## Reply to reviewer 1

**General Comments:**

This manuscript describes a set of idealized experiments using a coupled ice sheet and 2D sub-ice-shelf boundary layer model (IMAU-ICE and LADDIE), comparing the LADDIE results with those using simpler parameterizations that are widely used in the ice-sheet modeling community. The experiments are based on the MISMIP+, ISOMIP+ and MISOMIP1 setup, using the same bedrock topography and initialization along with a qualitatively similar forcing approach. The main findings of this work are that the 2D representation of the sub-ice-shelf flow can capture significant physical processes that are missing from simpler parameterizations, and that these processes have a major impact on both the timing and the nature of ice-sheet mass loss and retreat. The authors make a compelling case that 2D plume models like LADDIE can provide substantial improvements in physical accuracy at high model resolution (2 km in this study) compared with simpler parameterizations while remaining substantially more affordable than 3D ocean models.

I found the results compelling and the paper to be well structured and well written. The figures and tables do an excellent job of supporting the paper. The choices of what material to put in the main text and what to present in appendices also seemed great to me. The numerous experiments are well designed and the results are compelling, and complementary to ongoing work in realistic configurations with these components. These findings are also complementary with the MISMIP+, ISOMIP+ and MISOMIP1 results, which are in various stages of publication.

I have a number of more specific comments as well as a few suggested grammatical and typographical corrections below. After some minor corrections, I think this paper will be ready for publication.

Thank you for the positive feedback and for the constructive comments on how to further improve this manuscript. We agree with most of the comments and will implement them in the manuscript. Below, we provide a point-by-point response.

**Specific Comments:**

l. 17-19: “The parameterisations either inherently overestimate the persistence of margin thinning, leading to a sustained strong volume loss, or they underestimate margin thinning, delaying the onset of strong volume loss.” In the main manuscript, I think you do an excellent job of providing enough context that it is clear why LADDIE provides the higher physical fidelity and thus can be postulated to be closer to the “true” solution. I would be careful here in the abstract, though. This sentence in particular makes it seem like the results with LADDIE are the truth that can be used to evaluate the deficiencies of the parameterizations. I would strongly suggest you tone this down by adding something about “compared with LADDIE results” or something to that effect here.

Agreed, we will rephrase this part in the abstract as following:
“This results in a different transient volume loss **<u>between the parameterisations and LADDIE. Compared to LADDIE, the</u>** parameterisations either inherently overestimate the persistence of margin thinning, leading to a sustained strong volume loss, or they underestimate margin thinning, delaying the onset of strong volume loss.”

sec. 2.1.1: I think it would be really good to include the 2-km horizontal resolution here somewhere. I know it’s in Table B2 but that’s pretty buried and this is a fairly fundamental feature of the modeling in both IMAU-ICE and LADDIE.

In the submitted version of the manuscript, we mention the 2-km horizontal resolution in section 2.2.1. We agree, however, that it is good to mention this earlier in the model description so we will move the specification of horizontal resolution to 2.1.1. We will also add it to section 2.1.2, in which LADDIE is described. To avoid repetition, we will remove the comment on horizontal resolution from section 2.2.2.

l. 91-93: “First, it was demonstrated that the choice of sliding law has little effect on the results for perturbation experiments in an idealised setup.” I am concerned about this statement. It seems in direct contradiction to the findings of Cornford et al. (2020, https://doi.org/10.5194/tc-14-2283-2020), which found that the choice of basal friction led to the largest differences in model results in MISMIP+. There is an increasing consensus in the community (e.g. Joughin et al. 2019, https://doi.org/10.1029/2019GL082526) that Coulomb-limited friction laws like the Shoof parameterization that you use are both the most physically correct and the most consistent with observations. So I think it would be better to argue that you are using the “right” friction law, rather than that the friction law generally doesn’t matter.

By the way, Cornford et al. (2020) would be an important paper to cite in your paper, since your setup is similar to MISMIP+ as you comment.

