The Cryosphere

Calibrated sea level contribution from the Amundsen Sea sector, West Antarctica, under RCP8.5 and Paris 2C scenarios

Sebastian H. R. Rosier[1,2], G. Hilmar Gudmundsson[1], Adrian Jenkins[1], and Kaitlin A. Naughten[3]

[1]Department of Geography and Environmental Sciences, Northumbria University, Newcastle, UK
[2]Department of Geography, University of Zurich, Zurich, Switzerland
[3]British Antarctic Survey, Cambridge, UK

Correspondence: Sebastian H. R. Rosier (sebastian.rosier@northumbria.ac.uk)

Abstract. The Amundsen Sea region in Antarctica is a critical area for understanding future sea level rise due to its rapidly changing ice dynamics and significant contributions to global ice mass loss. Projections of sea level rise from this region are essential for anticipating the impacts on coastal communities and for developing adaptive strategies in response to climate change. Despite this region being the focus of intensive research over recent years, dynamic ice loss from West Antarctica and in particular from the glaciers of the Amundsen Sea represents a major source of uncertainty for global sea level rise projections. In this study, we use ice sheet model simulations to make sea level rise projections to the year 2100 and quantify the associated uncertainty using a comprehensive Bayesian approach aided by deep surrogates. The model is forced by climate and ocean model simulations for the RCP8.5 and Paris 2C scenarios, and it is carefully calibrated using measurements from the observational period. We find very similar sea level rise contributions of 19.0 ± 2.2 and 18.9 ± 2.7 mm by 2100 for Paris 2C and RCP8.5 scenarios, respectively. A subset of these simulations, extended to 2250, shows an increase in the rate of sea level rise contribution, and clearer differences emerge between scenarios, with increasing snow accumulation in RCP8.5 resulting in less cumulative mass loss. Our model simulations include both cliff-height- and hydrofracture-driven calving processes, and yet we find no evidence of the onset of rapid retreat that might be indicative of an unstable calving front retreat in any simulations within our modelled time frame.

1 Introduction

The Amundsen Sea region in Antarctica plays a pivotal role in influencing global sea levels, and the dynamics of its glaciers have garnered significant attention in the context of climate change. Since the advent of satellite observations, glaciers in this region, in particular Pine Island and Thwaites glaciers, have undergone periods of extensive thinning and retreat, believed to be largely driven by increased ocean melting of their floating ice shelves (Pritchard et al., 2012; Paolo et al., 2015; Gudmundsson et al., 2012; Jenkins et al., 2018). While the region already accounts for a substantial portion of current rates of global sea level rise (SLR), vulnerability to unstable retreat might lead to greatly increased rates of ice loss (Davison et al., 2023; Rosier et al., 2021; Pollard et al., 2015; Feldmann and Levermann, 2015). Modelling challenges, most notably the representation of ocean forcing, model initialisation, and the presence of possible future tipping points, mean that there is a wide spread in SLR projections (Robel et al., 2019) and the latest Intergovernmental Panel on Climate Change (IPCC) special report emphasises the "deep uncertainty" associated with ice sheet model projections (Pörtner et al., 2019).

Existing SLR projections for the Amundsen Sea Embayment (ASE) vary widely, with high-end scenarios suggesting contributions of ~ 0.3 m by 2200 under strong warming pathways (DeConto et al., 2021). The relative SLR contribution of Pine Island and Thwaites glaciers has varied considerably over the historical record, but recent research has focused increasingly on Thwaites Glacier due to its vulnerable configuration and greater SLR potential. Past studies have also identified tipping points where retreat past certain

thresholds could trigger irreversible loss via marine ice sheet instability (Rosier et al., 2021), potentially leading to eventual collapse of much of the West Antarctic Ice Sheet (Feldmann and Levermann, 2015). Due to its complexity and importance, the ASE region has attracted a lot of research interest in the past decade, and we include a more complete overview of previous modelling studies in Sect. 6.

A major challenge in modelling the response of Antarctic glaciers to future emissions scenarios lies in the representation of ocean processes within ice sheet models. Generally, ice sheet modelling studies use simple parameterisations of basal melting; however, this approach cannot hope to capture the complex response of the ocean and ice shelf cavities to atmospheric forcing. This is particularly true in the ASE, where increased ice flux through basal melting has not been driven by some simple increasing trend in ambient ocean temperature but rather by intermittent changes in thermocline depth, driven by complex non-local processes, leading to a spatially and temporally heterogeneous response in ice shelf melting (Holland et al., 2022; Jenkins et al., 2018). Increasingly, coupled ice-sheet–ocean models are becoming available that address this shortcoming (e.g. De Rydt and Naughten, 2024; Bett et al., 2024); however, these come with a computational cost that makes running large numbers of simulations impractical. This is currently a major downside of the coupled modelling approach, because being able to run a large number of simulations is necessary in order to be able to properly explore the potentially large uncertainty associated with a model prediction.

Here, we use an alternative approach by assuming that feedbacks between the ice sheet and ocean are sufficiently small, over the timescales of interest, that we can use standalone ocean model simulations to capture the response of ice shelf cavities to atmospheric forcing. We model melt rates with an adaptation of the 2D plume model of Lazeroms et al. (2018) that is driven by a stratified water column whose properties in each cavity are obtained from recent state-of-the-art model simulations of the region (Naughten et al., 2023). Both the ocean and ice sheet models are driven by atmospheric forcing from CESM1 for two emissions scenarios: RCP8.5 and Paris 2C. We use a comprehensive Bayesian framework to first calibrate our ice sheet model using observations and then evaluate the uncertainty arising from model parameters and internal climate variability, leading to SLR projections for the ASE region with corresponding uncertainty estimates.

2   Data and methods

In order to arrive at our final estimated sea level contribution and calculate the associated uncertainties, a number of different modelling components and methods need to come together. These can broadly be divided into two main steps. (1) Outputs from a large ensemble of forward model simulations for the period 1 January 1996–1 January 2021 (Úa-

obs) are compared to observations over the same period to arrive at a likelihood which, given priors for our model parameters, enables us to calculate posterior probabilities for uncertain model parameters. (2) A second set of model simulations (Úa-fwd) uses the calculated model parameter distributions to predict changes in ice volume between 1 January 2021 and 1 January 2100 for different emissions scenarios and quantifies the associated uncertainty. Since the forward model simulations are computationally expensive, the uncertainty quantification steps in (1) and (2) are performed on surrogates of the forward model, as described in more detail below. An overview of our methodology is presented in Fig. 1. In the description that follows, we only distinguish between the two sets of forward model simulations where they necessarily differ, e.g. in the model forcing. Where a difference is not stated explicitly, the details are the same. In Sect. 2.1, we give an overview of the ice-flow model, with a focus on highlighting the uncertain model parameters. These uncertain model parameters, related to ice flow, surface mass balance (SMB), basal mass balance (BMB), calving, and model initialisation, are listed in Table 1. Note that parameters related to surface mass balance and calving are not included in the Bayesian calibration step, as explained in Sect. 2.4 and Appendix C. Section 2.2 summarises the external forcing for all the simulations, Sect. 2.3 presents the observations used to constrain our model parameters, and Sect. 2.4 presents the model calibration methodology.

2.1   Ice-flow model

All simulations use the Úa community ice-flow model (Gudmundsson, 2020, 2024), which solves the dynamical equations for ice flow with the shallow ice stream approximation (SSTREAM or SSA; Hutter, 1983; MacAyeal, 1989), using the finite element method on an irregular triangular mesh. The momentum and mass conservation equations are solved simultaneously using a fully implicit forward time integration with respect to both velocities and thickness. Úa uses inverse methods to estimate the uncertain spatial distributions of the rate factor A and the basal slipperiness C, as described in Sect. 2.1.4. The rate factor determines the relationship between the strain rates $\dot{\varepsilon}$ and deviatoric stresses τ through the Glen–Steinemann flow law (Glen, 1955):

$$\dot{\varepsilon} = A \tau_E^{n-1} \tau, \tag{1}$$

where n determines the non-linearity of this constitutive law, and we treat it as an uncertain model parameter (Millstein et al., 2022). The relationship between basal traction, $\boldsymbol{t}_{\mathrm{b}}$, and basal sliding velocity, $\boldsymbol{v}_{\mathrm{b}}$, is given by a mixed Coulomb–Weertman sliding law of the following form:

$$\boldsymbol{t}_{\mathrm{b}} = \frac{\mu_k N}{\mu_k N + \beta^2 \|\boldsymbol{v}_{\mathrm{b}}\|} \beta^2 \boldsymbol{v}_{\mathrm{b}}, \tag{2}$$

where $\beta^2 = C^{-1/m} \|\boldsymbol{v}_{\mathrm{b}}\|^{1/m-1}$, N is the effective pressure, μ_k is the coefficient of friction, and m is an uncertain model

Table 1. Uncertain model parameters included in the Bayesian calibration step described in Sect. 2.4. Note that "Prior range" refers to the minimum and maximum values for parameters with uniform priors and the $\pm 3\sigma$ range for variables with Gaussian priors.

| Parameter	Name	Model component	Prior range
m	sliding exponent	ice flow	2–8
n	creep exponent	ice flow	2–5
p	precipitation factor	SMB	0.03–0.15
σ_T	temperature variability	SMB	-0.3277 to -0.3821
$C_d^{1/2}\Gamma_{TS}$	Stanton number	BMB	6.5×10^{-5}–4.6×10^{-4}
E_0	plume entrainment	BMB	7.9×10^{-3}–6.2×10^{-2}
$\dot{h}_e$	thickness change error	inversion	3×10^{-3}–3×10^3
γ_{sC}	C smoothness penalty	inversion	1.3×10^3–1.3×10^6
γ_{sA}	A smoothness penalty	inversion	0.25–250
γ_{aC}	C prior penalty	inversion	5×10^{-2}–5×10^5
γ_{aA}	A prior penalty	inversion	0.25–250
C_Q	calving runoff factor	calving	0–200
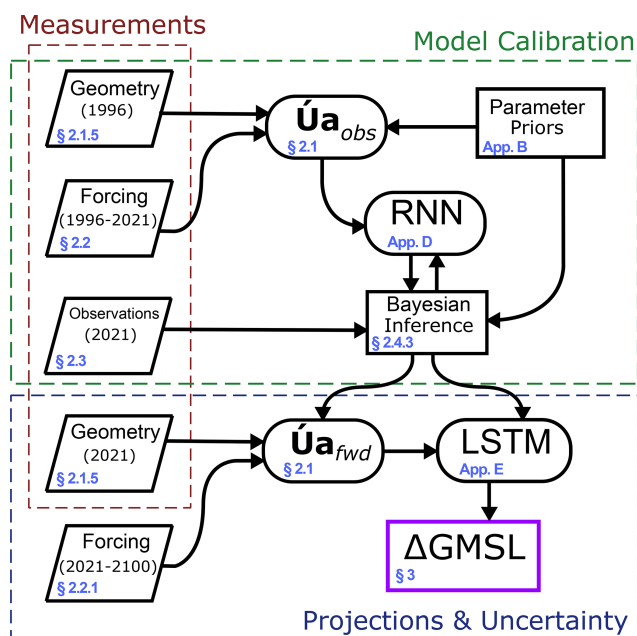
Figure 1. Schematic overview of the inputs, models, and processing steps that lead to projections of sea level contribution from the ASE up to the year 2100. GMSL stands for global mean sea level, RNN stands for residual neural network, and LSTM refers to a long short-term memory network type. We distinguish model calibration steps, resulting in posterior probability distributions for uncertain parameters (dashed green box) from projection/uncertainty steps (dashed blue box). Section or appendix labels for the relevant step are given in the lower left of each box.

parameter that controls the non-linearity of the sliding law in the Weertman term. The sliding law (Eq. 2) is arrived at by setting the basal traction, $\boldsymbol{t}_b$, equal to the reciprocal sum of

the Weertman, $\boldsymbol{t}_b^W$, and Coulomb, $\boldsymbol{t}_b^C$, tractions as follows:

$$\frac{1}{\|\boldsymbol{t}_b\|} = \frac{1}{\|\boldsymbol{t}_b^W\|} + \frac{1}{\|\boldsymbol{t}_b^C\|}, \tag{3}$$

where

$$\|\boldsymbol{t}_b^W\| = \beta^2 \|\boldsymbol{v}_b\|, \tag{4}$$

$$\|\boldsymbol{t}_b^C\| = \mu_k N. \tag{5}$$

This sliding law has been used previously in Barnes and Gudmundsson (2022) and has been implemented in the main Úa branch since 2020, along with the Weertman, Tsai, Coulomb, and Budd sliding laws.

We use the simplification that the effective pressure is the difference between ice overburden pressure and water pressure, assuming perfect hydraulic conductivity with the ocean. This results in $N = 0$ at the grounding line (GL) and N increasing inland as a function of ice thickness. In the limit as $N \to 0$, basal sliding follows a purely Coulomb behaviour, whereas as $N \to \infty$ basal sliding is purely Weertman. The use of a mixed sliding law such as this one is more frequently used by the community in recent years due to its flexibility in representing both plastic and hard-bed sliding, as well as an improved representation of conditions near the grounding line (e.g. Gudmundsson et al., 2023; Hill et al., 2024).

2.1.1 Surface mass balance

We use a positive degree day (PDD) model to capture two main processes: (1) changes in precipitation resulting from large adjustments of the ice sheet surface elevation and (2) production of runoff that contributes to both surface mass balance and the calving rate (Sect. 2.1.3). Note that other factors influencing changes in precipitation, such as those driven by temperature variations, are already accounted for in the CESM1 climate model, which provides our atmospheric forcing; therefore, these factors are not incorporated into the

PDD model. With this modification, the remaining formulation of our PDD model broadly follows the descriptions given in Huybrechts and de Wolde (1999) and Janssens and Huybrechts (2000). Precipitation (P) is given by

$$P = P_{\text{cesm}} e^{p\Delta T}, \tag{6}$$

where P_{cesm} is the bias-corrected (see Sect. 2.2) precipitation forcing from the CESM1 model, and $\Delta T = \gamma_p (s - s_{\text{init}})$, where γ_p is the lapse rate and where s_{init} is the initial ice sheet surface elevation. The parameter p captures the expected increase in snowfall arising from the increased moisture content of warmer air, as suggested by climate models (Aschwanden et al., 2019), and is treated as an unknown within our Bayesian framework.

Air temperature is modelled with a normal distribution whose mean is obtained from bias-corrected output from CESM1 (T_{cesm}), and the standard deviation (σ_M) is given by

$$\sigma_M = \sigma_T T_{\text{cesm}} + 1.66, \tag{7}$$

where σ_T is treated as an unknown parameter. This formulation for σ_M was chosen as it has been shown to provide a better fit to observations than using a constant standard deviation (Wake and Marshall, 2015), a result that was confirmed by our own testing. With this simple statistical model of air temperature, we can calculate the proportion of precipitation falling as rain, with the remainder therefore falling as snow. This snow accumulates and can be melted as a function of γ_{snow}PDD, where PDD is the number of positive degree days and where γ_{snow} is the degree day factor for snow. Any meltwater initially percolates down and refreezes as superimposed ice. If the amount of superimposed ice exceeds a threshold, no more superimposed ice is formed, and the meltwater contributes to a runoff term. We use the capillary retention model of Janssens and Huybrechts (2000) to calculate this threshold, which takes into account both the refreezing process and the capillary suction effect of the snowpack. The superimposed ice layer is melted if all the snow is melted (via a degree day factor for ice γ_{ice}), and finally the glacier ice itself can be melted if the ice layer is completely removed, both of which also contribute to a runoff term. Production of runoff is kept track of locally and is subsequently used in our calving law (Sect. 2.1.3), but we do not model the routing of meltwater across the ice shelf. Our PDD model is updated in monthly time increments in a separate time-stepping scheme to the ice sheet model (whose time steps adaptively change based on convergence).

2.1.2 Basal mass balance

Basal melting or ablation of ice at its base in these model simulations occurs as a result of transfer of heat from the ocean (on floating ice shelves) and frictional heating where ice is in contact with the bed (i.e. grounded). Ocean-induced basal melting accounts for over half of the mass loss of the

Antarctic Ice Sheet (Rignot et al., 2013) and is believed to be responsible for much of the measured changes during the observational period, so this component of the model requires particularly careful treatment. Complex feedbacks between ice shelf geometry, cavity circulation, atmospheric forcing, and conditions at the domain boundary mean that accurately predicting melt rates is a considerable challenge and generally requires computationally expensive general circulation models, an option that is not feasible for our purposes. Instead, we use an adaptation of the 2D plume model presented by Lazeroms et al. (2018) and Lazeroms et al. (2019) and originally proposed by Jenkins (1991) and Jenkins (2011). This retains much more physics than simpler but commonly used depth-dependent parameterisations while remaining computationally inexpensive to run. A detailed description of our modification to the implementation of Lazeroms et al. (2018) and Lazeroms et al. (2019) can be found in Appendix A, and what follows below is only a brief overview introducing the most relevant concepts and parameters.

The plume model considers a two-layer system in which an upper layer that is freshened by the addition of meltwater travels along the ice shelf base, driven by a buoyancy difference relative to the ambient ocean layer beneath that has temperature T_{a} and salinity S_{a}. The plume itself has temperature T_{p}, salinity S_{p}, thickness D_{p}, and velocity U_{p} and travels in a coordinate system orientated with X along the ice shelf base of slope α such that $X = 0$ at the grounding line. The plume equations result from conservation of mass, momentum, heat, and salt within the plume and are detailed in Lazeroms et al. (2019). As the plume travels in a positive X direction, its volume is increased through entrainment of ambient ocean water ($\dot{e}$) and a meltwater flux at the ice–ocean interface ($\dot{m}$). The entrainment rate is given by

$$\dot{e} = E_0 U_{\text{p}} \sin\alpha, \tag{8}$$

where E_0 is an uncertain model parameter. At the ice–ocean interface, melting arises from a balance between turbulent exchange of heat to the ice–ocean interface and latent heat of fusion in the ice:

$$\left(\frac{L}{c}\right)\dot{m} = C_d^{1/2}\Gamma_{\text{TS}}U_{\text{p}}(T_{\text{p}} - T_{\text{f}}), \tag{9}$$

where L is the latent heat of fusion in ice, c is the specific heat capacity of water, and the freezing temperature is a function of the plume salinity and the depth of the ice shelf base:

$$T_{\text{f}} = \lambda_1 S_{\text{p}} + \lambda_2 + \lambda_3 b, \tag{10}$$

where λ_1, λ_2, and λ_3 are constant parameters. The thermal Stanton number, $C_d^{1/2}\Gamma_{\text{TS}}$, is an uncertain model parameter that plays a key role in determining how effectively warm ocean temperatures can melt the ice shelf. While the original plume model of Jenkins (1991) solved a system of first-order differential equations to obtain the values of U_{p} and

$T_p - T_f$ needed to calculate the melt rate, Lazeroms et al. (2019) derived analytical expressions for those variables as a function of distance from the grounding line (detailed in Appendix A2), which we use in this study.

