS:

Please review our specific editorial comments and requests regarding your manuscript in the attached "Editorial Requests Table".

*****Please take care to match our formatting and policy requirements. We will check revised manuscript and return manuscripts that do not comply. Such requests will lead to delays. *****

Please outline your response to each request in the right hand column. Please upload the completed table with your manuscript files as a Related Manuscript file.

If you have any questions or concerns about any of our requests, please do not hesitate to contact me.

SUBMISSION INFORMATION:

In order to accept your paper, we require the files listed at the end of the Editorial Requests Table; the list of required files is also available at https://www.nature.com/documents/commsj-file-checklist.pdf .

OPEN ACCESS:

Communications Earth & Environment is a fully open access journal. Articles are made freely accessible on publication under a CC BY license (Creative Commons Attribution 4.0 International License). This license allows maximum dissemination and re-use of open access materials and is preferred by many research funding bodies.

For further information about article processing charges, open access funding, and advice and support from Nature Research, please visit https://www.nature.com/commsenv/article-processing-charges

At acceptance, you will be provided with instructions for completing this CC BY license on behalf of all authors. This grants us the necessary permissions to publish your paper. Additionally, you will be

asked to declare that all required third party permissions have been obtained, and to provide billing information in order to pay the article-processing charge (APC).

Please use the following link to submit the above items:
[Link redacted]
** This url links to your confidential home page and associated information about manuscripts you may have submitted or be reviewing for us. If you wish to forward this email to co-authors, please delete the link to your homepage first **

We hope to hear from you within two weeks; please let us know if you need more time.

Best regards,

Shin Sugiyama, PhD
Editorial Board Member
Communications Earth & Environment
orcid.org/0000-0001-5323-9558

Heike Langenberg, PhD
Chief Editor
Communications Earth & Environment

On Twitter: @CommsEarth

REVIEWERS' COMMENTS:

Reviewer #2 (Remarks to the Author):

Dear Authors,
thank you for the response. All remarks of the previous reviews have been addressed and, most crucially, the scope and limitations of your work have been made more clear. The manuscript shows how aleatory ocean forcing can be included in a state-of-the-art, yet computationally efficient way. This, together with the study design and derived statistics, is an important step towards the evaluation of the contribution of anthropogenic climate change to sea level rise from Antarctica. However, despite several improvements in the current version, the manuscript is still not printing a whole picture of how such an attribution study would have to look like (in order to be considered ‚sound'). This seems not to be the aim of the study, as can also be seen by statements of the authors, such as:

„[…] our use of a detailed ocean model to calibrate melt rates is the novel factor of our study, while there is already established literature on ice sheet model parameters; while these parameters are clearly important, we have chosen to focus on the novel aspect of our work."

„It is worth noting that this is just a brief overview of what really requires a dedicated review paper of its own!"

This is also shown by your discussion that the AER values are sensitive to parameter values with limited advice on what is required to derive robust estimates of the AER (I realize that this knowledge might simply not exist yet). This is what my comments of selecting parameters (and to a lesser

degree parameter value ranges) were aimed at. For an attribution study the value, not the big picture, is the main interest and if the value might depend on which and how many parameters are considered, ambiguities arise and the robustness is in peril.

I agree with the authors that outlining all aspects of an reliable attribution study is vast and outside what can be reached and expected here. Therefore I request that the expectations raised, in particular, in the title and abstract are adjusted accordingly.

The title should include the main novelty of this work (the melt rate part and/or aleatory forcing+sampling parametric unc.) together with the context of working towards a realistic attribution of the anthropogenic SLR contribution. I could imagine something in the spirit of „... to help evaluate the anthropogenic contribution to sea level rise from Antarctica."

Also the abstract and introduction should put the focus more on the novelties of the MS and should be very cautious not to imply that you provide a framework for a real world attribution study.
In separate notes:
-it is not clear from the abstract that this work is about future SLR contributions.
-in the sentence "accounting for uncertainties arising from both poorly-constrained model parameters and stochastic variations" the „model parameters" should be singular, right?
-as discussed before, I am not convinced the word ‚robust' in „demonstrating that robust attribution of SLR is possible" is justified.
- the y-axis of Fig.10 has no numbers or labels.

Other than this, I recommend the manuscript for publication.

All the best,
Andreas Wernecke

Thanks once again to the editor for convening the review process. We're extremely pleased to hear that the journal is happy, in principle, to publish our manuscript. Here, we respond to the outstanding comments of reviewer 2.

Reviewer #2 (Remarks to the Author):

Dear Authors,
thank you for the response. All remarks of the previous reviews have been addressed and, most crucially, the scope and limitations of your work have been made more clear. The manuscript shows how aleatory ocean forcing can be included in a state-of-the-art, yet computationally efficient way. This, together with the study design and derived statistics, is an important step towards the evaluation of the contribution of anthropogenic climate change to sea level rise from Antarctica. However, despite several improvements in the current version, the manuscript is still not printing a whole picture of how such an attribution study would have to look like (in order to be considered ‚sound'). This seems not to be the aim of the study, as can also be seen by statements of the authors, such as:

„[…] our use of a detailed ocean model to calibrate melt rates is the novel factor of our study, while there is already established literature on ice sheet model parameters; while these parameters are clearly important, we have chosen to focus on the novel aspect of our work."