We agree that the current phrasing may give the impression that the choice of sliding law has little influence on ice dynamics in general, which is not what we intended to communicate. To clarify, we will rephrase the paragraph as follows (or similar to this):

**"While the choice of sliding law has been shown to influence model results (Cornford et al., 2020), Berends et al. (2023), using the same ice sheet model and idealised domain as in our study, demonstrated that its impact is much smaller compared to that of the sub-shelf melt implementation. Given the focus on the latter, we conduct our experiments using a single sliding law: the Coulomb-limited modified power-law relation introduced by Schoof (2005). This choice is motivated by Joughin et al. (2019) which shows that Coulomb-limited laws best capture the dynamics at Pine Island Glacier."**

l. 121: I think "steady-state" needs a bit more explanation here even though you explain it later. In particular, LADDIE time-steps until it reaches steady state based on a given ice-sheet geometry and ambient ocean fields. It's important to make clear that the melt rates do evolve in time based on both ice-sheet geometry and (potentially, though not in this work) evolving ocean forcing.

Agreed, to clarify that this "steady-state" is connected to the combination of the ice shelf geometry and forcing at a given point in time, we will rephrase it as follows: "LADDIE is designed to simulate steady-state sub-shelf melt rates **for a given combination of ice shelf geometry and ambient ocean forcing.**"

l. 134-136: Could you say something about whether there is any coupling of heat fluxes?

The ice interior temperature influences ice shelf melt rates through the three-equation formulation. However, LADDIE does not modify the ice interior temperature, as there is no sensible heat flux into the ice at the ice-ocean interface. Consequently, the interior ice temperature remains constant throughout the simulation, following the MISOMIP1 protocol.

We will clarify this in Sect. 2.1.3 by adding something similar to:

**"Heat fluxes are coupled in one direction (output IMAU-ICE, input LADDIE). While the meltwater layer modelled by LADDIE is affected by the ice interior temperature through the three-equation formulation, the ice interior temperature itself is not affected by the meltwater layer, as there is no sensible heat flux into the ice at the ice-ocean interface."**

l. 144-145 and caption of Fig. 1: "The coupling between the two models is asynchronous, meaning that draft geometry and sub-shelf melt rates are exchanged at a certain frequency: the coupling frequency." "The models are coupled asynchronously, meaning that they exchange geometry and melt rates at a fixed coupling frequency (in this case: 8 times per year) which is independent of the time step in the individual models."

There are different understandings in the community about what "asynchronous" means but this is not a definition I am familiar with. I have heard "asynchronous" used to describe running one model, then the other (which is what you do, so in that sense the term would be appropriate). But I have also heard the term used to describe running one component with an accelerated time compared with the other (e.g. running the ocean for 1 year but the ice sheet for 10 years) each time you couple. To me, "synchronous" coupling would be to run each model simultaneously for a coupling interval, then exchanging information. This is not what you do, but also involves coupling at a fixed coupling frequency. I do not think "asynchronous" coupling is typically used to refer to coupling less frequently than every model time step, which I think might be your understanding of the term.

Thanks for pointing this out. Our intention was to convey that the models run sequentially rather than simultaneously while also clarifying that LADDIE is not necessarily computed at every ice model time step.

We will remove 'asynchronous' from l.144, and the caption of Fig. 1. We believe that both the schematic in Fig. 1 and the description in the second paragraph of Sect. 2.1.3 clearly convey that the models run sequentially (as you also inferred). Therefore, we will avoid using 'asynchronous coupling' anywhere in the text to prevent potential confusion.

l. 152-155: "The required runtime to reach a new quasi-steady-state depends on the flushing time...we run LADDIE for 4 days between each coupling step to ensure a near-stable meltwater layer thickness and velocity." In my experience with ISOMIP+ Ocean0 (which is qualitatively similar to the beginning of your high-melt experiments), it takes several months for a 3D ocean model to reach quasi-steady state, suggesting a flushing time on the order of months. I believe you when you say that LADDIE reaches a new quasi-steady state in 4 days for this setup, but that suggests that it doesn't require anywhere near the cavity flushing time to do so. Instead, it suggests that the previous quasi-steady state was close enough to the new quasi-steady state to require only minimal adjustment, presumably over a time far less than the cavity flushing time. This is great because it saves you a lot of computation! But I think it means the cavity flushing time isn't the relevant time scale after all and the paragraph maybe needs to be revised accordingly.