Beneath grounded ice, we calculate the frictional heating term resulting from basal sliding, which in areas of high basal sliding velocity can become non-negligible. We do not consider the contribution of geothermal heat flux since this term is more important for the ice sheet thermal state, which is inferred via our inverse approach, and in fast-flowing regions produces negligible melting compared to the frictional heating term. We calculate melting resulting from this frictional heating as

$$a_b = \frac{\boldsymbol{t}_b \cdot \boldsymbol{v}_b}{\rho L}, \tag{11}$$

where $\boldsymbol{t}_b$ and $\boldsymbol{v}_b$ are the basal traction and basal velocity, respectively. The resulting melt rates in the fastest-flowing regions are of the order of $1\,\mathrm{m\,yr^{-1}}$, which is 2 orders of magnitude larger than basin-wide average basal melting due to the geothermal heat flux (Joughin et al., 2009).

2.1.3 Calving law

We allow calving fronts to evolve dynamically during our simulations, using a calving law that depends on a parameterisation of local crevasse depth, as presented in Pollard et al. (2015) and DeConto et al. (2021). Total crevasse depth d_{tot} is a sum of terms depending on flow divergence, accumulated strain, ice thickness, and surface water. The overall calving rate is given by

$$R_c = 3000 \max[0, \min[1, (r - r_c)/(1 - r_c)]], \tag{12}$$

where $r = d_{\mathrm{tot}}/h$ is the proportion of crevasse depth to total ice thickness, and $r_c = 0.75$ is a critical value suggested by Pollard et al. (2015); i.e. calving occurs when a crevasse penetrates through 75 % of the ice thickness. Our primary motivation here is to include the possibility that a large increase in surface meltwater could lead to hydrofracturing and precipitate the breakup of ice shelves in the region. To this end, we explore the sensitivity of the calving law to the presence of surface meltwater via the parameter C_Q, such that the contribution to the total crevasse depth is $f_r C_Q R^2$ (DeConto et al., 2021). Here f_r is a factor that scales from 0 to a maximum of 1 for areas of low to moderate meltwater production (0–$1.5\,\mathrm{m\,yr^{-1}}$). Following DeConto et al. (2021), we use a range for C_Q of 0–$200\,\mathrm{m^{-1}yr^{-2}}$.

In addition to the calving law in Eq. (12) and also following Pollard et al. (2015), we implement the structural failure of ice cliffs necessary to allow for the possibility of a MICI (marine ice-cliff instability)-driven retreat. The cliff-based calving rate ramps up linearly from $0\,\mathrm{km\,yr^{-1}}$ at a cliff height of 80 m to $3\,\mathrm{km\,yr^{-1}}$ at a cliff height of 100 m. Our resulting final calving rate, defined everywhere as a field for

the level-set method, is then calculated as the maximum of either this cliff-based calving rate or the calving rate given by Eq. (12).

For a calving front to remain fixed, the calving rate must equal the ice velocity normal to the calving front. Generally, for a given calving law, those two quantities are not equal, resulting in migrating calving fronts over time. The Úa model uses the level-set method to solve for this moving boundary problem. The key idea is to introduce a new field variable, $\varphi(x, y, t)$, and define the calving front as the set of points in the x–y plane for which $\varphi(x, y, t) = 0$ for any time t. After initial initialisation at the start of the run, the level-set function (φ) is evolved by solving the level-set equation:

$$\partial_t \varphi + F \|\nabla \varphi\| = 0, \tag{13}$$

where the scalar F is the speed in the outward normal direction, $\hat{\boldsymbol{n}}$, to the calving front; i.e.

$$F = \boldsymbol{v} \cdot \hat{\boldsymbol{n}} - c,$$

where $\boldsymbol{v}$ is the material velocity and c the calving rate. A calving law provides c as a function of some other variables such as cliff height or state of stress, e.g. as $c = c(f, \sigma)$, where f is the cliff height and σ the Cauchy stress tensor. For the calving law to be physically admissible, the function f must be frame invariant, a condition that, as shown in Mitcham and Gudmundsson (2022), is not satisfied by all previously proposed calving laws.

The implementation details are provided in the Úa Compendium (Gudmundsson, 2024). A similar approach is used in the Ice-sheet and Sea-level System Model (ISSM), and further technical details for both the ISSM and the Úa ice sheet model are provided in Cheng et al. (2024).

2.1.4 Model initialisation

Úa uses the adjoint method to invert for C and A using observed surface ice velocity and the rate of change of ice thickness. It minimises the cost function $J = I + R$, where

$$R = \frac{1}{2\mathcal{A}} \int \left(\gamma_{sC}^2 \left[\nabla \log_{10}(C/\tilde{C}) \right]^2 + \gamma_{aC}^2 \log_{10}(C/\tilde{C})^2 \right) d\mathcal{A}$$
$$+ \frac{1}{2\mathcal{A}} \int \left(\gamma_{sA}^2 \left[\nabla \log_{10}(A/\tilde{A}) \right]^2 \right.$$
$$\left. + \gamma_{aA}^2 \log_{10}(A/\tilde{A})^2 \right) d\mathcal{A} \tag{14}$$

is the regularisation term where $\tilde{C}$ and $\tilde{A}$ are priors; $\mathcal{A}$ is the mesh area; and γ_{sC}, γ_{aC}, γ_{sA}, and γ_{aA} are regularisation parameters that determine how much the solution is penalised in terms of smoothness and deviation from the priors. I is the misfit term:

$$I = \frac{1}{2\mathcal{A}} \int ([u - u_{\mathrm{obs}}]/u_{\mathrm{err}}) d\mathcal{A}$$

$$+ \frac{1}{2\mathcal{A}} \int \left([v - v_{\text{obs}}] / v_{\text{err}} \right) \mathrm{d}\mathcal{A}$$

$$+ \frac{1}{2\mathcal{A}} \int \left([\dot{h} - \dot{h}_{\text{obs}}] / \dot{h}_{\text{e}} \right) \mathrm{d}\mathcal{A}, \tag{15}$$

where u_{obs}, v_{obs}, and $\dot{h}_{\text{obs}}$ are observations of surface ice velocity and changes in ice thickness, and u_{err}, v_{err}, and $\dot{h}_{\text{e}}$ are their respective observational errors.

By minimising the misfit to both velocities and thickness changes rather than just velocity (as has generally been done in the past), we can incorporate additional information that helps add further constraints on the inverted A and C fields. However, these two observational datasets are derived independently and generally not consistent with one another, meaning that our model will not be able to perfectly fit both the observed velocity and observed thickness change at the same time. A balance needs to be found such that our model matches both datasets as well as possible. To this end, we take the error terms as given for u_{err} and v_{err}, but we use a spatially uniform $\dot{h}_{\text{e}}$ whose magnitude is one of our uncertain model parameters. The effect of this is that $\dot{h}_{\text{e}}$ becomes a scaling parameter that determines to what degree the inverse procedure tries to match observed thickness changes against observed velocities. Details of the observations and errors described above are found in Sect. 2.3.

2.1.5 Ice sheet geometry

To define our ice sheet geometry for the two sets of simulations, we require data for the bedrock elevation (B), surface elevation (s), and ice thickness (h). We assume that B does not change throughout all our simulations, and here we use gridded data from BedMachine Antarctica version 3.4 (Morlighem, 2022). The Úa-*fwd* simulations, starting 2021, use the same dataset to define s and h. Measurements of these fields for the start of the Úa-*obs* simulations are more difficult to obtain, and so rather than using these highly uncertain products, we use a similar methodology to De Rydt et al. (2021). This approach makes use of the precise measurements by satellite altimeters to work backwards from the present-day geometry, subtracting the integrated surface elevation change trend over the period of interest, to calculate the ice sheet surface elevation, s, for 1996. This is fixed initially only for grounded ice that is further than 5 km from the coast in the present day, since altimeter measurements on floating ice are less reliable. By assuming all floating ice is at hydrostatic equilibrium, ice thickness can be solved for given s and B. Therefore, we make use of the position of the 1996 grounding line, for which we use the DInSAR-derived grounding line of Rignot et al. (2014a), which approximately corresponds to the period 1992–1996. With this constraint on the extent of freely floating ice for our starting period, we iteratively adjust s where it remains undefined until the grounding line position of our new geometry matches that 1996 grounding line. We add a regularisation term that penalises spatial gradients in s, ensuring that $s_{\text{BedMachine}} - s_{1996}$

does not vary by large amounts over small spatial scales and that there is a smooth transition between the fixed s field defined by altimetry and the s field that we solve for.

2.1.6 Model relaxation

Due to shocks in the transient solution immediately after initialisation, at least partly caused by insufficient knowledge of bed geometry, all our model runs go through a brief period of relaxation of a few years. The model is initialised, using the inverse methodology outlined in Sect. 2.1.4 to calculate A and C, and then forced with constant initial surface and basal mass balance for 3 years. The model is then re-initialised in an identical manner but with updated ice sheet geometry data that arise from the relaxation period. Following re-initialisation, the model is run for a further 3-year spinup with the same constant mass balance term, after which all model forcing terms evolve based on the CESM1 outputs for the period 1996–2021. A 3-year spinup was chosen as a balance between reducing large spurious thickness changes in the first few years of the simulation while also being as short as possible, and it is consistent with Nias et al. (2016). Without this approach, widespread regrounding occurs in shallow ice shelf cavities, and this approach greatly reduces this effect, although it persists somewhat in front of Pine Island Glacier, often leading to some grounding line advance in this sector.

2.1.7 Adaptive remeshing

We make use of the Úa ice-flow model's adaptive remeshing capabilities, such that the finite element mesh evolves with the changing ice in the domain. For each set of simulations, an initial coarse-resolution mesh is generated using mesh2d (Engwirda, 2014) and then immediately refined using nearest-vertex bisection. Elements are refined in regions of high strain rate and close to grounding lines and calving fronts. By starting with a coarse initial mesh and refining using this approach, the mesh can both refine and un-refine as, for example, the grounding line migrates inland. Target mesh resolution along the grounding line is 400 m and along the calving front is 1500 m. Remeshing is done every year and with the same refinement criteria for all simulations. Our mesh refinement criteria generally lead to maximum and minimum mesh sizes during simulations of approximately 12.5 km and 400 m, respectively. The size of each triangle is defined as the leg of an isosceles right triangle having the same area as the triangular element. An alternative definition of the size of a triangle-shaped finite element is the length of the side of a perfect square with an area equal to that of the triangular element. Using this alternative definition, size estimates are a factor of $\sqrt{2}$ smaller, and this smaller element size estimate is arguably more comparable to those used by finite-difference models.

2.2 Model forcing

Our ice sheet model is forced by changes at both the ocean and atmospheric boundaries. Motivated by the challenges in accurately representing ocean conditions, we drive our ocean melt parameterisation with the aid of recent state-of-the-art ocean model simulations of the region using MITgcm (Naughten et al., 2023). Rather than directly using the melt rates calculated in Naughten et al. (2023), which would not be representative of the state of our evolving ice shelf cavity geometry, we use modelled ocean conditions within the three main cavities to force our plume model parameterisation of basal melting. Specifically, we find the depth of winter water from the ocean model by searching for the coldest depth level in the top two-thirds of the water column, and we extract temperature and salinity from this depth. Then, masking everything above the winter water core, the depth of the base of the thermocline is found by calculating dT/dz for modelled temperature profiles and then searching for the first depth level (starting from the base) that exceeds a threshold value of 3×10^{-3} (determined empirically). Temperature and salinity at this thermocline base depth are then extracted, meaning that for each ocean model profile we calculate a thermocline depth and two temperature and salinity values. These are spatially averaged over three pre-defined catchments (Pine Island, Thwaites, and Dotson and Crosson) and sampled 10 times per simulation year for each scenario and ensemble member. In our Úa simulations, we identify the catchment for each node and then find the closest corresponding point in time from which to sample these five scalar values. We take the temperature and salinity values and create vertical profiles that vary linearly from the thermocline depth to the surface values and are constant below the thermocline. Our catchments are defined everywhere in the ice-flow model domain so as to consistently apply the correct ocean conditions to newly floating nodes as the grounding line retreats. More details on the implementation of the plume parameterisation can be found in Appendix A.

The ocean model simulations described above are themselves forced with outputs from the CESM1 model (Kay et al., 2015; Sanderson et al., 2017, 2018). Results exist for five core scenarios representing different anthropogenic and natural forcing pathways. For our Úa-*obs* simulations, we combine results from the *historical* scenario (1920–2005) with the RCP4.5 scenario (2005–2080) in order to span our desired period of 1996–2021. For this purpose, we select the RCP4.5 scenario that most closely matches observed atmospheric conditions in the model domain by calculating a total root-mean-square error (RMSE) in space and time compared to ERA5. For our Úa-*fwd* simulations, we include two scenarios: RCP8.5 and Paris 2C, which broadly represent more pessimistic or optimistic scenarios for anthropogenic forcing, respectively. These scenarios are further divided into 10 ensemble members each, which sample different realisations of internal climate variability within CESM1. Capturing this

variability is important, since it is likely to have a strong influence on the region and may have contributed considerably to observed trends (Holland et al., 2019, 2022). Therefore, all of our simulations randomly sample from these different ensemble members by sampling from a uniform distribution of ensemble IDs (E_{ID}), which are then one of the inputs to our surrogate model described in Sect. 2.4.2 and thus contribute to our uncertainty estimates.

We use the same CESM1 atmospheric forcing as in Naughten et al. (2023) to provide changes in air temperature and precipitation during the course of our simulations. These two fields then drive the PDD component of our model (Sect. 2.1.1), leading to changes in accumulation and surface melting. CESM1 modelled temperature and precipitation exhibit clear and systematic biases in the region which must be corrected for before they can be used to force the PDD model. Temperature shows an overall cold bias and differences in seasonal variability, which we correct for using processed data from 14 automatic weather station (AWS) instruments in the region (Janet, Kohler Glacier, Noel, Kominko-Slade, Kathie, Ferrigno, Up Thwaites Glacier, AUstin, Toney Mountain, Lower Thwaites Glacier, Inman Nunatak, Bear Peninsula, Evans Knoll, and Backer Island), obtained from the AntAWS dataset (Wang et al., 2023). For precipitation bias correction, we compared CESM1 precipitation in the period 1998–2018 to precipitation from the ERA5 reanalysis (Hersbach et al., 2017) and the MAR regional climate model (Marion et al., 2019; Donat-Magnin et al., 2020). In both cases, we create time-averaged but spatially varying interpolants from these data and calculate the difference to the time-averaged CESM1 field. This bias field is then added to the CESM1 forcing field in our model simulations, and this same correction is done for all simulations. Note that this is also the same methodology as in Naughten et al. (2023), meaning the atmospheric forcing in our ice sheet model is consistent with that of the ocean model.

2.2.1 Extended forcing

The model forcing described above, derived from climate and ocean modelling, only extends to the year 2100. To model ice sheet behaviour beyond this point in time, we require some method to extend our model forcing. Many different approaches could be used here, but our motivation in running these extended simulations stems from the inherently long response times of the ice sheet to perturbations in climate, and so we aim to reveal a more complete picture of the ice sheet's response to the climate perturbations imposed in our main set of runs (2021–2100). Thus, in our extended simulations we detrend and then repeat the forcing from 2080–2100. For every point in the domain and for temperature, precipitation, and MITgcm forcing fields, a linear trend is calculated for the 2080–2100 period and then removed, and this new 20-year forcing is then applied every 20 years cyclically. By detrending the forcing, we keep the

climate relatively constant to allow the ice sheet more time to adjust to its new state; furthermore, we avoid large discontinuous jumps in the forcing terms after each 20-year repeat cycle. We prefer to extend our simulations with a cyclical forcing rather than driving the model with constant conditions in order to preserve potentially important natural variability (Jenkins et al., 2018).

2.3 Observations

Key to our uncertainty quantification approach is the use of observations to update our prior estimates of uncertain model parameters. We also require observations to initialise the ice-flow model for simulations starting in 1996 and 2021, as described in Sect. 2.1.4. For both purposes, we use satellite observations of surface ice velocity and surface elevation change over the entire ASE region. While satellite products have the significant advantage of providing unparalleled spatial coverage, ascribing a precise date to a particular field is often difficult, since velocity and surface elevation change maps are mosaics made up of many repeat satellite passes. Based on the variable and sometimes not-well-defined dates of the remote sensing products we use in this study, we have selected simulation start dates of 1996 and 2021 as the dates that align most closely with all these inputs taken as a whole.

Velocity observations for Úa-*obs* (i.e. initialised for 1996) are from the MEaSUREs InSAR data for the Amundsen Sea Embayment (Rignot et al., 2014b). These data originate from the ERS-1 and are dated from 1 January to 31 December 1996. Details of the error sources and estimation are provided in Mouginot et al. (2012), and the resulting term is the square root of the sum of the independent errors squared. Velocity observations for Úa-*fwd* are from the MEaSUREs annual Antarctic ice velocity map (Mouginot et al., 2017a). These observations are dated from 1 July 2019 to the 30 June 2020, and details of the errors can be found in Mouginot et al. (2017b). In addition to their use in model initialisation step described in Sect. 2.1.4, the two velocity maps described above are converted to surface ice speed and then differenced to provide one component of our observations with which we calibrate the forward model (Sect. 2.4).

Satellite-derived $\dot{h}_{obs}$ values are obtained from ITS_LIVE, which combines a number of satellite sensor observations from 17 April 1985 to 16 December 2020 (Nilsson et al., 2022, 2023). These data are provided as a change in ice sheet elevation relative to 16 December 2013, at a monthly resolution. For initialisation of each forward model we use an average of this field in the 3 years preceding the simulation start date, whereas for model calibration we calculate the difference in ice sheet elevation between 1 January 1996 and 16 December 2020. As described in Sect. 2.1.4, the $\dot{h}_e$ error term is treated as an uncertain model parameter for the purposes of model initialisation. For model calibration, we use the supplied error term, which is a root-mean-square error as described in Nilsson et al. (2022).

2.4 Model calibration

The set of Úa-*obs* simulations is used to calibrate our model parameters, assigning probability distributions to these parameters based on how well they fit our observations of the region during the simulation period (1996–2021, Sect. 2.3). Model parameters contributing to uncertainty in our final estimates of SLR contribution are listed in Table 1. Nine of these parameters are included in the Bayesian inference step outlined in Sect. 2.4.3. Probabilities for the remaining three parameters, related to SMB (p and σ_T) and calving (C_Q), are specified separately due to there being insufficient model sensitivity to these parameters in the present day or over the relatively short observational period, as described in Appendix C.