„It is worth noting that this is just a brief overview of what really requires a dedicated review paper of its own!"

This is also shown by your discussion that the AER values are sensitive to parameter values with limited advice on what is required to derive robust estimates of the AER (I realize that this knowledge might simply not exist yet). This is what my comments of selecting parameters (and to a lesser degree parameter value ranges) were aimed at. For an attribution study the value, not the big picture, is the main interest and if the value might depend on which and how many parameters are considered, ambiguities arise and the robustness is in peril.

I agree with the authors that outlining all aspects of an reliable attribution study is vast and outside what can be reached and expected here. Therefore I request that the expectations raised, in particular, in the title and abstract are adjusted accordingly.

We'd like to thank the reviewer for another thorough review and for their comments, which have helped to improve the manuscript. In particular, in response to these comments, we have made the novelty of our work much clearer and adjusted the expectations, in particular making it much clearer that there are still significant challenges to overcome in translating the work here to a real-world attribution study. We outline these changes in response to the individual comments below.

The title should include the main novelty of this work (the melt rate part and/or aleatory forcing+sampling parametric unc.) together with the context of working towards a realistic attribution of the anthropogenic SLR contribution. I could imagine something in the spirit of „... to help evaluate the anthropogenic contribution to sea level rise from Antarctica."

In our revised manuscript, we have adopted the editor's suggested title: "A framework for estimating the anthropogenic part of Antarctica's sea level contribution in a synthetic setting".

Also the abstract and introduction should put the focus more on the novelties of the MS and should be very cautious not to imply that you provide a framework for a real world attribution study.

In our updated manuscript, we have adjusted the abstract in the following ways to focus on the novelties and not imply that the we provide a framework for a real-world attribution study:

- We have been more careful to write that our results are applicable to the setup we consider, thereby reducing the implication that this is a framework for a real-world attribution study. (*"by simulating the retreat of a synthetic ice sheet setup using an ice sheet model", "We identify characteristic effects of climate change on sea level rise distributions in this setup…"*).
- We have also removed the final sentence about the implications of a real-world attribution study *("This paves the way for real-world attribution studies, which have significant implications for climate damage reparations and greenhouse-gas emissions policy"*), removing any connotation with a real world attribution study which might be drawn from this.
- We are quite constrained by the word count, but have made the novelties of the manuscript clearer by adding that *"These sources of uncertainty [parametric and aleatory] have only previously been considered in isolation."* In the updated manuscript, we now also expound upon the key novelties in the introduction (see below).

In our updated manuscript, we have adjusted the introduction in the following ways to focus on the novelties and not imply that the we provide a framework for a real-world attribution study:

- Remove the mention of framework [e.g. removing "the framework (i.e. an outline of the necessary experiments and analysis)"] to avoid the implication that we provide a direct framework for a real-world attribution study.
- We have briefly explained our novel methods in the introduction section, where we describe the overview of the study, as outlined above in the response to the editor: Lines 126-131: *"We demonstrate how uncertainties associated with poorly constrained model parameters interact with uncertainties associated with stochastic climate forcing, identifying that it is necessary to consider both, a feature that is lacking in current SLR projections. To the best of our knowledge, this is the first time such uncertainties have been considered simultaneously in an ice sheet modelling exercise."*
- We have added a sentence at the end of the introduction stating that we discuss the practical challenges of such an assessment in the discussion section of the manuscript. We explicitly mention that these challenges are not considered in our synthetic configuration. This makes it clear that our work does not directly translate into a framework for a real word attribution study: Lines 143-147: *"We finish with a brief discussion of the challenges associated with determining the role of anthropogenic forcing on SLR contributions from the WAIS, which are avoided in our use of a synthetic configuration. These include uncertainty in other model parameters, uncertainty in the initial state, and uncertainties in climatic forcing."*

In separate notes:
-it is not clear from the abstract that this work is about future SLR contributions.
We have left this deliberately unspecified because our work is appropriate to both past and future: we can construct an anthropogenic enhancement ratio for any time throughout the simulation, meaning that past and future are equally applicable – the present has no special status. We have added a mention of this in the abstract to avoid confusion:

"From these findings, we construct metrics quantifying the role of climate change on both past and future sea level rise, suggesting that its attribution is possible even for unstable marine ice sheets."

We also note this in the main text, where the AER is introduced, writing (lines 376-379):

"Note that, because the AER can be constructed for any time throughout the simulation, past and future SLR are equally applicable -- the present has no special status. Therefore, attribution statements may be made for either past or future SLR contributions (or both)."

-in the sentence "accounting for uncertainties arising from both poorly-constrained model parameters and stochastic variations" the „model parameters" should be singular, right?
That's correct – thanks for highlighting. We've corrected this in the updated manuscript.

-as discussed before, I am not convinced the word ‚robust' in „demonstrating that robust attribution of SLR is possible" is justified.
We have removed this word from the abstract in the updated manuscript.

- the y-axis of Fig.10 has no numbers or labels.
Thankyou for spotting this, we have fixed this in the updated manuscript.

Other than this, I recommend the manuscript for publication.

All the best,
Andreas Wernecke