We agree that the cavity flushing time is not the relevant time scale to refer to in this context. As you suggest, a more appropriate reference is the flushing time of the meltwater layer, which represents a much smaller volume and typically involves higher velocities than the full ocean cavity. This results in a substantially shorter flushing time.

Using the total volume of the meltwater layer divided by total entrainment, we estimate flushing times of approximately 23 days in the moderate-melt scenario and 18 days in the high-melt scenario — consistent with the 40-day spinup phase required from rest.

As you point out, subsequent 4-day LADDIE runs do not need to span a full flushing time. Rather, the adjustment to a new quasi-steady state is much faster, due to only minor changes in geometry between coupling steps. To clarify, we will remove the flushing time from l. 152–155. We will reword this as follows:

**“The time required for LADDIE to reach a new quasi-steady state after each coupling step depends on the magnitude of changes in geometry and external forcing. Since we keep the oceanic forcing constant throughout our experiments, the adjustment time is driven by changes in geometry. In our setup, we found that running LADDIE for 4 days between coupling steps of 0.125 years is sufficient for the meltwater layer thickness and velocity to reach near-stable conditions.”**

We will add the meltwater layer flushing time as the relevant time scale to spin up LADDIE from rest (see our response to your comment below).

l. 158-159: “To address this discrepancy, we use nearest neighbour averaging to extrapolate the resulting sub-shelf melt field to include the grounding line cells.” Here is the part of the review where I tell you, perhaps unhelpfully, that this is not how I would have done things. I’m going to do that nonetheless because maybe we can have a discussion about it sometime. First, I’m a bit skeptical of the FCMP approach (even though I’m a co-author on the Leguy paper you cited for that). It seems like a low order choice from a mathematical perspective. But if that seems to be what works best, it is hard to argue with successful results. But regarding the extrapolation approach, wouldn’t an alternative (maybe a preferable one) be to just have the LADDIE domain cover all cells with centers that are floating? And maybe even cover all cells that are even partially floating. You can always compute a melt rate in LADDIE but then use the fractional area that is floating as part of computing the total melt flux in the cell that you pass to IMAU-ICE. This could be done even if you stick with the FCMP approach.

We see that we incorrectly stated this in the manuscript. LADDIE considers grid cells to be within the ice shelf mask when the ice shelf draft at the cell center exceeds the bedrock height at the cell center, thus satisfying the FCMP condition. We will correct this in the revised manuscript.

The nearest neighbour averaging is still implemented to address a practical issue that arises in our coupled setup: between coupling time steps, grounding line retreat can cause new cells to become afloat. While the sub-shelf melt parameterisations are immediately updated to reflect this change, LADDIE only updates its domain at the next coupling step. Without extrapolating melt rates to these newly floating cells, they would temporarily lack any melt input, potentially affecting the ice dynamics. However, in our experiments, the time step of the ice sheet model was equal to the coupling time step during the largest part of the simulation. As a result, the extrapolated melt was rarely applied in practice. We will clarify this in the revised manuscript.

l. 164: “and an upper limit of 0.125 years”: Do you stick with 1/8 of a year even for your 50,000 year initialization simulation? If so, why (since you’re not coupling)? If no, maybe state that since you are talking about initialization in this paragraph.