2.4.1 Dimensionality reduction

Approximating probability distributions for large numbers of parameters and a complex model may require evaluating the forward model sequentially hundreds of thousands of times, and the computational cost of our forward model makes this intractable. As an alternative, we use the Úa-*obs* simulations as the training set for a surrogate model that is orders of magnitude faster to run than the forward model and therefore allows for proper sampling of the probability distributions. Therefore, we require a surrogate model that takes as input a set of model parameters $\boldsymbol{\theta}$ and outputs an expected change in ice speed and surface elevation that we can compare to observations.

Our set of training simulations suffers from the curse of dimensionality, since any model output we are seeking to emulate is defined at every node in the domain, but the information within these simulations is relatively sparse. We use singular value decomposition (SVD) to decompose our set of training simulations into principal components that make use of the strong spatial correlations in the model results. We build an $m \times n$ matrix $\mathbf{Y}$ of Úa-*obs* simulation results, where each row is an individual model run $\boldsymbol{y}_i$ in the m-member ensemble for a particular combination of model parameters $\boldsymbol{\theta}_i$ from the full sample of model parameters $\boldsymbol{\Theta}$, i.e. $\boldsymbol{y}_i = \mathcal{F}(\boldsymbol{\theta}_i)$, where each column represents a particular node and where $\boldsymbol{y}$ is either the modelled surface elevation change or surface speed change during the course of the simulation. Since the finite element mesh of each simulation evolves dynamically depending on the transient state of that simulation, the column entries of matrix $\mathbf{Y}$ arise from mapping model results to a single finite element mesh, consisting of n nodes, using the underlying shape functions. SVD involves finding the matrices $\boldsymbol{U}$, $\boldsymbol{S}$, and $\boldsymbol{V}$ such that

$$\mathbf{Y} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathrm{T}}, \tag{16}$$

where the columns of $\boldsymbol{U}$ and $\boldsymbol{V}$ are the left and right singular vectors, and they are both orthogonal such that $\mathbf{U}^{\mathrm{T}}\mathbf{U} = \mathbf{V}^{\mathrm{T}}\mathbf{V} = \mathbf{I}$. $\boldsymbol{S}$ is a diagonal matrix where the diagonal entries

are the singular values of $\mathbf{Y}$ and are ordered from largest to smallest. We define the matrix $\mathbf{B} = \mathbf{US}$, which represents the new basis, and each row of $\mathbf{B}$ represents the main modes of change between model simulations. The first four rows of $\mathbf{B}$ are plotted for Y_{sec} and Y_{vel} in Fig. 2, showing that the largest principal component captures an overall thinning and retreat signal for the entire ASE, while subsequent principal components represent variation in this response between the main catchments. By using only the first k columns of these matrices, we end up with a truncated SVD. The proportion of total variance captured for a given integer k is

$$f(k) = 1 - \frac{\sum_{i=1}^{k} S_{ii}}{\sum_{i=1}^{m} S_{ii}}, \tag{17}$$

and for each of the two $\mathbf{Y}$ matrices we choose k such that $f(k) > K$, where K is a hyperparameter as described in Sect. 2.4.2, and our final surrogate used $k = 8$ for surface elevation change and $k = 11$ for change in surface ice speed. This results in an approximation for $\mathbf{Y}$, given by $\tilde{Y} = \tilde{B}\tilde{V}^{\mathrm{T}}$, where $\tilde{B}$ and $\tilde{V}$ contain the first k rows of B and V.

In order to conduct our Bayesian inference in this newly defined framework, we need a way to project fields defined on our finite element mesh into the reduced k-dimensional space. Premultiplying by $\tilde{B}^{\mathrm{T}}$, followed by $(\tilde{B}^{\mathrm{T}}\tilde{B})^{-1}$, yields the desired mapping from the n-dimensional field y to the k-dimensional field $\hat{y}$:

$$\left(\tilde{B}^{\mathrm{T}}\tilde{B}\right)^{-1}\tilde{B}^{\mathrm{T}}y = \hat{y}. \tag{18}$$

2.4.2 Surrogate model

The SVD methodology outlined above provides a way to drastically reduce the dimensionality of our problem by encapsulating each model simulation as a linear combination of $\tilde{B}$ and $\tilde{V}^{\mathrm{T}}$. The matrix $\tilde{B}$ represents the main modes of change as exhibited in the set of training simulations in $\mathbf{Y}$, and each row of $\tilde{V}^{\mathrm{T}}$ represents the proportion of each of these k components for a particular set of model parameters. Since all our Bayesian inference is done in principal component space, our surrogate model is trained so that, for a particular combination of θ, it can accurately predict the corresponding row of $\tilde{V}^{\mathrm{T}}$ for both speed and surface elevation change (simultaneously). Outside of the existing simulations in Úa-obs, we do not know this mapping a priori. Thus, the input data used to train our surrogate model consist of the matrix of sampled model parameters Θ, and the network targets are the corresponding rows of $\tilde{V}^{\mathrm{T}}$. Following a similar approach to Brinkerhoff et al. (2021), we train a deep residual neural network for this purpose. A detailed description of the design and training of this surrogate model can be found in Appendix D. Once trained, the computationally efficient surrogate model can be used in combination with our parameter priors for Bayesian inference, as described in Sect. 2.4.3.

Before moving forwards with the Bayesian inference, we must verify that the truncated SVD and the surrogate model

(reprojected back to model space using the SVD) are able to properly represent the observations. It is entirely possible, if there are no model simulations in $\mathbf{Y}$ that show similarities to the observations, that there does not exist a k-dimensional $\hat{y}$ that matches closely to the observations. Similarly, although the surrogate model is trained to agree well with the forward model across the parameter space, there is no guarantee that the output of the model for the optimal set of model parameters agrees well with observations. To verify these two points, in Fig. 3 we compare the observations of change in surface elevation and surface ice speed with their equivalents, having been reprojected using the truncated SVD, and with the surrogate model output for the maximum a posteriori parameters (as described further below). This confirms that both the magnitude and spatial distribution of thinning and acceleration in the region are well approximated using the truncated SVD and our surrogate model.

2.4.3 Bayesian inference

We assume that within the complete parameter space Θ^* there exists a combination of parameters θ' that leads to a good agreement between model output $\mathcal{G}(\theta')$ and our chosen set of observations in the basis representation $\hat{y}$ (note that $\Theta \subset \Theta^*$ and presumably $\theta' \not\subseteq \Theta$).

Our search for θ' and its associated uncertainty is complicated by the presence of noise in the observational data and errors in our forward model. We use a Bayesian framework to infer probability distributions for our model parameters by updating our prior beliefs with observations to obtain a posterior probability. Using Bayes' theorem, we can express this posterior probability of certain model parameters given the observations as

$$\pi(\theta \mid y) = \pi(\theta) \cdot \mathcal{L}(y \mid \theta), \tag{19}$$

where $\pi(\theta)$ is the prior probability for θ, and $\mathcal{L}$ is the likelihood function. The choice of priors is an important step in any Bayesian analysis, and a more detailed description is given in Appendix B. For model parameters that do not have a directly measurable impact within our model, we choose uniform priors bounded by upper and lower limits chosen to span physically plausible values (see Appendix B). For parameters related to the surface and basal mass balance models which calculate measurable quantities independently from the ice-flow model, we use priors constrained on those observations.

The output of our surrogate model $\mathcal{G}$ is related to the observed changes in the basis representation by

$$\hat{y} + e_{\mathrm{obs}} = \mathcal{G}(\theta) + e_{\mathrm{model}}, \tag{20}$$

where e_{obs} and e_{model} are the measurement and model discrepancy terms, respectively. Furthermore, the error of our forward model can be broken down into two terms, i.e. $e_{\mathrm{model}} = e_{\mathrm{Ua}} + e_{\mathrm{rnn}}$, which are the errors in the forward ice
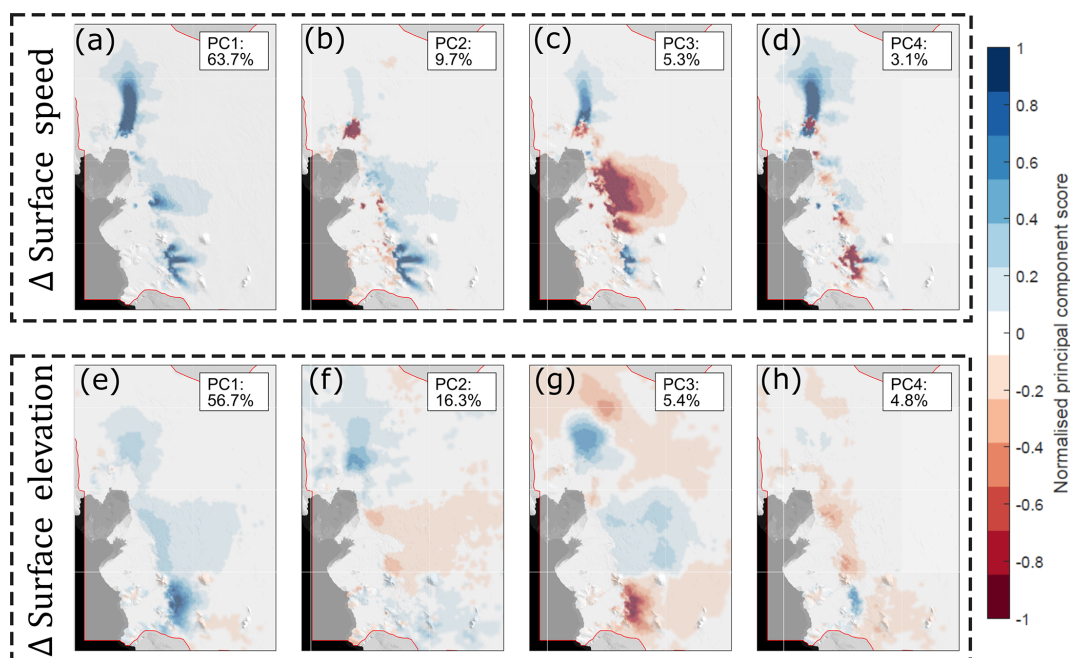
Figure 2. First four principal components for changes in surface ice speed **(a–d)** and surface elevation **(e–h)**. The percent of the total variance captured by each component is given in the top right corner of each panel. The shaded region surrounded with a red border indicates the extent of the model domain. Background image is from the MODIS mosaic of Antarctica 2008–2009 (Haran et al., 2014).
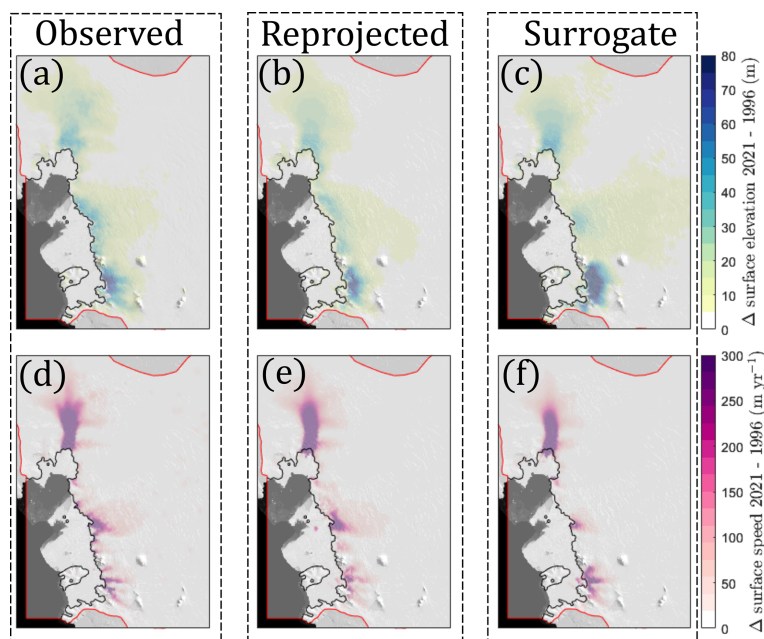

Figure 3. Comparison between observed **(a, d)**, reprojected via truncated SVD **(b, e)**, and surrogate **(c, f)** changes in surface elevation **(a–c)** and surface ice speed **(d–f)** for the period 1996–2021. Positive values indicate thinning and acceleration. Background image is from the MODIS mosaic of Antarctica 2008–2009 (Haran et al., 2014).

sheet model and the discrepancy between the forward model and the surrogate model, respectively. We assume that the three discrepancy terms (e_{obs}, e_{Ua}, and e_{rnn}) are statistically independent and Gaussian distributed, and we require an expression for the covariance matrix for each of these terms.

Observational errors are assumed to be uncorrelated, i.e. $e_{obs} \sim \mathcal{N}(0, \sigma_{obs}^2 \mathbf{I})$, where $\mathbf{I}$ is the identity matrix. Observations of surface elevation change and speed are derived from satellite measurements, and processing of these data includes (spatially varying) estimates of uncertainty. These supplied error estimates are derived from various sources including instrumental accuracy and the number of data points in a particular grid cell.

We can generally expect errors in the ice sheet model to be spatially correlated, since it is likely that the neighbouring nodes will also be similarly inaccurate if the model is not able to accurately capture ice flow at a given node. Thus, the discrepancy covariance will not be diagonal, and so $e_{Ua} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{Ua})$. We use an exponential covariance function to represent the spatial correlation in $\boldsymbol{\Sigma}_{Ua}$; i.e.

$$\boldsymbol{\Sigma}_{Ua} = \sigma_{Ua}^2 \exp\left(-\frac{\|x_i - x_j\|^2}{2\ell^2}\right), \tag{21}$$

where ℓ is a length scale, and σ_{Ua}^2 is the model discrepancy variance. This variance term, representing how well the model is able to match observations, is hard to quantify, but one reasonable argument is that we are not able to model the ice sheet more accurately than we can observe it; i.e. $\sigma_{Ua}^2 > \sigma_{obs}^2$. Here, we conservatively set the model discrepancy variance to be 4 times the mean observational error, i.e. $200\,\text{m yr}^{-1}$ in surface ice speed and $20\,\text{m yr}^{-1}$ in dh/dt. We can estimate the model error correlation length scale ℓ by fitting an exponential semi-variogram to the model results, whereby the range of the semi-variogram is equivalent to ℓ in Eq. (21).

Finally, the discrepancy between the ice sheet model and its surrogate can be estimated directly using the test set, which is hidden from the network during training. In this case, since the network is making a prediction in the reduced k-dimensional space and since each component is orthogonal, we can expect these models to be spatially uncorrelated; thus, $e_{rnn} \sim \mathcal{N}(0, \sigma_{rnn}^2 \mathbf{I})$. With these assumptions,

$$\sigma_{rnn}^2 = \frac{\sum_{i=1}^{n_{test}} (\mathcal{F}(\theta_i) - \mathcal{G}(\boldsymbol{\theta}_i))^2}{n_{test} - 1}, \tag{22}$$

where n_{test} is the number of simulations in the test set (566).

Importantly, the ice sheet model and observation discrepancy terms defined above are defined in model space, whereas the surrogate model discrepancy is defined in the reduced k-dimensional space. Thus, the final total covariance matrix, defined in the reduced k-dimensional space, is given by

$$\hat{\boldsymbol{\Sigma}} = \mathbf{A}\left(\sigma_{obs}^2 \mathbf{I} + \boldsymbol{\Sigma}_{Ua}\right) A^{\mathrm{T}} + \sigma_{rnn}^2 \mathbf{I}, \tag{23}$$

where $\mathbf{A} = (\tilde{\boldsymbol{B}}^{\mathrm{T}} \tilde{\boldsymbol{B}})^{-1} \tilde{\boldsymbol{B}}^{\mathrm{T}}$ follows from Eq. (18). With this estimate of the covariance matrix, we can calculate the likelihood term in Eq. (19), which is given by

$$\mathcal{L} = \prod_{i=1}^{N} \frac{1}{\sqrt{(2\pi)^{N_{out}} \det(\hat{\boldsymbol{\Sigma}})}}$$
$$\times \exp\left(-\frac{1}{2}\left(\hat{\boldsymbol{y}}_i - \mathcal{G}(\boldsymbol{\theta})\right)^{\mathrm{T}} \hat{\boldsymbol{\Sigma}}^{-1} \left(\hat{\boldsymbol{y}}_i - \mathcal{G}(\boldsymbol{\theta})\right)\right). \tag{24}$$

In practice, the posterior probability given in Eq. (19) does not have a tractable closed-form solution for such a complex model, so we use a Markov chain Monte Carlo (MCMC) algorithm to approximate its form numerically. For this, we use UQLab (Marelli and Sudret, 2014), which provides an extensive suite of uncertainty quantification tools for use within MATLAB. MCMC algorithms construct a Markov chain that explores the parameter space, whereby a chain is a particular set of parameters in a given iteration. In each iteration, a new set of parameters is proposed in each chain, and these are either accepted or rejected. After a sufficient number of iterations, the Markov chain converges such that the distribution of samples closely approximates the true posterior distribution. We used the affine invariant ensemble (AIES) algorithm to determine whether a proposed step for each chain should be accepted, due to its strength in dealing with correlation between parameters and only requiring one tuneable parameter (a, which we set to 1.5). Further details on the AIES algorithm and the choice of a can be found in Goodman and Weare (2010) and Allison and Dunkley (2013). We ran the AIES algorithm with 20 chains for 5000 iterations and verified convergence using the Gelman–Rubin diagnostic (Gelman and Rubin, 1992).

