

Response to Reviewer 2:

We would like to thank the reviewer for the thoughtful comments and constructive feedback. Below, reviewer comments (RC) are followed by our responses (AC).

**RC:** The authors relax sea ice concentration and thickness to PIOMAS, which is a reanalysis product. That is understandable given the lack of observational thickness data. Yet, I'd still be interested in seeing the authors consider examining the sensitivity of their results to other products. Maybe another reanalysis like GLORYS? But more importantly, as well as assimilating to PIOMAS, the authors also evaluate their results in comparison with PIOMAS. I would definitely like to see some evaluation using independent products, both reanalysis, but also observational. There are some satellite sea ice thickness fields available, as well as lots for concentration.

**AC:** We agree that comparisons with independent observational and reanalysis products are important when assessing the realism of a sea ice simulation or the performance of a data-assimilation system. However, the objective of this manuscript is different. This study was submitted to GMD as a model development paper, with the primary goal of describing, implementing, and verifying the Newtonian relaxation capability in SIS2.

The numerical experiments are designed to demonstrate that the relaxation algorithm behaves as intended in a coupled ocean–sea ice model. Specifically, we assess whether the relaxation can robustly constrain the model state toward a prescribed target field, whether the solution approaches the target as the relaxation timescale decreases, whether the response is smooth and monotonic, and whether the procedure introduces undesirable effects such as numerical instability, model shocks, or substantial drift in the coupled ocean state. For these objectives, the choice of target dataset is not critical. During development, the method was tested using a variety of target fields, including fields derived from other products and idealized synthetic ice distributions, and the algorithm behaved robustly in all cases. We have added a statement to the revised manuscript noting that the methodology has been successfully applied using alternative target products (Section 3.1).

For the demonstration experiments presented here, we selected PIOMAS because it is a widely used Arctic sea ice reanalysis that has been extensively evaluated and validated in previous studies cited in the manuscript. In the context of the algorithm validation study, PIOMAS serves as the prescribed “truth” (or reference) dataset. The relevant measure of success is how closely the relaxed model solution approaches the prescribed PIOMAS state. Consequently, all evaluation metrics are computed relative to PIOMAS, regardless of potential differences between PIOMAS, observations, or other reanalysis products. A comprehensive assessment of the relative accuracy of PIOMAS and alternative datasets would constitute a separate scientific investigation and is beyond the scope of this model development study.

**RC:** The authors state there is an initialization step, where the grid points where nudging will occur are identified? Does that mean these grid points are fixed in time? They don't change

between seasons? Or years, given there can be significant inter-annual variability in the Arctic sea-ice.

**AC:** We apologize if this point was not sufficiently clear in the manuscript. As described in Section 4.2.1, the initialization routine identifies the ocean grid points where nudging will be applied based on the user-specified relaxation mask and relaxation rates. Grid points with relaxation rates greater than zero are included in the nudging domain. The resulting set of grid points is fixed throughout the simulation and does not vary seasonally or interannually. The purpose of the relaxation mask is to define where relaxation is applied and to allow spatially varying relaxation strength (e.g., stronger relaxation in one region and weaker relaxation in another, or a spatially uniform relaxation rate as used in the ARC10k experiments). Temporal variability is introduced through the target sea-ice fields rather than through changes in the nudging domain itself. At each time step, the model solution at the selected grid points is relaxed toward the corresponding target values. Therefore, seasonal and interannual variability in the target sea-ice concentration and thickness fields is naturally reflected in the nudged solution, even though the nudging mask remains fixed in time. We have added a clarifying sentence in section 4.2.1 that states that the relaxation mask does not change in time.

**RC:** If I am following section 4.2.2 correctly, with the multiple categories and figure 2, it seems like most of the ice goes in within one category, with small changes in the thinner ice categories. Is this optimal? Might a smoother distribution work? Or addition of some ice in all categories? I think some more discussion and potentially some sensitivity analysis would help convince readers the author's approach is optimal. And for the thickness assimilation, does it go in at the maximum thickness in a category, the mean, or just whatever value you take from PIOMAS? Given this is a key aspect of the work, I think more detail and evaluation is justified.