We used an upper limit of 0.125 years for the 50,000-year initialisation as well. While this was not strictly necessary, it remained in the configuration file and was used in the simulation. Since this detail is not specific to the initialisation, we have decided to move it—along with the horizontal resolution—to sect. 2.1.1 for better clarity.

l. 158: “To obtain a stable central grounding line position at X = 50 km…” Please state somewhere in the text that you have defined your coordinate system differently (offset by 400 km in X and 40 km in Y) compared with MISMIP+, ISOMIP+ and MISOMIP1. Otherwise, the locations you refer to will be confusing to colleagues who are familiar with the original protocols and geometry for those experiments.

Thanks for pointing this out. We will add the following to the beginning of this paragraph (sect. 2.2.1):

“The ice sheet model is initialised following the MISMIP+ protocol (Asay-Davis et al., 2016). **Compared to the geometry presented in the protocol paper, the coordinate system is offset by 400 km in the X-direction and 40 km in the Y-direction (Fig. 2).**”

We will also clarify it in the caption of Fig. 2.

l. 175-176; “We run the model for 40 days to obtain a quasi-steady-state melt pattern that aligns with the initial ice sheet geometry and forcing data.” To me, this reinforces the comment above about the flushing time. To me, it seems like the flushing time must be on the order of 40 days for this setup with LADDIE, and that the subsequent 4-day LADDIE runs to quasi-steady state do not need to run for a full flushing time.

Yes, that is correct — the meltwater layer flushing time is on the order of 40 days (see our response to the comment above). This time scale is relevant during the initial spin-up from rest, while subsequent runs require much shorter adjustments due to only minor changes in geometry. To clarify, we will revise lines 175–176:

“**We run the model for 40 days to obtain a quasi-steady-state melt pattern that aligns with the initial ice sheet geometry and forcing data. This time scale of 40 days equals approximately twice the flushing time of the meltwater layer, which depends on both the size of the ice shelf and the oceanic forcing. This meltwater layer flushing time is considerably shorter than the flushing time of the entire cavity due to the smaller volume of the meltwater layer and the higher velocities within it (Holland et al., 2017).**”

l. 190-191: “The salinity profiles are determined such that the density profiles of the moderate-melt and high-melt scenario are identical.” I don’t think you say anything in the paper about the equation of state you are using. Is it the same linear EOS as in Asay-Davis et al. (2016)? I think it may be necessary to mention that since you mention density profiles here.

Yes, we used the same linear EOS coefficients as in Asay-Davis et al. (2016), which we will list in Table A1 and reference in the text. These coefficients were applied consistently for computing the forcing in our experiments.

It is important to emphasise that the original PICO paper (Reese et al., 2018) used different EOS coefficients. For the PICO parameterisation, we retained those coefficients to maintain consistency with Reese et al. (2018). We will also add them to Table A1.

While double-checking, we discovered that the salinity profiles shown in Fig. 3 were based on the PICO linear EOS coefficients from Reese et al. (2018). However, the forcing applied in the experiments was correctly computed using the EOS coefficients from Asay-Davis et al. (2016). We will update Fig. 3 to show the salinity profiles corresponding to the actual forcing used.

l. 212-213 and 216: "…over the first period of the simulation…" "On longer time scales…": Especially because you only introduce time-series plots quite a bit later in Fig. 7, I think it would be important to give the reader a sense for what these times are – the first 200 to 400 years, and then the final 600 to 800 years would be my interpretation from Fig. 7.

We will follow your suggestion and rephrase:

"over the first **300 years** of the simulation", "On longer time scales **(600 to 800 years)**"

Fig. 4: I found the gray pixels in panels e-h confusing at first. I would explain them in the figure caption as areas of melt-through.

To clarify this, we will add the following to the figure caption:

"In all panels, the **ice-free** ocean is masked by dark grey. **Dark grey pixels enclosed by ice shelf pixels indicate areas of melt-through.**"

l. 244 and at least a dozen other places in the text: I believe when you use the term "collapse" here and in most other places in the text, you are referring to what I would call "melt-through". To me (and I think to the broader ice-shelf community), the term "collapse" implies larger sections of the ice shelf breaking up by fracture, and is considered distinct from melt-through (though fracture would likely occur in reality before melt-through is possible). I would ask you to change "collapse" to "melt-through" or a similar term except where you are referring to a larger break-up of the ice shelf by fracture.