Once the Markov chains have converged, we take the final 1000 iterations for each chain and infer the marginal distributions and copula, once again using the uncertainty quantification tools within UQLab. Histograms show the converged samples for each chain, along with the fitted marginals plotted along the main diagonal of Fig. 4. The off-diagonal elements show the joint probability distribution for each pair of model parameters, helping to reveal correlations between certain model parameters (correlation coefficient between the pairs of parameters is included in the bottom left corner of the upper diagonal panels). An interesting side note here is that our probabilistic approach finds a distribution for n centred around 4. This lines up well with other recent studies which have used different approaches to suggest that a value closer to 4 than 3 may be more appropriate (Qi and Goldsby, 2021; Millstein et al., 2022). Having inferred the probability distributions for each parameter, we use Latin hypercube sampling to generate our sample of model parameters, with which we can run our forward model.
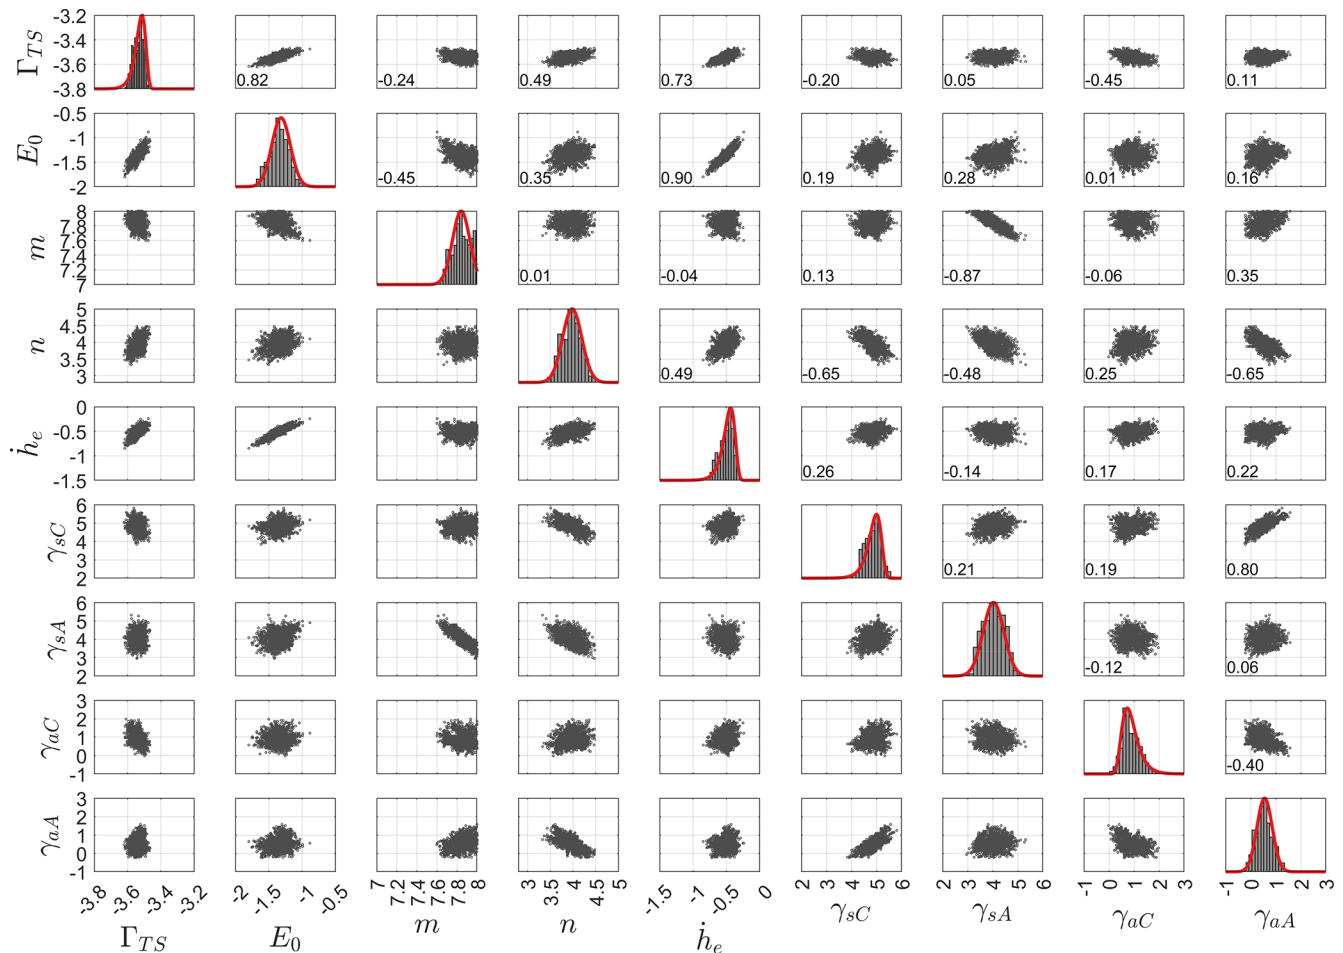
Figure 4. Parameter distributions for each converged chain in the final 1000 iterations of the MCMC algorithm. Red lines in the panels of the main diagonal show the fitted probability distribution for each parameter; blue lines show the priors (not shown if a uniform prior was used). The number in the bottom left corner of the upper diagonal plots shows the correlation coefficient between the two sets of parameters.

3   Sea level rise projections

We run our forward model until 2100, sampling from parameter probability distributions as calculated in Sect. 2.4. Sea level contribution is calculated as change in ice volume above flotation, assuming an area of the ocean of $3.625 \times 10^{14} \, \text{m}^2$, meaning that 1 mm of sea level contribution is equivalent to ~ 362.5 Gt of water equivalent ice added to the ocean. Once again, our uncertainty quantification approach requires a large number of forward model evaluations, and so we train a deep surrogate model that takes as input a vector of model parameters along with the forcing ensemble ID (E_{ID}) and that outputs cumulative change in ice sheet volume above flotation in each modelled year from 2021 to 2100. This is a different surrogate to that described in Sect. 2.4.2, and details of the network architecture, its training, and validation can be found in Appendix E. The surrogate is trained separately to make predictions for the two emissions scenarios, from a total training set of 2290 simulations, of which 20 %

are held back for validation and testing. The RMSE between forward model (target) and surrogate (predicted) cumulative change in volume above flotation by the year 2100 was 1.65 and 1.35 mm for the RCP8.5 and Paris 2C scenarios, respectively. Expressed as a percentage error between targets and predictions, the surrogate model error was 5.7 % and 9.5 % for Paris 2C and RCP8.5, respectively.

Using the surrogate model, we calculate a time series of sea level contribution until 2100 for one million samples drawn from the posterior parameter probability distributions, allowing for dense sampling of the parameter space and robust estimation of the probability distribution for each year. We show the resulting sea level curves for the RCP85 and Paris 2C scenarios, along with selected confidence intervals (denoted C.I.), in Fig. 5. This shows that there is almost no difference in sea level contribution by the year 2100 for the two scenarios and that the uncertainties are almost identical, although RCP8.5 has a slightly wider range of extremes than Paris 2C. Results from a one-million-member

ensemble of the Paris 2C scenario show a median SLR contribution of 19.0 ± 2.2 mm from the ASE by 2100 and 18.9 ± 2.7 mm for RCP8.5. The 5 %–95 % percentiles (the so-called "very likely range" in IPCC reports) are 15.6–22.9 and 15.1–24.0 mm for Paris 2C and RCP8.5, respectively. These projections represent a similar but generally slightly lower estimate than most previous estimates, as discussed in more detail in Sect. 6.

To explore the response of each glacier in more detail, we look to the Úa-*fwd* simulations directly rather than the surrogate model that only provides information on the sea level contribution. In Fig. 6, we show the final grounding line and calving front positions for all simulations in the year 2100, compared to their starting positions in 2021. To generate this plot, the domain is divided into cells, and the colour map represents how frequently a grounding line (brown) or calving front (green) lies in a given cell as a percentage of all grounding lines or calving fronts; i.e. 10 % would mean that a grounding line or calving front was present in this cell for 10 % of the total number of simulations. This figure shows extensive calving front retreat across the entire ASE, with an almost complete collapse of the Dotson and Crosson ice shelves and only very small ice shelves remaining in front of Pine Island and Thwaites glaciers. Grounding line movement in the region is generally more limited, with very few model simulations finding either significant re-advance or retreat by 2100. Some limited re-advance occurs frequently at the grounding line of Pine Island Glacier. This is most likely a result of inaccurate bed topography rather than a physical response of the glacier to climatic forcing, and this may lead to a slightly negative bias in terms of SLR contributions from this glacier.

4 Sources of uncertainty

We can explore the relative contribution to total uncertainty arising from different model parameters and internal climate variability using Sobol indices. These are based on the principle that the forward model can be expanded into summands of increasing dimension, such that the total variance is then given by the sum of the variance of these summands. The first-order Sobol indices ($\mathcal{S}_i$) represent the effect that only parameter $\boldsymbol{\theta}_i$ has on the variability of the model response ($\mathbf{Y}$); i.e.

$$\mathcal{S}_i = \frac{\text{Var}\left[\mathbb{E}(\mathbf{Y}|\boldsymbol{\theta}_i)\right]}{\text{Var}(\mathbf{Y})}. \tag{25}$$

Higher-order Sobol indices give the interactions between parameters, and the total Sobol index for a parameter is then the sum of all Sobol indices. We plot the total Sobol indices for each model input in Fig. 7. The first 12 Sobol indices for each scenario represent parametric uncertainty, while E_{ID} represents uncertainty resulting from internal climate variability.

Overall, the $\dot{h}_e$ parameter related to our inversion is the single most important for both scenarios, and all inversion

parameters together contribute 51 % and 66 % of total uncertainty for RCP8.5 and Paris 2C, respectively. Parameters m and n, related to ice flow, contribute 17 % and 19 %, while parameters related to basal melting (Γ_{TS} and E_0) contribute 18 % and 8 % for RCP8.5 and Paris 2C, respectively. Generally, Sobol indices for the RCP8.5 experiments show a stronger sensitivity to parameters related to external forcing (e.g. mass balance and internal climate variability), whereas the Paris 2C simulations are more sensitive to parameters related to internal ice dynamics and initialisation. This makes sense since the clearest difference between the two scenarios is changes in mass balance and particularly an increase in precipitation in the RCP8.5 scenario.

The most notable finding in terms of Sobol indices is the strong sensitivity of our sea level projections on parameters related to model initialisation. It implies a strong sensitivity to the resulting A and C fields whose spatial distribution these parameters alter. This in itself is not surprising, since uniform A and C fields would yield model simulations that have no bearing on the present state of the ice sheet and would have no meaningful predictive power. Additionally, our Sobol indices are calculated for predictions spanning a 79-year period, i.e. a relatively short-term forecast in terms of ice sheet modelling. We would expect that for longer simulations spanning hundreds of years this sensitivity to initialisation would become weaker, although – since the physical interpretation behind A and C is processes such as hydrology that generally vary with time – using fields that are constant in time becomes harder to justify over longer timescales.

It is important to emphasise that with this methodology we can only explore certain sources of uncertainty, mostly related to model parameters. Given the important role that surface mass balance appears to play in our simulations in offsetting increased grounding line flux, resulting in sea level projections at the lower end of published results (Sect. 5.1), it is clear that atmospheric forcing is an important consideration. We can explore uncertainty related to our PDD model but not the precipitation and temperature forcing from CESM1 that drives this. Clearly, an important avenue for future research should be to include forcing from other models and indeed the structural uncertainty associated with our choice of ice sheet and ocean model.

5 Extended simulations

We conduct two sets of simulations that extend a subset of the Úa-*fwd* simulations from their end date of 2100 to either 2250 or 2300. The first set of these extended simulations, to the year 2250, continues a large subset (608) of Úa-*fwd* model runs with the goal of exploring how the ASE sea level contribution evolves over a longer period of time in response to the changes in forcing imposed until 2100. The second set, extended to the year 2300 and saving all model fields at the end of each year, was conducted to enable a more detailed
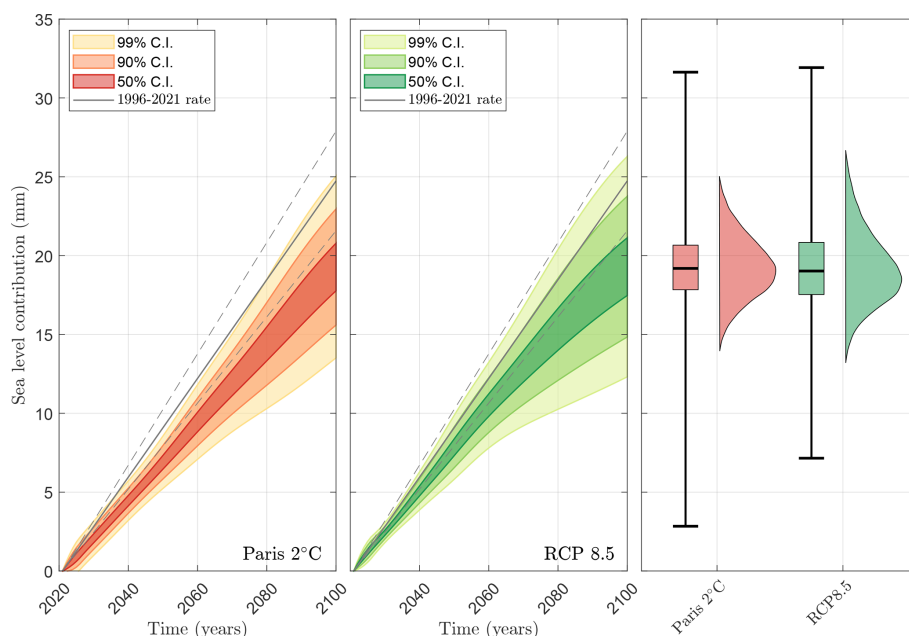
Figure 5. Sea level contribution from the ASE for the Paris 2C and RCP8.5 emissions scenarios, as calculated by the surrogate model. Also plotted is the mean observed 1996–2021 rate for the region (solid black line) and its propagated uncertainty (dashed black lines) as calculated using the input–output method (Davison et al., 2023). Shading indicates the 50 %, 90 %, and 99 % confidence intervals. Box plots in the right panel show the mean, interquartile range, and full range of SLR contributions for each emissions scenario.
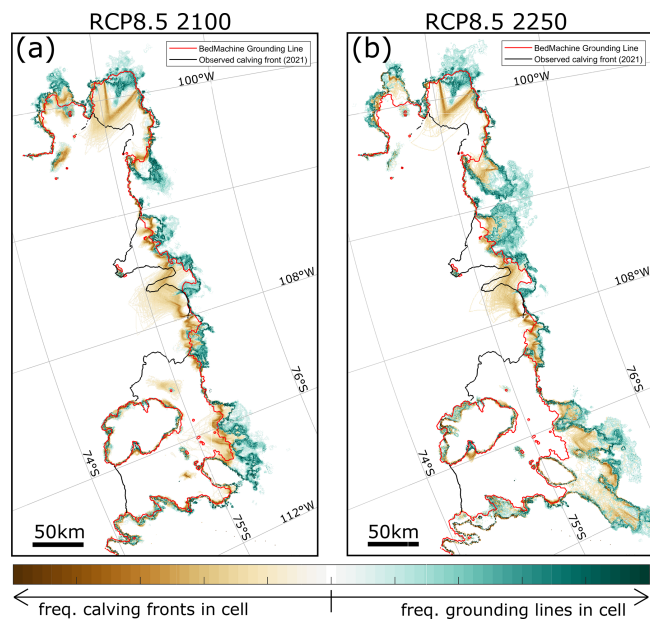
Figure 6. Grounding line and calving front positions for all Úa-*fwd* simulations **(a)** and Úa-*extended* simulations **(b)** for the RCP8.5 scenario. Locations of the grounding line (red line) and calving front (black line) in 2021 are also shown. Model results are transferred to a uniform grid, and the colour map indicates the proportion of simulations for which a grounding line (blue-green colour map) or calving front (brown colour map) is present in a particular grid cell.

analysis of our simulations and was necessarily much smaller (60 simulations) due to computational constraints. Note that since our extended ocean forcing is detrended and repeated from 2100, SLR estimates from these extended simulations may be conservative compared to studies where forcing continues to evolve with time.

5.1 Extended simulations to 2250

We randomly selected a total of 608 simulations from the Úa-*fwd* model runs and continued them from where they stopped in 2100, using the extended forcing described in Sect. 2.2.1. We show the cumulative sea level contribution for the ASE for the RCP8.5 and Paris 2C emissions scenarios in Fig. 8a. By 2250 there is a more substantial difference between the two scenarios, with a median sea level contribution of 6.0 and 7.1 cm for RCP8.5 and Paris 2C, respectively. These results are similar to recent modelling work using a calibrated melt rate parameterisation that found a maximum of 8 cm of sea level contribution by 2300 (Reese et al., 2023), but this is substantially lower than the ~ 0.3 m by 2200 in response to a $+3\,°C$ global warming scenario (DeConto et al., 2021).

Interestingly, our results show a generally greater sea level contribution under the Paris 2C emissions scenario. To explore the reason for this in more detail, we compare the area-integrated precipitation, total grounding line flux, and change in ice volume above flotation in the final year of these extended simulations. This shows very little difference between scenarios in terms of dynamic ice loss, whereas there
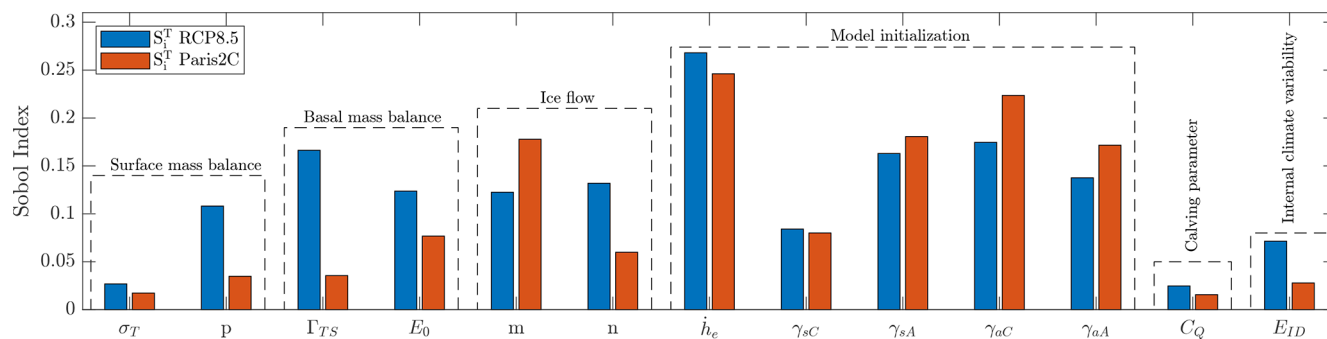
Figure 7. Sobol indices for each model input, representing the fractional contribution of each input on the uncertainty in our projections of sea level contribution for the ASE. Each input consists of two columns, showing the Sobol index as calculated for either the RCP8.5 or Paris 2C scenarios.
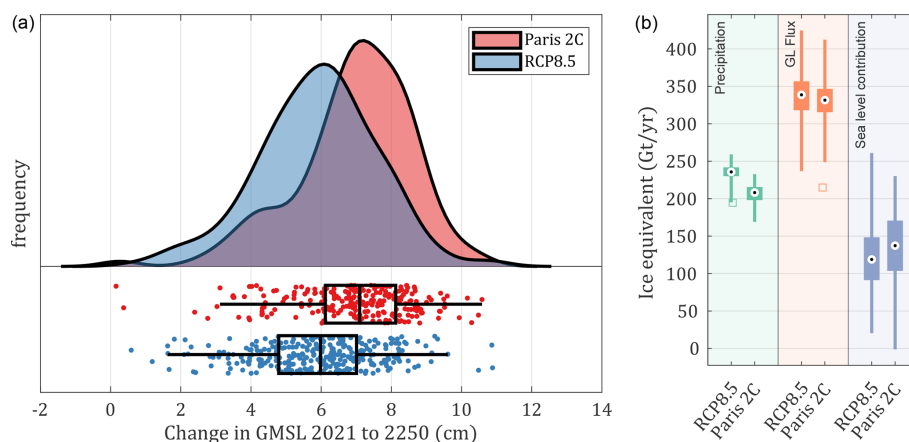

Figure 8. Extended simulation results, showing change in global mean sea level from the ASE region between 2021 and 2250 for the RCP8.5 and Paris 2C emissions scenarios **(a)** and the relative contribution of different ice mass terms for the final year of simulations **(b)**.

is a relatively large difference in terms of total precipitation. These results suggest that over longer timescales increases in precipitation for RCP8.5 may outweigh the relatively minor differences in ocean forcing, relative to the Paris 2C scenario, although with the caveat that our ocean forcing does not evolve beyond 2100 and may itself show a stronger scenario dependency over longer timescales.