**AC:** We thank the reviewer for this comment. The objective of the redistribution procedure is not to reconstruct the true subgrid-scale ice thickness distribution (ITD) from the target ice concentration and thickness fields. Because only bulk ice concentration and volume are prescribed (Eqs. 7–8), the problem is inherently underdetermined, and many different category distributions can satisfy the same constraints.

For this reason, we do not view the redistribution problem in terms of finding a unique or optimal ITD. Instead, the procedure is designed to satisfy two practical requirements: (1) conservation of the prescribed bulk ice concentration and volume, and (2) generation of an ice state that can be robustly integrated by SIS2. The adopted approach satisfies both requirements and has proven stable in all experiments performed with the relaxation scheme.

The reviewer is correct that most of the ice is assigned to a primary thickness category, while smaller amounts are distributed among thinner categories. The extent of this redistribution is controlled by the parameter  $\alpha_C$  (Eq. 14 and Fig. 2). The purpose of this redistribution is not to reproduce a realistic ITD, but rather to avoid highly singular states in which all ice resides in a single category. During development, we found that assigning all ice to a single category

occasionally resulted in failures of the SIS2 category redistribution procedure, whereas the adopted multi-category initialization did not exhibit such behavior.

A smoother distribution, or one that populates all categories, could certainly be constructed. However, the available target fields provide no information about the underlying ITD shape, and therefore no objective basis exists for selecting one smooth distribution over another. We explored alternative approaches during development, including methods that produced smoother category distributions, but found no evidence that they improved the performance of the relaxation algorithm. Consequently, we did not pursue a systematic sensitivity analysis of alternative redistribution schemes.

Regarding the target thickness, the bulk ice thickness from PIOMAS is not assigned directly to an individual SIS2 category. Instead, the target ice concentration and thickness are redistributed among the SIS2 categories while conserving the prescribed bulk ice concentration and ice volume. The primary category is selected based on the mean ice thickness over the ice-covered fraction of the grid cell (Eq. 11), and the remaining categories are populated according to the redistribution procedure described in Section 4.2.2.

We have added clarifying text to Section 4.2.2 and refer the reviewer to our response to a related comment from Reviewer 1.

**RC:** Is the relaxation approach the same in the sea ice pack interior as within the marginal ice zone?

**AC:** The relaxation approach is the same, however relaxation rate may vary depending on the relaxation mask that the user designs for the simulation. In the ARC10k experiments, the relaxation rate is spatially uniform over the relaxation region (north of 60°N). Consequently, the same relaxation is applied within the marginal ice zone (MIZ) and the pack ice whenever both lie within the relaxation region. In contrast, the relaxation mask used in the NEP10k experiments is spatially varying (Fig. 3c). As a result, the relaxation rate depends on location rather than on the sea ice regime itself. For example, when the MIZ is located in the Bering Sea, the applied relaxation is weaker than over the Arctic pack ice. Later in the season, when the MIZ retreats northward into the Arctic Ocean, it may be subject to relaxation rates similar to those applied within the pack ice. Thus, the relaxation algorithm does not explicitly distinguish between pack ice and MIZ conditions. Any differences arise solely from the spatial structure of the relaxation mask.

**RC:** For ARC10K and looking at figure 3c shows the 60N nudging domain going through the middle of the Hudson Bay Complex, across the Labrador Current and into the Baltic, all areas of active sea-ice. Is this latitude the best? Wouldn't a slightly lower latitude work better?

**AC:** We agree that the 60°N relaxation mask used in the ARC10k experiments is not necessarily optimal for all regions of the domain. In particular, the relaxation boundary intersects several seasonally ice-covered areas, including Hudson Bay, the Labrador Sea, and the Baltic Sea, where a different relaxation mask might be more appropriate for operational or process-

oriented applications. However, the primary objective of the ARC10k experiments was to evaluate the implementation and behavior of the relaxation algorithm rather than to optimize the relaxation mask for a particular application. The spatially uniform relaxation north of 60°N was intentionally chosen as a simple test configuration that allows the effects of the relaxation timescale to be assessed without introducing additional complexity associated with a spatially varying mask.