Yes, we agree, we will follow your suggestion to replace all mentions of "collapse" by "**melt-through**", except for two occasions where we refer to possible break-up of the ice shelf not solely caused by melt (l. 36, l. 534).

l. 245-246: "When the western boundary current encounters one of these gaps in the ice shelf…" As written, I was confused by what was meant by "these gaps", since I did not understand "localized collapse" to cause gaps in the ice shelf but rather loss of larger sections. In revising this paragraph to reword the "localized collapse" phrase above, please make sure that it is clear what "gaps" are being referred to here.

We will change "gaps" to "**areas of melt-through**" to clarify this.

l. 246: “(assumption in LADDIE)”: I think this could use a bit more explanation here.

We agree that this explanation should be expanded. Areas of melt-through are treated the same as ice-free cells at the calving front. At these boundaries, the pressure gradient force drives a strong outflow of the meltwater across the calving front. Any momentum, heat, or salt advected beyond this boundary is lost to what is treated as an infinite open ocean.

To clarify this, we will add a paragraph to section 2.1.2 (LADDIE description) providing a more detailed explanation of how these conditions are handled in the model.

l. 490-491: “This leads to phenomena such as collapse of the ice shelf as a result of peak melt rates at locations with insufficient ice shelf thickness (Wearing et al., 2021).” In addition to being another place where I would replace “collapse” with “melt-through”, I would suggest rephrasing or explaining “insufficient”. I guess it seems like a tautology to me to say that melt-through happens at locations with insufficient ice-shelf thickness.

We agree, so we will rephrase it as follows:

“**This leads to phenomena such as melt-through of the ice shelf, which occurs when the applied melt rate exceeds the ice shelf thickness divided by the ice sheet model time step (Wearing et al., 2021).**”

l. 503-504 and Appendix D: “...for LADDIE, as well as for the Quadratic and Plume parameterisations, calving front retreat does not impact melting closer to the grounding line.” I was quite surprised by this. My results for MISOMIP1 IceOcean2, which has a similar approach to calving, though it used a 3D ocean model, showed qualitatively different results with and without calving (see e.g. Fig. 12 in Asay-Davis et al. 2016). I would have thought that 100-m thick ice would provide some fairly non-negligible buttressing in IMAU-ICE so that, even if melt rates didn’t change appreciably, the flow of ice would be affected more substantially. I just want to make sure the calving was applied in IMAU-ICE and not just used as a masking of melt in LADDIE. Maybe IMAU-ICE just behaves very differently than BISICLES, the ice-sheet model we were using.

The calving is indeed applied in IMAU-ICE, not just used as a masking in LADDIE. We will clarify this in the manuscript.

We can distinguish between two main potential impacts of calving in these simulations:
(1) Changes in buttressing, and
(2) Changes in melt rates near the deep grounding line.

Regarding (1), we observe minimal impact of calving on buttressing in the coupled LADDIE–IMAU-ICE simulations. This is because the 100-meter thickness threshold causes the loss of ice primarily along the western margin, which already contributes very little to buttressing in the default (no calving) configuration. Removing it entirely does not substantially affect the upstream ice.

For (2), calving can influence melt rates depending on how the sub-shelf melt implementation handles the geometry. For PICO, calving significantly alters the box configuration, consequently impacting melt rates near the grounding line. In contrast, LADDIE, Quadratic and Plume show minimal/no change in deep melt rates under calving. In a 3D ocean model, the ocean circulation could respond more strongly to calving front retreat – affecting melt rates near the grounding line.