To better understand how the regional mass loss evolves during the course of our simulations, we plot the rate of volume above flotation loss (or equivalently sea level contribution rate) for all simulations and both scenarios (Fig. 9). In both scenarios, there is a clear increase in the rate of ice loss after 2100, with the rate approximately doubling from 2100 to 2250 in both cases. In addition, a clear difference emerges between the Paris 2C and RCP8.5 scenarios up to approximately the year 2100. RCP8.5 shows a substantial decrease in the rate of mass loss that is not present in the Paris 2C scenario, but following this both scenarios follow very similar trajectories. By the year 2250, average annual rates of SLR contribution from the ASE reach 0.34 and 0.41 mm yr^{-1} for RCP8.5 and Paris 2C, respectively. This represents a similar

rate of SLR to the observed rate of 0.37 mm yr^{-1} observed in the period 1996–2021 (Davison et al., 2023). The majority of our model simulations start with a slightly lower SLR contribution rate than the observed rate (dashed black line in Fig. 9). This is in spite of the fact that our model is calibrated based on observed thinning and acceleration, and one of the terms that the inverse initialisation seeks to minimise is the observed thinning rate $\dot{h}_{\mathrm{obs}}$. However, as explained in Sect. 2.1.4, simultaneously matching both observed thinning rate and observed velocities was not possible, given the ice geometry from BedMachine. Trying to minimise both, in combination with our relaxation step, generally leads to an initial grounding line flux that is less than observed estimates, however both these steps lead to greatly improved model behaviour over longer timescales (see also Sect. 2.1.6).

5.2 Extended simulations to 2300

From the set of simulations described in Sect. 5.1, we evenly sampled 30 simulations from the sea level contribution distribution for both the Paris 2C and RCP8.5 forcing scenarios, resulting in a total of 60 additional simulations extended to
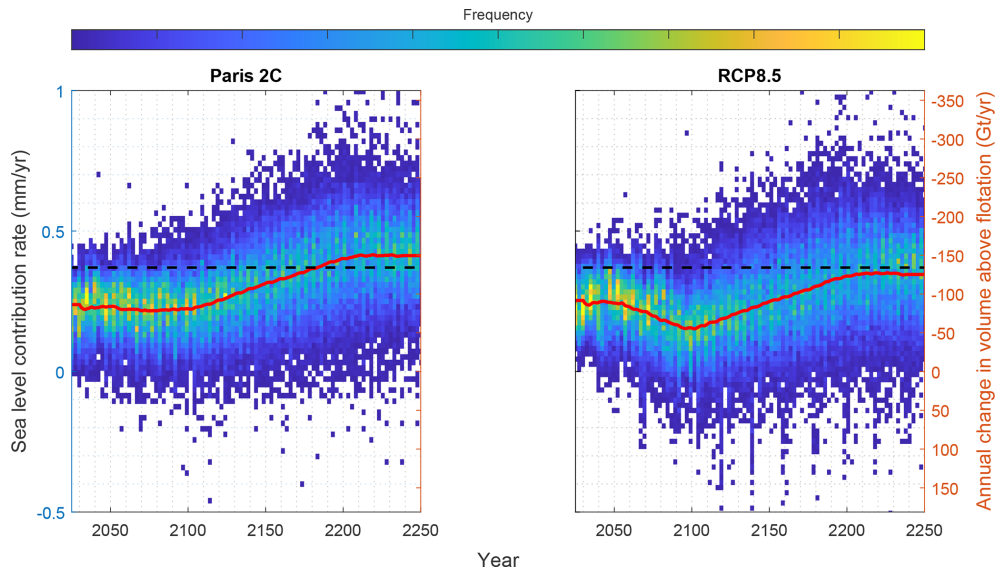
Figure 9. Annual rate of sea level contribution from the Paris 2C and RCP8.5 scenarios for all extended simulations. The red line shows the average across all simulations for each scenario, the dashed black line shows the observed rate from 1996 to 2021 (Davison et al., 2023), and the background colour map intensity represents how many simulations fall within a particular year/sea level bin.

the year 2300. With such a small subset of simulations, we do not capture the full uncertainty in our simulations, but this sample allows us to analyse simulations in much more detail to extract information that can help understand our model behaviour. In Fig. 10, we plot a time series of important metrics for each of the main catchments in our model domain.

In terms of overall change during these extended simulations, the Dotson and Crosson catchment shows the largest reduction in both grounded and floating area, with a corresponding increase in grounding line flux. Total ocean-induced melt follows the same trajectory as ice shelf area, with a very significant decrease until 2100, at which point most of the Dotson and Crosson ice shelves are completely gone, and melt is forcibly limited to relatively small ice shelves in front of the major outlet glaciers (see also Fig. A2). This catchment also shows a decreasing trend in accumulation for both scenarios, and an initial increase in surface melting only lasts until 2100 before reducing again, most likely as a result of the loss of floating shelf area where most of the melting takes place.

Turning our attention to Pine Island Glacier, we find an initial increase in grounding line flux generally levels out by 2100, at which point the two scenarios diverge slightly with an average reduction in flux for Paris 2C. Although the average grounding line flux time series appears relatively smooth, looking more closely at individual simulations we can see these are punctuated by periods of greatly enhanced flux, whose timing and magnitude are likely a result of periods of grounding line retreat. The clearest overall signal in terms of mass balance is in the accumulation term, which initially increases until 2100, particularly for RCP8.5, but then

levels out as a result of our extended forcing approach (see Sect. 2.2.1). Ocean melting for both scenarios decreases by $\sim 20\,\mathrm{Gt\,yr^{-1}}$ by 2100 as the floating area is reduced, at which point it levels off as the increasing grounding line depth offsets any continued loss in ice shelf.

Thwaites Glacier shows only relatively minor changes in grounding line flux, despite considerable grounding line retreat (see also Fig. 6). The Thwaites eastern ice shelf almost completely disappears in most simulations, but this loss in floating area is offset by the regrowth of the western ice shelf, promoting a large increase in ocean melt rates for this area. Accumulation shows a similar trend to Pine Island Glacier, with an increase until 2100 that is stronger in RCP8.5 than the Paris 2C scenario. Arguably, the most noteworthy feature in our simulations is that generally Thwaites Glacier does not appear to show strong acceleration in response to either grounding line retreat or loss of floating ice shelves. Grounding line retreat continues steadily until 2300, so this is not the result of the grounding line becoming stuck on some portion of the bed. One possible explanation is the formation of a new branch of Thwaites Glacier in almost all simulations, flowing out into northwest Pine Island bay (Fig. 6). As the grounding line retreats further and this new branch grows, it redirects much of the previous Thwaites Glacier flow and accounts for almost 50 % of the total grounding line flux for this catchment. The key difference to the present-day situation is the emergence of a buttressed ice shelf in front of this outlet, counteracting the increased mass loss that might be expected from such a retreated configuration.

Overall, our simulations show a complex response that varies considerably across each of the three regions. Taken
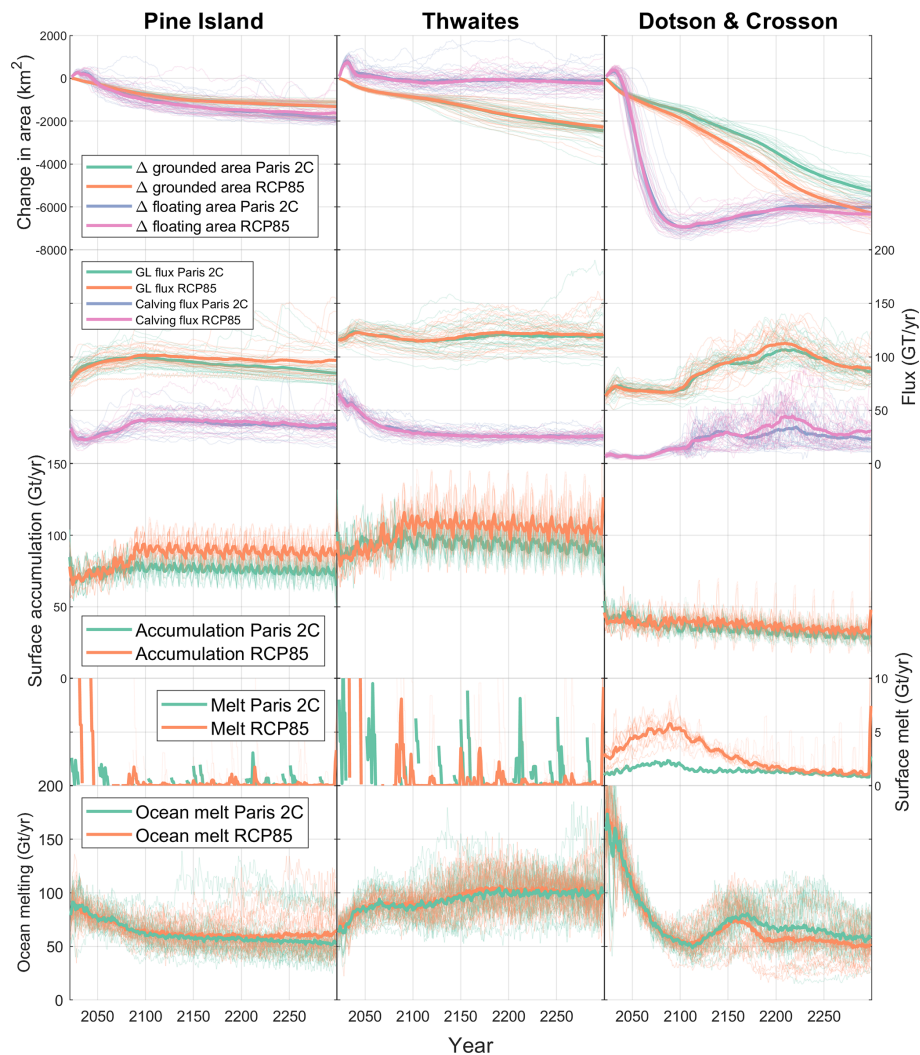
Figure 10. Time series of model metrics for our simulations extended to the year 2300, showing changes in the area of grounded and floating ice (first row), grounding line and calving flux (second row), surface accumulation (third row), surface melt (fourth row), and ocean-induced melt (fifth row). Each metric is calculated by integrating finite element mesh fields over pre-defined catchments consisting of Pine Island (first column), Thwaites (second column), and Dotson and Crosson (third column). Lighter lines show metrics for all simulations individually, and heavier lines show the average for each emissions scenario (i.e. Paris 2C vs. RCP8.5).

together, these result in little change or even a decrease in the rate of SLR contribution until 2100 (Fig. 9), at which point mass loss accelerates, reaching or exceeding present-day observed rates by 2200. Mass loss due to ocean melting decreases substantially by 2100, from ~ 300 to $\sim 200\,\mathrm{Gt\,yr^{-1}}$, averaged across simulations and scenarios, due mostly to a large reduction in floating area during this time. Conversely, accumulation increases over the same period. The acceleration in SLR contribution after 2100 is largely due to increased grounding line flux in the Dotson and Crosson region, together with a cessation of the increasing surface accumulation as we remove trends in the atmospheric forcing after 2100. Increased rates of SLR contribution persist until

2200, at which point they stagnate once again as grounding line flux from the Dotson and Cross region decreases.

A crucial factor to consider when investigating future sea level contribution of the ASE is whether or not a tipping point may be crossed, in which a positive feedback yields a greatly increased rate of ice loss. Robustly identifying tipping points with this set of simulations is not possible since we would require reversibility experiments or other tipping point analysis methods (e.g. Rosier et al., 2021); however, looking at results from these longer simulations may provide some indirect insight. Crossing a major marine ice sheet instability tipping point would be expected to manifest itself in a marked increase in grounding line flux that stands out from other experiments, and we do not see any behaviour of this kind in

any of our simulations. Similarly, a marine ice-cliff instability tipping point should lead to a very large increase in calving flux and although some simulations show quite a large three to fourfold increase in the Dotson and Crosson region, these increases do not appear to be self-reinforcing and leading to a runaway retreat. An important caveat is that, since the ocean forcing is held quasi-constant after 2100, our interpretation on tipping points in these extended simulations is limited to possibly delayed responses to perturbations in forcing up to 2100. If we were to force our model with a continued increase in ocean thermal forcing, it may be that a tipping point could be crossed. In the ISMIP6-2300 experiments, for example, where ocean thermal forcing increases considerably after 2100 (Fig. A1), some simulations show a complete collapse of the West Antarctic Ice Sheet by 2300 (Seroussi et al., 2024).

6　Projections in the context of previous studies

A number of previous ice sheet modelling studies have produced SLR projections for the Amundsen Sea Embayment. Direct comparison with many of these studies is made challenging since they cover a range of time periods, domains, and forcings; however, limited comparison helps contextualise the results presented here.

The recent study of Wernecke et al. (2020) is arguably the most similar to ours in terms of methodology, using a combination of model emulation and Bayesian calibration of parameters. For 50-year simulations, the spatially calibrated model ensemble resulted in a median of 16.8 mm sea level equivalent (SLE) with 90 % confidence intervals of 13.9 and 24.8 mm. Several other studies have produced SLR projections for the region using a model calibrated using Bayesian methods. In Nias et al. (2019), simulations with a duration of 100 years lead to a median SLR contribution of 55.7 mm, increasing to 139.7 mm after 200 years. Even after calibration, there was significant spread in SLR contributions after 200 years, with 5th and 95th percentiles of 56.2 and 424.3 mm, respectively. In another recent study, Bevan et al. (2023) found a median sea level contribution of 16.2 mm after 50 years, with 5th and 95th percentiles of −0.2 and 39.1 mm, respectively.

A number of other studies have explored the future sea level contribution of the ASE region under various climate scenarios but without the use of Bayesian methods to calibrate their models. Alevropoulos-Borrill et al. (2020) used a subset of CMIP5 simulations to force model simulations of the ASE and found an SLR contribution of between 20 and 45 mm by 2100 under RCP8.5. A study including the results of 16 ice sheet models found a median SLR contribution for the ASE of ∼ 20 mm by 2100 under the RCP8.5 scenario with CMIP5 model forcing (Levermann et al., 2020).

The studies described above cover a wide range of timescales, domains, and forcings, but in order to make a rough comparison, we convert the results listed above to a time-averaged rate and then average these rates across all studies to arrive at a mean rate of sea level contribution of 0.35 mm yr^{-1}, remarkably close to the observed value of 0.37 mm yr^{-1} for the period 1996–2021 (Davison et al., 2023). Global mean sea levels have risen by 3.7 ± 0.5 mm yr^{-1} in the period 2006–2018, and further acceleration is considered very likely (Fox-Kemper et al., 2021). So in the context of the recent observed global rate, this averaged rate represents ∼ 10 % of the current total. In contrast, in this study we find SLR contribution for RCP8.5 of 0.23 mm yr^{-1} between 2021 and 2100, with a "very likely" range of 0.14–0.30 mm yr^{-1}. Thus, our model projections to 2100 result in a median SLR contribution that is at the low end of the range from previous modelling studies of the region.

Our study differs from the studies outlined above in a number of important aspects. Firstly, we explore uncertainty related to a larger number of model parameters (12), including parameters related to calving, initialisation, ice dynamics, basal, and surface mass balances. Secondly, using adaptive mesh refinement on an unstructured mesh allows us to use finer resolution at the grounding line than most other similar studies (400 m in this study). On top of this, by solving velocity and thickness evolution fully implicitly, we expect the simulations presented to provide a more accurate representation of grounding line behaviour. We calibrate our model parameters using observations of changes in both surface elevation and surface ice speed (the model of Bevan et al., 2023, is the only other one to do this for the ASE). Our model simulations also include a dynamically evolving calving front position, rather than a fixed calving front as most previous studies have done. Finally, we use a parameterisation of the plume model to calculate melt rates in response to our evolving cavity geometry and forced by temperature and salinity from an uncoupled ocean model. While a fully coupled approach would clearly be preferable, this methodology includes more physics than the simple melt rate parameterisations used by other studies that do large ensemble simulations.

Although we cannot directly attribute the reason for our modelled SLR projections being at the low end of previously published results, one plausible explanation is differences in the treatment of surface mass balance. Our study makes use of a PDD model combined with atmospheric forcing from a climate model, including feedbacks in snow accumulation with temperature and surface elevation, that together lead to a strong increase in surface mass balance during the first 80 years of our simulations. In contrast, this feedback is either completely missing or partially absent in all the model studies discussed above, many of which place a heavy emphasis on the response due to ocean-induced melting instead (Wernecke et al., 2020; Nias et al., 2019; Bevan et al., 2023; Alevropoulos-Borrill et al., 2020; Levermann et al., 2020).

A weakness of our approach is that we only include results from one ice sheet model, and our forcing is derived

from only one climate and ocean model; thus, a large component of structural uncertainty is not included and would undoubtedly lead to considerably wider confidence intervals than presented here. This is perhaps most notable in terms of ocean forcing, where there is a relatively small spread in thermal forcing arising from the various ensemble members and emissions scenarios which is in contrast to, for example, ASE ocean properties in the ISMIP6 experiments (Fig. A1). In particular, there is little difference between the RCP8.5 and Paris 2C MITgcm forcing, which was one of the main findings of that study (Naughten et al., 2023). Once again, this is in contrast to ocean forcing in ISMIP6 and presumably the similarity in SLR contribution that we find between the two scenarios arises due to our choice of ocean forcing. We have also not included in our uncertainty framework any representation of uncertainties in bedrock geometry which are also likely to contribute to the overall model spread (Sun et al., 2014; Wernecke et al., 2022; Castleman et al., 2022).