We agree that a different mask, potentially including spatially varying relaxation rates similar to those used in the NEP10k experiments, would likely be more appropriate for realistic simulations in the subpolar seas. The design and optimization of such masks, however, is beyond the scope of the present study. We have added clarification to the manuscript to make this distinction more explicit (section 4.3).

**RC:** No multi-year hindcasts with ARC10K? In any case, for the single-year hindcasts, why those specific years? Might the relaxation work differently in high vs low sea-ice years?

**AC:** Unlike the NEP10k configuration, no multi-year ARC10k hindcast simulations were available for this study. The ARC10k experiments were designed primarily to test and evaluate the implementation of the relaxation algorithm, and we considered one-year integrations sufficient for this purpose. Long-term NEP10k hindcast simulations were used instead to assess whether relaxation introduces slow-evolving drift in the coupled ocean–ice system. No evidence of such drift was found.

The specific year selected for the ARC10k experiments was not intended to represent a particular sea ice state. In principle, the effectiveness of the relaxation algorithm should not depend on whether a year is characterized by relatively high or low sea ice extent, provided that the relaxation mask encompasses the regions of interest. However, the impact of relaxation can depend on the spatial overlap between the sea ice cover and the relaxation mask. For example, in the NEP10k configuration the relaxation strength decreases southward. During years with extensive ice cover, portions of the marginal ice zone may extend into regions where the relaxation is weaker, whereas during low-ice years the marginal ice zone may be located farther north where relaxation is stronger. This illustrates the importance of carefully designing the relaxation mask for a particular application.

A systematic assessment of the sensitivity of the relaxation algorithm to interannual variations in sea ice conditions was beyond the scope of the present study. The primary objective here was to demonstrate and evaluate the implementation of the relaxation capability in SIS2 rather than to optimize relaxation parameters for specific sea ice regimes.

**RC:** In terms of evaluation, would like to see how the relaxation impacts the sea-ice velocity. And for the upper ocean, how the stratification and mixed layer depths respond.

**AC:** We agree that the response of sea ice velocity, upper-ocean stratification, and mixed-layer depth to the relaxation procedure would provide useful additional diagnostics. However, the primary objective of this study is to describe, implement, and validate the sea ice relaxation

capability in SIS2, with emphasis on its ability to constrain sea ice concentration and thickness while maintaining stable and physically consistent coupled model behavior. Accordingly, our evaluation focuses on variables that are directly affected by the relaxation algorithm, namely sea ice concentration and thickness. Sea ice velocity is not directly modified by the relaxation procedure and therefore falls outside the scope of the present validation. Similarly, while upper-ocean stratification and mixed-layer depth may respond to changes in sea ice conditions induced by relaxation, these diagnostics are more closely related to the physical response of the coupled ocean–ice system than to validation of the relaxation algorithm itself.

The manuscript already includes an assessment of Arctic surface temperature and salinity in the multi-decadal NEP10k simulations to evaluate whether the relaxation introduces spurious ocean drift. We have also examined impacts on additional oceanic properties and fluxes, which demonstrate that the relaxation influences the coupled system through its effects on sea ice distribution and associated air–sea interactions. However, presenting a comprehensive analysis of the dynamical response of the sea ice and upper ocean would substantially broaden the scope of the manuscript beyond its intended focus on development and validation of the relaxation capability. We agree that a detailed investigation of the response of sea ice dynamics, upper-ocean stratification, and mixed-layer processes to sea ice relaxation would be valuable and is a natural direction for future work.

**RC:** For the discussion, would the authors expect this to work for other models? What about for sea-ice in the southern hemisphere? In a changing/warming climate?

**AC:** The relaxation framework is not model-specific and can, in principle, be applied to any sea ice model that allows external forcing or state adjustment of ice variables. The algorithm itself can be used in both hemispheres as well as in global or regional configurations, provided the relevant sea ice state variables are available for adjustment. Regarding application in a changing or warming climate, the effectiveness of the method depends primarily on the fidelity of the target fields (e.g., observations, reanalyses, or externally generated model fields) used in the relaxation. The relaxation procedure itself simply constrains the model toward these prescribed states and does not introduce additional climate information. Therefore, its performance under climate change conditions will reflect the ability of the target dataset to represent evolving sea ice conditions, rather than any inherent limitation of the relaxation approach.