We will clarify this distinction between the two effects in the revised manuscript to improve understanding of our results.

l. 513-514 and Appendix E: "...adjusting the deep melt rates to match those from the LA_H experiment (Appendix E)" It wasn't clear to me either here or in Appendix E what exactly is different in this retuning. My understanding was that the original tuning was done for the initial geometry such that the melt rate below -300 m was 30 m/yr. Since this procedure was used for LADDIE, a retuning using the initial geometry wouldn't change anything. So is the retuning for a different period of time? Or what am I missing? Could you describe the procedure and how it different from the initial tuning in more detail?

The original LADDIE tuning was performed for the initial geometry, with the averaged melt rates below -300 m set to 30 m/yr in the high-melt scenario, resulting in 16.75 m/yr in the moderate-melt scenario. The parameterisations were then adjusted to match the deep melt rates in the moderate-melt scenario, as this is the scenario we discuss most extensively.

This is explained in section 2.2.3, l. 203: "We then take the resultant averaged deep melt rates in the moderate-melt scenario to tune the parameterisations."

However, the scenario used for tuning the parameterisations may seem somewhat arbitrary. To address this, we also conducted retuning experiments, where the parameterisations were tuned using the 30 m/yr average melt rate from the high-melt scenario. These experiments are presented in Appendix E. We understand this could be confusing, so we will clarify it further in the Appendix by specifying in Table E that the default tuning is based on the moderate-melt scenario.

l. 520: "We want to emphasise that our coupled setup is not intended to fully replace coupled ice sheet-ocean models." I really appreciate you including this paragraph and this sentence in particular. I think it's an important caveat for the work.

Yes, we agree that this is important to mention.

l. 525-526: "In that context, our setup can be used to employ LADDIE's downscaling functionality by feeding it 3D ocean data to produce physically advanced 2D melt patterns on the ice sheet model grid." I really like this. This seems like a very powerful potential use for LADDIE and similar 2D boundary-layer models.

We fully agree and we are looking forward to working on this application in the future!

l. 539: Unless I missed it, ISMIP7 should be introduced as the "Ice Sheet Model Intercomparison Project for CMIP7". I think you can probably get away with not introducing CMIP7 but I will defer to the editors on that.

Agreed, we will include the full name for ISMIP7.

**Typographical, Grammatical and Formatting Suggestions:**

l. 15: I think "the parameterization's limitations" should probably be "the parameterizations' limitations" (i.e. both plural and possessive).

Agreed, we will follow your suggestion.

l. 113: Just a pet peeve of mine but maybe don't have back to back ") (" but instead use a semicolon.

We will change it, and we will move the name to the first occurrence of the LADDIE abbreviation in the introduction.

l. 147: "…resulting into a coupling interval of 0.125 years." I think this phrase is obvious from the frequency of 8 times per year and can be removed.

Agreed, we will remove it.

l. 189: "tangent hyperbolic" should be "hyperbolic tangent"

Thanks, we will change this.

l. 456-458: "Hence, we believe that despite the schematic character of our simulations, these observations of Crosson ice shelf can likely be explained by the 2D horizontal meltwater flow." I don't have any problem with this sentence scientifically – it's speculative but you make that clear. But I think there's a missing intermediate logical step here. May I suggest a rewording such as the following? "Although our simulations are schematic in character, we see qualitatively similar behavior. Hence, we believe that these observations of Crosson ice shelf can also likely be explained by the 2D horizontal meltwater flow." Still needs work but hopefully you get the point.

Thanks for this suggestion, we agree that this intermediate step improves the readability. We will rephrase this sentence as follows in the updated manuscript:

"**Although our simulations are schematic in terms of geometry and forcing, we see qualitatively similar behaviour. Hence, we believe that these observations of Crosson ice shelf can likely be explained by the 2D horizontal meltwater flow.**"

l. 460-461: "…characterised by a rapid retreat followed by suppressed volume loss." I think this phrase is redundant and can be removed.

Agreed, we will remove it.