Another simplification in our Bayesian approach is to only compare total changes in surface speed and elevation over the entire observational period rather than more detailed temporal variability that has been observed over shorter timescales. While it may initially seem appealing to use as much data as possible, the ability of standalone ice sheet models to capture short-term variability in processes such as calving and ocean melting remains poor and so attempting to match higher frequency temporal variability is arguably a mistake at this time. Such an approach could ascribe a weight that is too high to simulations that are getting the right answer for the wrong reasons and lead to overconfidence in terms of model uncertainty. Along similar lines, since we do not use a coupled ice–ocean model, our ocean forcing does not change in response to changes in cavity geometry, a process that could play an important role in the evolution of melt rates during our transient simulations (De Rydt and Naughten, 2024). That being said, melt rates calculated by the plume model do change in response to the cavity geometry, and so at least part of this feedback is captured by our modelling approach.

Our implementation of the calving laws described in Sect. 2.1.3 is intended to include the possibility of triggering a rapid MICI-type retreat. Increased runoff on ice shelves can induce an increase in the calving rate through the C_Q parameter and parameters related to surface mass balance. If the runoff and calving rate become sufficiently large, a rapid hydrofracture-driven breakup of the ice shelf could occur, exposing tall ice cliffs which would then rapidly retreat. If this point is reached in a simulation, a MICI tipping point is crossed, and rapid retreat occurs without the possibility of stopping until the majority of the drainage basin is ice free. We include this process to compare our results with those of two modelling studies that find that very large SLR contributions of over a metre can occur in 125 years or less (DeConto and Pollard, 2016; DeConto et al., 2021). Although these projections are for all of Antarctica, a large proportion of the modelled mass loss occurs in the ASE region,

and these SLR projections are generally considerably greater than other modelling studies.

In contrast to the studies of DeConto and Pollard (2016) and DeConto et al. (2021), we find no evidence that a MICI-type retreat occurs in any of our simulations; thus, including the cliff failure calving criterion has no meaningful impact on our SLR projections. This is in line with another recent study that found no evidence of a MICI retreat in the 21st century (Morlighem et al., 2024). Due to, at least in part, the absence of a MICI in our simulations, our SLR projections are considerably lower than those of DeConto et al. (2021) for the ASE only, where ~ 0.3 m of SLE mass loss was modelled by 2200 under a $+3\,°C$ global warming scenario. This is in spite of the fact that a very large proportion of ice shelf area has been lost in most of our model simulations by 2100 (Fig. 6a), implying that loss of these ice shelves has very little effect on future ice loss in the region, as previously shown by Gudmundsson et al. (2023). However, it is important to note that our ocean forcing is held quasi-constant after 2100, and so effectively we can only assess whether tipping points occur for the perturbation up to that date; we cannot rule out the possibility that a continued climate forcing would lead to more significant retreat within a similar time frame to DeConto et al. (2021). There are also some differences in our modelling approach compared to DeConto et al. (2021); we do not use the flux formula which is known to fail for buttressed ice streams (Reese et al., 2018), and details of the implementation of calving are necessarily different since the model presented here uses finite elements and the level-set method.

7 Conclusions

In this study, we conduct calibrated ensemble simulations, with forcing derived from a state-of-the-art ocean model and CESM1, yielding SLR projections together with uncertainty estimates for the Amundsen Sea Embayment. We investigate two emissions scenarios, RCP8.5 and Paris 2C, and find the sea level contribution until 2100 to be almost identical for both scenarios (19.3 mm), although with a slightly higher uncertainty for RCP8.5 vs. Paris 2C (± 2.7 vs. ± 2.2 mm). These results are generally slightly lower than previously published SLR projections for the region. We explore how the sea level contribution of the ASE might evolve further into the future by conducting a subset of extended simulations until 2250. Here, the difference between RCP8.5 and Paris 2C is slightly larger (median SLR contributions of 6.0 vs. 7.1 cm, respectively), with Paris 2C on average leading to a slightly higher sea level contribution as a result of the increased snow accumulation found in RCP8.5.

The similarity in SLR contribution between the two emissions scenarios that we investigate is unsurprising, given that the major driver of change in this region is the ocean, and both thermocline depth and temperatures from the ocean

model that we use are very similar in both scenarios (Naughten et al., 2023). For both scenarios, we find the largest fraction of the total uncertainty comes from parametric uncertainty related to model initialisation, although we do not consider structural uncertainty related to our choice of ocean and atmospheric forcing or the ice sheet model itself. After this, parameters related to ice dynamics and basal melting contribute the most to our calculated uncertainty. Other sources of uncertainty, such as calving and internal climate variability, are relatively small in comparison. The importance of initialisation is particularly noteworthy. Although we carefully calibrate the model to best fit measured changes over the observational record, the initial model state plays an important role in our SLR projections, and ongoing work towards new approaches such as time-dependent data assimilation (e.g. Choi et al., 2023) should be a focus for the ice sheet modelling community.

Our model includes the same surface melt and MICI processes as the model of DeConto and Pollard (2016), updated in DeConto et al. (2021). Inclusion of these processes in those studies led to very substantial and rapid mass loss, and these high-end projections have been the focus of much discussion in the following years. In our simulations, as a result of oceanic melting, calving, and increased runoff, the majority of ice shelf area in our simulations is lost by 2100. In spite of this, we find no evidence of a MICI-type retreat, and in fact the loss of these ice shelves seems to have very little effect on the dynamic mass loss in the region. As a result, our SLR projections by the year 2100 are almost an order of magnitude less than those of DeConto et al. (2021) despite using a more pessimistic emissions scenario.

We have not conducted reversibility experiments that would be required to establish whether any of our simulations have crossed a tipping point with regard to ice loss. While we do not see clear evidence of accelerated ice loss towards the end of our simulations, we do not rule out the possibility that Thwaites Glacier or Pine Island Glacier may enter a period of self-enhanced retreat, leading to far larger sea level contribution beyond the timescale considered here.

Appendix A: Plume model description

A1   Plume routing algorithm

Plume theory was originally conceived to model a system in one horizontal dimension, a situation for which the source of a plume is clear. In order to adapt plume theory to work for two-dimensional models, Lazeroms et al. (2018) incorporated an algorithm that determines the origin of the plume melting each location on an ice shelf. This is important, because melting at a point is not just a function of local conditions but also the depth at which the plume originates (z_{GL}) and the average slope between the origin and the location being considered (θ_b). Since the work of Lazeroms et al.

(2018), further refinements to the plume model parameterisation were described in Burgard et al. (2022) but retaining the original routing algorithm. Here, we describe a modified plume routing algorithm that provides a better fit to melt rate observations in the region, followed by a description of the plume model parameterisation itself.

As a first step, we only calculate ocean-driven melting for nodes considered strictly downstream of a grounding line. This means only for nodes belonging to elements for which no other node is upstream of a grounding line to avoid ocean-driven melting leaking upstream of the grounding line. Note that some recent studies suggest ocean water may indeed intrude upstream of the grounding line (e.g. Bradley and Hewitt, 2024; Rignot et al., 2024), and establishing where and when this process may be important is crucial for future ice sheet modelling studies due to their sensitivity to melting in this region. For each of these ice shelf nodes, the routing algorithm must determine an origin location for the plume that then drives melting at that point. A list of candidate nodes is generated from all nodes along the main ice sheet grounding line, meaning that plumes cannot originate from pinning points or inland subglacial lakes. For each list of nodes for which a melt rate must be calculated and candidate plume origin nodes, we create an array of x, y, and z coordinates (where the z coordinate is given by the ice shelf draft b for each node), and we refer to these arrays as N_m and N_{GL}, respectively. The z coordinate of N_m is then offset by a large negative number, and the z coordinate of N_{GL} is multiplied by 2. Finally, we use a standard nearest-neighbour algorithm to find the k-nearest points in N_{GL} for each point in N_m. The effect of the modifications to the z coordinate is to make deeper candidate nodes in N_{GL} appear closer to nodes in N_m and therefore preferred as plume origin nodes. We choose $k = 10$ for our simulations, leading to 10 candidate plume origin nodes for each node in N_m.

We calculate the depth of the origin plume (z_{GL}) for each melting node as the average ice draft of the k plume origin nodes. The average slope (θ_b) at melting node i is calculated as

$$\theta_{b,i} = \sum_{n=1}^{n=k} \frac{b_i - b_{i,n}}{L_{i,n}}, \tag{A1}$$

where b_i is the ice draft at node i, $b_{i,n}$ is the ice draft at the plume origin node n, and $L_{i,n}$ is the horizontal Euclidian distance between node i and plume origin node n.

In addition to the non-local quantities calculated above via the plume routing algorithm, melt rate at each node is a function of local slope (θ_l), ice draft (b), and a depth dependent local temperature and salinity. Vertical profiles of temperature and salinity are generated as a function of ocean model thermocline depth (z_T) and temperature and salinity averaged above (T_s, S_s) and below (T_d, S_d) the thermocline as
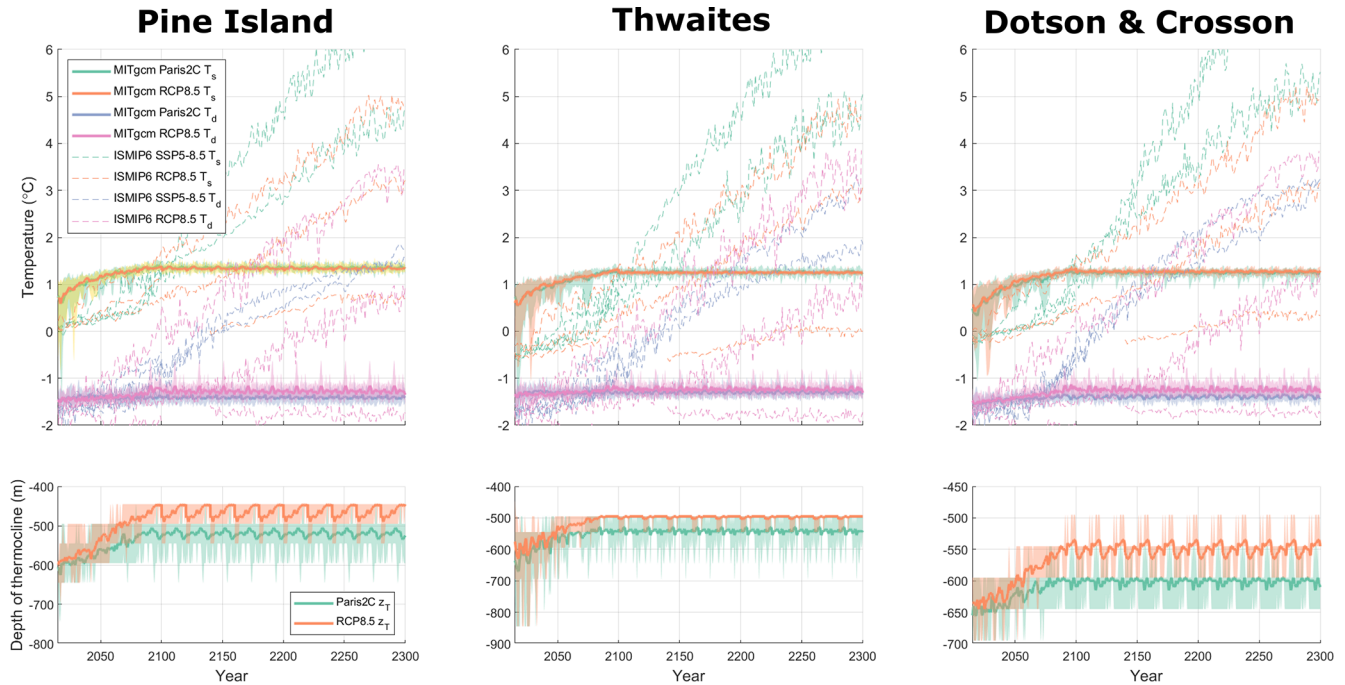
Figure A1. Ocean forcing in the three catchments of our model, extracted from MITgcm simulations for the RCP8.5 and Paris 2C scenarios (Naughten et al., 2023). The shaded area shows the range for each value emerging from different realisations of internal climate variability E_{ID}, while the solid line shows the ensemble mean. Also shown are temperatures extracted from ISMIP6-2300 ocean forcing for different climate models and the SSP5-8.5 and RCP8.5 scenarios (Nowicki and Team, 2024; Seroussi et al., 2024). The temperatures plotted are extracted from all model outputs in the same way, as described in Sect. 2.2. The temperatures (T_s, T_d) above and below the thermocline (z_T) yield a vertical temperature profile as defined in Eq. (A3).

follows:

$$T(z) = \begin{cases} T_s + b\left(\frac{T_d - T_s}{z_T}\right) & \text{if } z \geq z_T, \\ T_d & \text{otherwise,} \end{cases} \quad (A2)$$

resulting in a linear variation in temperature (salinity) from the base of the thermocline ($b = z_T$), where $T = T_d$, to the surface ($b = 0$), where $T = T_s$. These quantities are extracted from MITgcm for each basin (Pine Island, Thwaites, and Dotson/Crosson), and the resulting values are shown in Fig. A1 including the range of internal climate variability (E_{ID}) and emissions scenario. Also shown for comparison are ocean temperature extracted with the same method from the ISMIP6-2300 experiment forcing (Seroussi et al., 2024; Nowicki and Team, 2024) for contrasting scenarios and different climate models (although our method to calculate melt rates from the ocean forcing differs significantly from those experiments).

Finally, the plume parameterisation requires a mean temperature of water entering the cavity ($\bar{T}$), which is calculated as the weighted average of temperatures in the cavity for each node; i.e.

$$\bar{T} = \begin{cases} (T_z + T_d)/2 & \text{if } z_{GL} \geq z_T, \\ T_d & \text{if } b \leq z_T, \\ \frac{(b - z_T)(T_z + T_d)/2 + (z_T - z_{GL})T_d}{b - z_{GL}} & \text{otherwise,} \end{cases} \quad (A3)$$

where T_z is the local temperature at the ice base $T(z = b)$.

A2 Ocean-induced melt rate calculation

To calculate the melt rate, we implement the plume parameterisation of Lazeroms et al. (2018) in a modified form that partially follows Burgard et al. (2022). We briefly summarise the equations as there are some subtle differences in the definition and use of average properties required by our revised plume routing algorithm (A1).

Using the unmodified Lazeroms et al. (2018) parameterisation, the melt rate can be expressed as a function of scaled distance from the grounding line, x':

$$\dot{m} = \left(\frac{C_d^{1/2} \Gamma_{TS}}{L/c}\right) \left[U' f_U(x')\right] \left[\Delta T' f_{\Delta T}(x')\right], \quad (A4)$$

where the velocity and thermal forcing scales are

$$U' = \left[\frac{\beta_T g E_0 \sin \alpha}{\lambda_3 (C_d + E_0 \sin \alpha)}\right]^{\frac{1}{2}}$$

$$\times \left\{\frac{C_d^{1/2} \Gamma_{TS}\left[\frac{\beta_S}{\beta_T}\left(\frac{S_a c}{L}\right) - 1\right]}{C_d^{1/2} \Gamma_{TS}\left[1 - \lambda_1\left(\frac{S_a c}{L}\right)\right] + E_0 \sin \alpha}\right\}^{\frac{1}{2}}$$

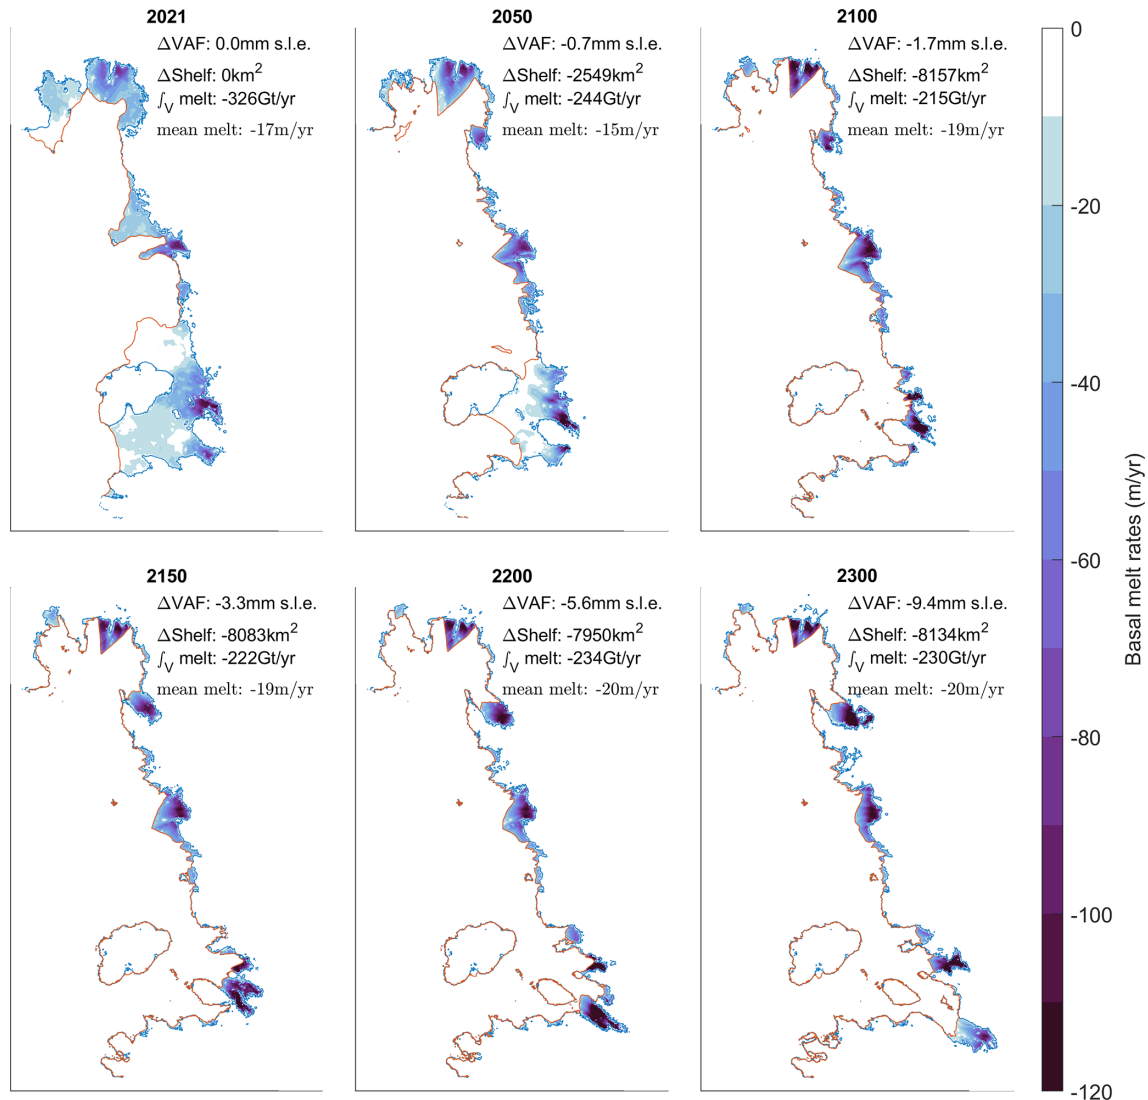$$\times (T_a - T_{f,gl}), \quad (A5)$$

Figure A2. Basal melt rates at six time slices for a single extended Paris 2C simulation, whose final change in volume above flotation by 2300 was closest to the median of all Paris 2C simulations. The solid red line is the calving front position, and the solid blue line is the grounding line. ΔVAF and ΔShelf are change in volume above flotation and change in area of the floating ice shelf with respect to the modelled 2021 values. "$\int_V$ melt" is the integrated basal melt rates across the whole domain, and "mean melt" is total melt divided by ice shelf area.