**RC:** Figure 1 caption – For ARC10K, since the caption mentions relaxation north of 60N, it would be good to indicate that latitude with a line on the figure.

**AC:** The figure and the figure caption have been edited.

**RC:** Table 1 – Is the background kinematic viscosity really zero?

**AC:** This is correct; in the present MOM6 configuration the background kinematic viscosity is set to zero. The MOM6 numerical formulation reduces the need for a prescribed constant

viscosity lower bound, as stability is primarily ensured through flow-dependent parameterizations and conservative numerics. Importantly, the effective viscosity is not zero. In this configuration, a biharmonic Smagorinsky-type nonlinear scheme is used to compute the lateral viscosity, which provides spatially and temporally varying dissipation and remains nonzero in regions of resolved shear. The effective viscosity is determined as the maximum of the background viscosity, the Smagorinsky viscosity, and other contributing terms, ensuring nonzero dissipation where shear exists.

**RC:** Section 3.2 title – Model validation metrics or Model evaluation metrics? I’m assuming the model code is valid, so this should more likely be the latter option.

**AC:** We have revised the section title accordingly to avoid ambiguity.

**RC:** Line 449 – “This supports the assumption that the Bering Sea ice state is weakly sensitive to Arctic LBCs.” Is this surprising given the flow through Bering Strait is generally to the north. What about during a northerly wind event when ice is exported to the south (which has been shown to occur in the literature).

**AC:** We agree that the result is not unexpected given that the mean ocean and sea ice transport through Bering Strait is northward into the Arctic Ocean. Nevertheless, this behavior cannot be assumed a priori, particularly in the NEP10k configuration where the northern lateral boundary is located relatively close to the Bering Sea. As the reviewer notes, observations have shown episodic southward ice transport through Bering Strait during northerly wind events, which can influence sea ice and ocean conditions in the northern Bering Sea. This possibility was one of the motivations for testing relaxation along the northern lateral boundary. In unconstrained simulations, unrealistically thick ice is often produced near Bering Strait and in the northern Bering Sea during heavy ice years. Although 2001 was not an extreme case, positive ice thickness anomalies are still evident in this region (Fig. 5b revised version). Our results indicate that relaxation applied only along the lateral boundary is insufficient to effectively constrain the northern Bering Sea ice state. In contrast, domain-wide relaxation, even with a weaker relaxation rate over the Bering Sea, substantially improves the simulated ice conditions. These findings support the conclusion that the Bering Sea ice state is only weakly sensitive to Arctic lateral boundary conditions in the context of the present experiments.

**RC:** Line 571 – Why might the relaxation being having no effect during the freeze-up?

**AC:** There are several possible reasons why the relaxation has a weaker impact on sea ice concentration during the freeze-up season. As discussed in the manuscript, sea ice concentration is inherently more difficult to constrain than sea ice thickness because it responds to oceanic and atmospheric forcing on much shorter timescales. During freeze-up, when sea surface temperatures are near the freezing point, relatively small variations in surface heat fluxes can rapidly promote or suppress ice formation, leading to changes in ice concentration that may partially counteract the relaxation tendency. In addition, the ARC10k simulations are forced with

JRA-55 atmospheric fields, whereas the PIOMAS target fields are generated using NCEP/NCAR Reanalysis 1 forcing. Differences between these atmospheric datasets may produce local thermodynamic conditions that are not fully consistent with the prescribed target ice state, reducing the effectiveness of the relaxation. This effect is expected to be more pronounced for ice concentration than for ice thickness, which evolves on longer timescales and is therefore more readily constrained. Finally, the relaxation experiments presented here use monthly target fields. Higher-frequency target fields, together with stronger relaxation, may provide improved control of ice concentration during the freeze-up period. Additional discussion of these factors has been added to Sections 7.2.1 and 9 of the revised manuscript.

**RC:** Figure 13 – Looks like the different experiments have different run length. I.e. one goes to 2025. Unless there is something relevant to be discussed in that extended period, just truncate in the plots and show all results for the same time frame.

**AC:** Figure 13 has been revised to show all experiments over a common time period.