$$\Delta T' = \left[\frac{E_0 \sin\alpha}{C_d^{1/2} \Gamma_{TS} \left[1 - \lambda_1 \left(\frac{S_a c}{L} \right) \right] + E_0 \sin\alpha} \right]$$
$$\times (T_a - T_{f,\text{gl}}). \tag{A6}$$

The dimensionless velocity and thermal forcing functions are given by

$$f_U(x') = \frac{1}{\sqrt{2}} (1 - x')^{\frac{1}{3}} \left[1 - (1 - x')^{\frac{4}{3}} \right]^{\frac{1}{2}}, \tag{A7}$$

$$f_T(x') = \frac{1}{2} \left[3(1 - x') - \frac{1}{(1 - x')^{\frac{1}{3}}} \right], \tag{A8}$$

and the scaled distance is defined as

$$x' = \frac{\lambda_3(z - z_{\text{gl}})}{(T_a - T_{f,\text{gl}})}$$
$$\times \left\{ 1 + 0.6 \left[\frac{E_0 \sin\alpha}{C_d^{1/2} \Gamma_{TS} \left[1 - \lambda_1 \left(\frac{S_a c}{L} \right) \right] + E_0 \sin\alpha} \right]^{\frac{3}{4}} \right\}^{-1}. \tag{A9}$$

In the original derivation of the above equations, the ambient ocean properties (T_a, S_a) are considered uniform, while the 1 d ice shelf base shallows monotonically with a constant basal slope ($\sin\alpha$) from its maximum depth at the grounding line depth (z_{GL}). Lazeroms et al. (2018) demonstrated that the parameterisation could be extended to non-uniform

slopes simply by using the local value of slope in the above scales, while Burgard et al. (2022) introduced further modifications to account for depth variation in the ambient ocean properties. Here, we follow Lazeroms et al. (2018) in using local slope, except in the definition of the scaled distance (x'), where we make use of the mean slope defined in Appendix A1. In the length scale definition, we further make use of the mean grounding line depth and mean ambient ocean temperature, also defined in Appendix A1. We then follow Burgard et al. (2022) in making use of the mean temperature in the expression for $\Delta T'$ and the local temperature to calculate U'. Note that our plume routing algorithm focuses the origins along deeper parts of the grounding line, unlike the uniform sampling of the cavity inherent in the original algorithm of Lazeroms et al. (2018) that was subsequently used by Burgard et al. (2022). We therefore do not make use of a cavity-wide average of either temperature or slope. In Fig. A2, we show basal melt rates varying in response to forcing and cavity geometry for an example simulation, as calculated by the plume parameterisation described above at various times between 2021 and 2300.

Appendix B: Parameter priors

As a first step to evaluating a posterior probability, we require priors for our selected parameters of the ice sheet, plume, and PDD models (referred to as our three forward sub-models hereafter). One choice would be to use uniform (uninformed) priors or to base our priors on some distribution of the parameter choice in previous studies (e.g. Hill et al., 2021). However, this does not take advantage of the extensive observations that exist in the ASE, particularly since the late 1990s when the satellite record becomes more complete. In the Bayesian framework, these observations can be used to update our probabilistic beliefs. More specifically, we can calculate the posterior probability for our model parameters, given the observations, and then use this posterior as our prior for the model parameters in the future simulations.

B1   Model inversion priors

Our forward model is initialised using inverse methods so that our simulations agree with observed velocity, and surface elevation changes at their starting date. As described in Sect. 2.1.4, this results in five uncertain model parameters related to regularisation of the inverted A and C fields (γ_{aC}, γ_{aA}, γ_{sC}, and γ_{sA}) and the relative weighting of the velocity observations vs. the surface elevation change observations ($\dot{h}_e$). A standard approach when selecting the magnitude of these regularisation parameters is to conduct an L-curve analysis. The motivation behind this is that there exists a tradeoff between the size of the regularisation term (R, Eq. 14) and how well the resulting model agrees with the data (i.e. the misfit term, I, in Eq. 15). By running the inversion for dif-

ferent amounts of regularisation and creating a log–log plot of these two terms, an optimal balance can be determined as the point of greatest curvature in the plot. We conducted this analysis for the five parameters, and the L curves for each one are plotted in Fig. B1.

To define the prior probability distributions for each inverse model parameter, we selected Gaussian distributions whose mean is the optimal point determined by the L-curve analysis (red points in Fig. B1). It is not possible to attribute a variance based on the L-curve plots, and instead we assign a value such that 99.75 % of the distribution falls within 1 order of magnitude either side of this mean. This conservative estimate is based on the fact that the L-curve analysis is an ad hoc strategy to tune the regularisation parameters, and so the most appropriate values may lie some distance away from the points selected by this method.

B2   Plume model priors

We assign prior probabilities to the two plume model parameters based on how well the plume parameterisation matches with melt rates as simulated by the MITgcm ocean model simulations of Naughten et al. (2023). This additional Bayesian inference step uses the same framework as described in Sect. 2.4.3, and the posterior probabilities calculated as described in this section can then be used as prior probabilities for the main model calibration. This ensures that the plume model parameters used in our forward simulations are broadly consistent with the MITgcm model that drives our ocean forcing while also leading to changes in ice velocity and surface elevation in the ASE that are consistent with observations. To do this, we must first define priors for the plume model parameter calibration that follows. We choose normal distributions whose mean is centred around the values proposed by Lazeroms et al. (2018) and with a standard deviation such that 99.7 % of the distribution lies within 1 order of magnitude of the mean.

Since the MITgcm simulations do not include feedbacks on the ice shelf geometry, we run the plume model in a standalone configuration, with the same ice shelf cavity geometry as prescribed in Naughten et al. (2023) and driven by changes in ocean temperature, salinity, and stratification provided by the same model. This is done for all model ensemble members and for both the RCP8.5 and Paris 2C scenarios. As outlined in Sect. 2.4.1 and with the same motivation, we greatly reduce the dimensionality of the melt rate fields that we compare by constructing a SVD of the plume model melt rates. We find that a truncation of $k = 10$ enables us to account for > 95 % of the variance in a large ensemble of standalone plume model simulations spanning the parameter space. In this case, however, the standalone plume model is sufficiently computationally efficient that we do not require a surrogate model. Other than this, the calculation of posterior probabilities follows the same procedure as in Sect. 2.4.3, whereby the likelihood is the difference in modelled melt rates in princi-
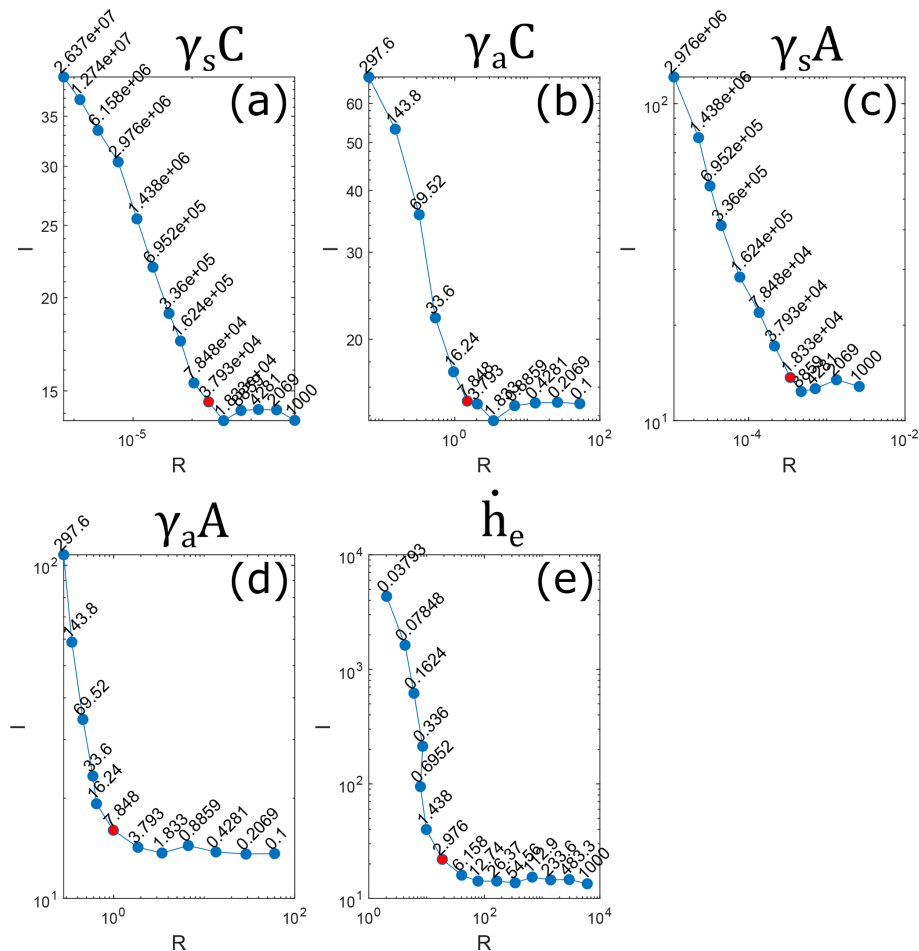
Figure B1. L-curve analysis for model initialisation hyperparameters, showing change in the regularisation term (R, Eq. 14) vs. change in misfit term (I, Eq. 15) for different values of the regularisation parameters (where the value of the parameter is given as a label on each point in the plot.

pal component space (Eq. 24). The maximum a posteriori estimates resulting from this analysis are $\Gamma_{TS} = 1.75 \times 10^{-4}$ and $E_0 = 2.2 \times 10^{-2}$, compared with previously published values of $\Gamma_{TS} = 6.0 \times 10^{-4}$ and $E_0 = 3.6 \times 10^{-2}$ given by Lazeroms et al. (2018).

B3   Ice-flow model priors

We include two ice-flow parameters in our uncertainty analysis; related to the ice rheology (n, Eq. 1) and the basal sliding (m, Eq. 2). Historically, these have typically both been set equal to three, but increasingly this paradigm is being challenged. In the case of Pine Island Glacier, for example, two recent studies have found that using a more strongly non-linear basal sliding parameter is more consistent with observations (Gillet-Chaulet et al., 2016; De Rydt et al., 2021). Similarly, increasing availability and accuracy of remote sensing data along with new methodologies have contributed to the modelling community's adoption of $n = 3$, with some lines of evidence suggesting a greater sensitiv-

ity of ice viscosity to stress is necessary (Millstein et al., 2022; Wang et al., 2025). Despite an increasing amount of research into both of these parameters, no clear consensus has emerged on the best choice in either case. For this reason, we choose uniform priors for both, bracketed by upper and lower bounds of $n = 2$ to $n = 5$ and $m = 2$ to $m = 8$. Note that at the upper end of our prior for m, basal sliding becomes effectively plastic.

Appendix C: Uncalibrated model parameters

Parameters of the PDD model, whose role is to include changes in surface mass balance under warming scenarios, cannot be readily calibrated on present-day ASE observations, where these processes are currently very limited. Similarly, the parameter C_Q controlling the sensitivity of calving rate to surface meltwater is included to introduce the possibility that increased atmospheric warming could lead to accelerated ice shelf collapse, but this process has not been

observed in the region. Without relevant observations or in absence of sensitivity to these parameters under present-day conditions, assigning posterior probabilities is not possible as it is with the other parameters in Table. 1.

The σ_T parameter of the PDD model largely determines the amount of surface meltwater produced as a result of changes in temperature; thus, we can use future simulations of the Regional Atmospheric Model (MAR), that include this process at a point in the future where it becomes relevant, to provide a posterior probability. For this purpose, the PDD model can be decoupled from the ice sheet model, driven by atmospheric forcing and solving for meltwater production with no feedbacks on ice sheet geometry. A recently published modelling study ran the MAR until the year 2100 for the Amundsen sector, driven by CMIP5 anomalies for the RCP8.5 emissions scenario, and found an order of magnitude increase in production of surface meltwater on ice shelves in the region (Donat-Magnin et al., 2021; Marion et al., 2019). By using the same atmospheric forcing as the MAR simulations and comparing meltwater production calculated by the two models in the final 20 years (2080–2100), we can constrain σ_T by finding the value that leads to the best agreement with the state-of-the-art MAR. Defining the decoupled PDD model as our forward model and the MAR meltwater production as our observations, we can calculate posterior probability using the same framework described in Sect. 2.4.3. We repeated this analysis using $\sigma_M = \mathrm{const}$ as a replacement for Eq. (6), but we found that this provided a worse fit to the MAR model data.

A different approach is required for the two remaining parameters, p and C_Q, where analogous model simulations in the region that can inform our parameter choices are not readily available. As an alternative, we use published values in the literature to quantify uncertainty. The parameter p, representing the sensitivity of precipitation to changes in air temperature, is widely used within the ice sheet modelling community, and a recently published study by Nicola et al. (2023) provides an overview of published values from a variety of models. We use these values to define a Gaussian prior with mean and standard deviation calculated directly from the spread of previously published values. The parameter C_Q is unique to the recent study by DeConto et al. (2021), who vary this parameter between 0 and 195 m^{-1} yr^2, and we adopt the same conservative approach, sampling this parameter from a uniform distribution with the same limits.

Appendix D: RNN surrogate

As outlined in Sect. 2.4.2, we require a computationally cheap surrogate model to map from any given θ to a prediction in principal component space ($\tilde{V}^T$) in our Bayesian inference. We choose to use a residual neural network as our surrogate model, due to its simplicity and ability to learn complex non-linear relationships within the training set. Our

goal is that, once training is complete, our surrogate model can predict changes in surface ice speed and elevation over the duration of our chosen observational period for a previously unseen combination of model parameters.

As a first step, we define a flexible architecture consisting of an input layer, a series of N residual blocks, and an output layer. Each residual block consists of a series of D sequences of a fully connected layer with H neurons, a batch normalisation layer, and dropout layer. The output of a residual block is given by the sum of the output of the last dropout layer in the sequence and its input. In all fully connected layers apart from the output layer, we use swish activation functions. The loss function, defined as the Huber loss between modelled and predicted speed and surface elevation change and reprojected from the reduced SVD representation to the model space, is minimised with the Adagrad optimiser.

We train this network using a large ensemble of 5670 model simulations with sampled model parameters Θ and modelled surface ice speed and surface elevation change, reprojected using Eq. (18). We split our m model simulations into training, validation, and test sets with proportions of 70 %, 15 %, and 15 %, respectively. The model is built using TensorFlow, and network optimisation is automated using the Bayesian optimisation option in the Keras Tuner module to find a learning rate l_R, depth D, number of residual blocks N, number of neurons in each layer H, and dropout fraction d that yield the best validation loss. We run this optimisation with various choices for K so that the choice in truncation of the SVD is informed by what results in the best surrogate model. Our final network used $D = 1$, $N = 4$, $H = 115$, $l_R = 0.25$, $d = 0.1$, and $K = 0.9$ (resulting in a truncation of $k = 8$ for surface elevation change and $k = 11$ for change in speed). Figure D1 shows the results of the network evaluated on the test set, at every node in the model domain, showing good agreement between the surrogate model and the forward model with RMSE values of 4.77 m and 35.33 m yr^{-1} for surface elevation and speed change, respectively.
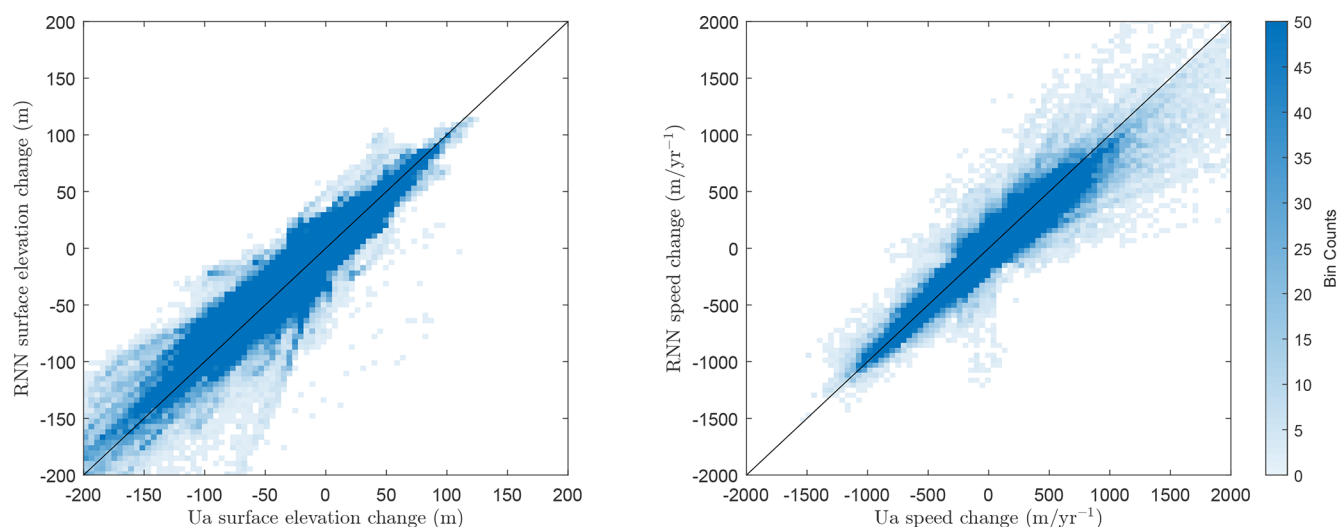
Figure D1. Change in surface elevation and ice surface speed between 1998 and 2021 as predicted by the Úa-*obs* simulations and the RNN surrogate model (Sect. 2.4.2).

Appendix E: LSTM surrogate

We train a surrogate model to make predictions of SLR contribution for a combination of model parameters and forcing ensemble member. Although our primary motivation is to obtain projections for the year 2100, we can make better use of our training simulations by leveraging the inherent time dependency, and so the network is trained to make predictions for every year between 2021 and 2100. For this purpose, a long short-term memory (LSTM) model is appropriate, as it can learn longer time dependencies than a standard recurrent neural network (not to be confused with a residual neural network, as used in Appendix D). An LSTM unit consists of three gates that regulate the flow of information in and out of its current state (analogous to the network's memory). Predictions are made sequentially, based on the input to the LSTM and its current state. The first layer of the network is a sequential input layer that takes normalised input parameters from a training subset of the Úa-*fwd* simulations. These are passed to a series of blocks consisting of an LSTM unit and a dropout layer. Finally, a fully connected layer takes the output of the final block and returns a scalar value representing sea level contribution for each time interval that the model is run for. The network loss was calculated using the Huber loss function on the predicted time series compared to the result calculated from the Úa-*fwd* simulations. As with the surrogate described in Sect. 2.4.2, we optimised the network's hyperparameters to minimise the validation loss, resulting in a network with two blocks of 34 hidden units, a dropout of 0.5, an initial learning rate of 0.007 that was reduced every 10 epochs, and a mini batch size of 32.

The network was trained separately for the two scenarios using 1183 and 1107 simulations for the Paris 2C and RCP8.5 simulations, respectively, of which 20 % were reserved for testing and validation. Figure E1 shows the performance of the two trained surrogate models on the test set of each scenario compared to the outputs from the Úa-*fwd* simulations.

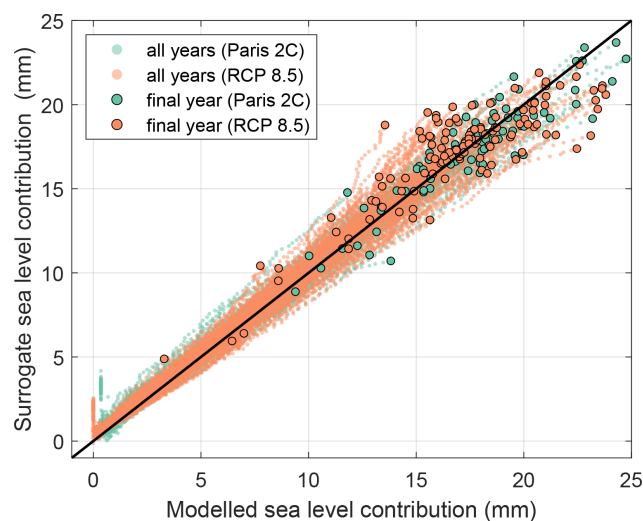**Figure E1.** Cumulative sea level contribution from the ASE for each year, between 2021–2100, as predicted for the test subset of Úa-*fwd* simulations and the LSTM surrogate model. The cumulative sea level contributions for the final year of each simulation, which make up the bulk of our uncertainty quantification and analysis, are highlighted with a black edge.

References

Alevropoulos-Borrill, A. V., Nias, I. J., Payne, A. J., Golledge, N. R., and Bingham, R. J.: Ocean-forced evolution of the Amundsen Sea catchment, West Antarctica, by 2100, The Cryosphere, 14, 1245–1258, https://doi.org/10.5194/tc-14-1245-2020, 2020.

Allison, R. and Dunkley, J.: Comparison of sampling techniques for Bayesian parameter estimation, Mon. Not. R. Astron. Soc., 437, 3918–3928, 2013.

Aschwanden, A., Fahnestock, M. A., Truffer, M., Brinkerhoff, D. J., Hock, R., Khroulev, C., Mottram, R., and Khan, S. A.: Contribution of the Greenland Ice Sheet to sea level over the next millennium, Science Advances, 5, eaav9396, https://doi.org/10.1126/sciadv.aav9396, 2019.

Barnes, J. M. and Gudmundsson, G. H.: The predictive power of ice sheet models and the regional sensitivity of ice loss to basal sliding parameterisations: a case study of Pine Island and Thwaites glaciers, West Antarctica, The Cryosphere, 16, 4291–4304, https://doi.org/10.5194/tc-16-4291-2022, 2022.

Bett, D. T., Bradley, A. T., Williams, C. R., Holland, P. R., Arthern, R. J., and Goldberg, D. N.: Coupled ice–ocean interactions during future retreat of West Antarctic ice streams in the Amundsen Sea sector, The Cryosphere, 18, 2653–2675, https://doi.org/10.5194/tc-18-2653-2024, 2024.

Bevan, S., Cornford, S., Gilbert, L., Otosaka, I., Martin, D., and Surawy-Stepney, T.: Amundsen Sea Embayment icesheet mass-loss predictions to 2050 calibrated using observations of velocity and elevation change, J. Glaciol., 69, 1–11, https://doi.org/10.1017/jog.2023.57, 2023.

Bradley, A. T. and Hewitt, I. J.: Tipping point in icesheet grounding-zone melting due to ocean water intrusion, Nat. Geosci., 17, 631–637, https://doi.org/10.1038/s41561-024-01465-7, 2024.

Brinkerhoff, D., Aschwanden, A., and Fahnestock, M.: Constraining subglacial processes from surface velocity observations using surrogate-based Bayesian inference, J. Glaciol., 67, 385–403, https://doi.org/10.1017/jog.2020.112, 2021.

Burgard, C., Jourdain, N. C., Reese, R., Jenkins, A., and Mathiot, P.: An assessment of basal melt parameterisations for Antarctic ice shelves, The Cryosphere, 16, 4931–4975, https://doi.org/10.5194/tc-16-4931-2022, 2022.

Castleman, B. A., Schlegel, N.-J., Caron, L., Larour, E., and Khazendar, A.: Derivation of bedrock topography measurement requirements for the reduction of uncertainty in ice-sheet model projections of Thwaites Glacier, The Cryosphere, 16, 761–778, https://doi.org/10.5194/tc-16-761-2022, 2022.

Cheng, G., Morlighem, M., and Gudmundsson, G. H.: Numerical stabilization methods for level-set-based ice front migration, Geosci. Model Dev., 17, 6227–6247, https://doi.org/10.5194/gmd-17-6227-2024, 2024.

Choi, Y., Seroussi, H., Morlighem, M., Schlegel, N.-J., and Gardner, A.: Impact of time-dependent data assimilation on ice flow model initialization and projections: a case study of Kjer Glacier, Greenland, The Cryosphere, 17, 5499–5517, https://doi.org/10.5194/tc-17-5499-2023, 2023.

Davison, B. J., Hogg, A. E., Rigby, R., Weldhuijsen, S., van Wessem, J. M., van den Broeke, M. R., Holland, P. R., Selley, H. L., and Dutrieux, P.: Sea level rise from West Antarctic mass loss significantly modified by large snowfall anomalies, Nat. Commun., 14, 1479, https://doi.org/10.1038/s41467-023-36990-3, 2023.

De Rydt, J. and Naughten, K.: Geometric amplification and suppression of ice-shelf basal melt in West Antarctica, The

Cryosphere, 18, 1863–1888, https://doi.org/10.5194/tc-18-1863-2024, 2024.

De Rydt, J., Reese, R., Paolo, F. S., and Gudmundsson, G. H.: Drivers of Pine Island Glacier speed-up between 1996 and 2016, The Cryosphere, 15, 113–132, https://doi.org/10.5194/tc-15-113-2021, 2021.

DeConto, R. M. and Pollard, D.: Contribution of Antarctica to past and future sea-level rise, Nature, 531, 591–597, https://doi.org/10.1038/nature17145, 2016.

DeConto, R. M., Pollard, D., Alley, R. B., Velicogna, I., Gasson, E., Gomez, N., Sadai, S., Condron, A., Gilford, D. M., Ashe, E. L., Kopp, R. E., Li, D., and Dutton, A.: The Paris Climate Agreement and future sea-level rise from Antarctica, Nature, 593, 83–89, https://doi.org/10.1038/s41586-021-03427-0, 2021.

Donat-Magnin, M., Jourdain, N. C., Gallée, H., Amory, C., Kittel, C., Fettweis, X., Wille, J. D., Favier, V., Drira, A., and Agosta, C.: Interannual variability of summer surface mass balance and surface melting in the Amundsen sector, West Antarctica, The Cryosphere, 14, 229–249, https://doi.org/10.5194/tc-14-229-2020, 2020.

Donat-Magnin, M., Jourdain, N. C., Kittel, C., Agosta, C., Amory, C., Gallée, H., Krinner, G., and Chekki, M.: Future surface mass balance and surface melt in the Amundsen sector of the West Antarctic Ice Sheet, The Cryosphere, 15, 571–593, https://doi.org/10.5194/tc-15-571-2021, 2021.

Engwirda, D.: Locally optimal Delaunay-refinement and optimisation-based mesh generation, PhD Thesis, School of Mathematics, University of Sydney, https://core.ac.uk/download/pdf/41240536.pdf (last access: 22 October 24), 2014.

Feldmann, J. and Levermann, A.: Collapse of the West Antarctic Ice Sheet after local destabilization of the Amundsen Basin, P. Natl. Acad. Sci. USA, 112, 14191–14196, https://doi.org/10.1073/pnas.1512482112, 2015.

Fox-Kemper, B., Kopp, R. E., Frölicher, T. L., Cassou, C., Hemer, M., Krinner, G., Kamae, Y., Jourdain, N., Notz, D., Ruiz, L., Berger, S., Muir, L. C., Pearson, B. J., Palmer, M. D., Sweet, W., Mix, A., Hermans, T. J., Jackson, L., Menary, M., Garner, G. G., and Dörr, J.: Ocean, cryosphere and sea level change, in: Climate Change 2021: The Physical Science Basis, Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, edited by: Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., Cambridge University Press, Cambridge, UK and New York, NY, USA, 1211–1362, https://doi.org/10.1017/9781009157896.011, 2021.

Gelman, A. and Rubin, D. B.: Inference from iterative simulation using multiple sequences, Stat. Sci., 7, 457–472, 1992.

Gillet-Chaulet, F., Durand, G., Gagliardini, O., Mosbeux, C., Mouginot, J., Rémy, F., and Ritz, C.: Assimilation of surface velocities acquired between 1996 and 2010 to constrain the form of the basal friction law under Pine Island Glacier, Geophys. Res. Lett., 43, 10311–10321, https://doi.org/10.1002/2016GL069937, 2016.

Glen, J. W.: The creep of polycrystalline ice, P. Roy. Soc. A-Math. Phy., 228, 519–538, 1955.

Goodman, J. and Weare, J.: Ensemble samplers with affine invariance, Commun. App. Math. Com. Sc., 5, 65–80, 2010.

Gudmundsson, G. H.: GHilmarG/UaSource: Ua2019b (Version v2019b), Zenodo [code], https://doi.org/10.5281/zenodo.3706623, 2020.

Gudmundsson, H.: GHilmarG/UaSource: Ua2023b, Zenodo [code], https://doi.org/10.5281/zenodo.10829346, 2024.

Gudmundsson, G. H., Krug, J., Durand, G., Favier, L., and Gagliardini, O.: The stability of grounding lines on retrograde slopes, The Cryosphere, 6, 1497–1505, https://doi.org/10.5194/tc-6-1497-2012, 2012.

Gudmundsson, G. H., Barnes, J. M., Goldberg, D. N., and Morlighem, M.: Limited impact of Thwaites Ice Shelf on future ice loss from Antarctica, Geophys. Res. Lett., 50, e2023GL102880, https://doi.org/10.1029/2023GL102880, 2023.

Haran, T., Bohlander, J., Scambos, T., Painter, T., and Fahnestock, M.: MODIS Mosaic of Antarctica 2008–2009 (MOA2009) Image Map, Version 1, NASA National Snow and Ice Data Center Distributed Active Archive Center [data set], https://doi.org/10.7265/N5KP8037, 2014.

Hersbach, H. B., Berrisford, B., Hirahara, P., Horányi, S., Muñoz Sabater, A., J. Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N.: Complete ERA5 from 1979: Fifth generation of ECMWF atmospheric reanalyses of the global climate, Copernicus Climate Change Service (C3S) Data Store (CDS) [data set], https://doi.org/10.24381/cds.143582cf, 2017.

Hill, E. A., Rosier, S. H. R., Gudmundsson, G. H., and Collins, M.: Quantifying the potential future contribution to global mean sea level from the Filchner–Ronne basin, Antarctica, The Cryosphere, 15, 4675–4702, https://doi.org/10.5194/tc-15-4675-2021, 2021.

Hill, E. A., Gudmundsson, G. H., and Chandler, D. M.: Ocean warming as a trigger for irreversible retreat of the Antarctic ice sheet, Nat. Clim. Change, 14, 1165–1171, https://doi.org/10.1038/s41558-024-02134-8, 2024.

Holland, P., Bracegirdle, T., and Dutrieux, P.: West Antarctic ice loss influenced by internal climate variability and anthropogenic forcing, Nat. Geosci., 12, 718–724, https://doi.org/10.1038/s41561-019-0420-9, 2019.

Holland, P. R., O'Connor, G. K., Bracegirdle, T. J., Dutrieux, P., Naughten, K. A., Steig, E. J., Schneider, D. P., Jenkins, A., and Smith, J. A.: Anthropogenic and internal drivers of wind changes over the Amundsen Sea, West Antarctica, during the 20th and 21st centuries, The Cryosphere, 16, 5085–5105, https://doi.org/10.5194/tc-16-5085-2022, 2022.

Hutter, K.: Theoretical Glaciology: Material Science of Ice and the Mechanics of Glaciers and Ice Sheets, Springer, ISBN 13 9789401511698, 10 9401511691, 1983.

Huybrechts, P. and de Wolde, J.: The dynamic response of the Greenland and Antarctic ice sheets to multiple-century climatic warming, J. Cli-

mate, 12, 2169–2188, https://doi.org/10.1175/1520-0442(1999)012<2169:TDROTG>2.0.CO;2, 1999.

Janssens, I. and Huybrechts, P.: The treatment of melt-water retention in mass-balance parameterizations of the Greenland ice sheet, Ann. Glaciol., 31, 133–140, https://doi.org/10.3189/172756400781819941, 2000.

Jenkins, A.: A one-dimensional model of ice shelf-ocean interaction, J. Geophys. Res.-Oceans, 96, 20671–20677, https://doi.org/10.1029/91JC01842, 1991.

Jenkins, A.: Convection-driven melting near the grounding lines of ice shelves and tidewater glaciers, J. Phys. Oceanogr., 41, 2279–2294, https://doi.org/10.1175/JPO-D-11-03.1, 2011.

Jenkins, A., Shoosmith, D., Dutrieux, P., Jacobs, S., Kim, T. W., Lee, S. H., Ha, H. K., and Stammerjohn, S.: West Antarctic Ice Sheet retreat in the Amundsen Sea driven by decadal oceanic variability, Nat. Geosci., 11, 733–738, https://doi.org/10.1038/s41561-018-0207-4, 2018.

Joughin, I., Tulaczyk, S., Bamber, J. L., Blankenship, D., Holt, J. W., Scambos, T., and Vaughan, D. G.: Basal conditions for Pine Island and Thwaites Glaciers, West Antarctica, determined using satellite and airborne data, J. Glaciol., 55, 245–257, https://doi.org/10.3189/002214309788608705, 2009.

Kay, J. E., Deser, C., Phillips, A., Mai, A., Hannay, C., Strand, G., Arblaster, J. M., Bates, S., Danabasoglu, G., and Edwards, J.: The Community Earth System Model (CESM) large ensemble project: a community resource for studying climate change in the presence of internal climate variability, B. Am. Meteorol. Soc., 96, 1333–1349, 2015.

Lazeroms, W. M. J., Jenkins, A., Gudmundsson, G. H., and van de Wal, R. S. W.: Modelling present-day basal melt rates for Antarctic ice shelves using a parametrization of buoyant meltwater plumes, The Cryosphere, 12, 49–70, https://doi.org/10.5194/tc-12-49-2018, 2018.

Lazeroms, W. M. J., Jenkins, A., Rienstra, S. W., and van de Wal, R. S. W.: An analytical derivation of ice-shelf basal melt based on the dynamics of meltwater plumes, J. Phys. Oceanogr., 49, 917–939, https://doi.org/10.1175/JPO-D-18-0131.1, 2019.

Levermann, A., Winkelmann, R., Albrecht, T., Goelzer, H., Golledge, N. R., Greve, R., Huybrechts, P., Jordan, J., Leguy, G., Martin, D., Morlighem, M., Pattyn, F., Pollard, D., Quiquet, A., Rodehacke, C., Seroussi, H., Sutter, J., Zhang, T., Van Breedam, J., Calov, R., DeConto, R., Dumas, C., Garbe, J., Gudmundsson, G. H., Hoffman, M. J., Humbert, A., Kleiner, T., Lipscomb, W. H., Meinshausen, M., Ng, E., Nowicki, S. M. J., Perego, M., Price, S. F., Saito, F., Schlegel, N.-J., Sun, S., and van de Wal, R. S. W.: Projecting Antarctica's contribution to future sea level rise from basal ice shelf melt using linear response functions of 16 ice sheet models (LARMIP-2), Earth Syst. Dynam., 11, 35–76, https://doi.org/10.5194/esd-11-35-2020, 2020.

MacAyeal, D. R.: Large-scale ice flow over a viscous basal sediment: theory and application to ice stream B, Antarctica, Geophys. Res.-Sol. Ea., 94, 4071–4087, https://doi.org/10.1029/JB094iB04p04071, 1989.

Marelli, S. and Sudret, B.: UQLab: A Framework for Uncertainty Quantification in MATLAB, in: The 2nd International Conference on Vulnerability and Risk Analysis and Management (ICVRAM 2014), University of Liverpool, UK, 13–16 July 2014, https://doi.org/10.1061/9780784413609.257, 2554–2563, 2014.

Marion, D.-M., Nicolas, J., Cécile, A., Hubert, G., Charles, A., Christoph, K., and Xavier, F.: Amundsen Sea MAR simulations forced by ERAinterim, Zenodo [code], https://doi.org/10.5281/zenodo.4308510, 2019.

Millstein, J. D., Minchew, B. M., and Pegler, S. S.: Ice viscosity is more sensitive to stress than commonly assumed, Commun. Earth Environ., 3, 57, https://doi.org/10.1038/s43247-022-00385-x, 2022.

Mitcham, T. and Gudmundsson, G. H.: On the validity of the stress-flow angle as a metric for ice-shelf stability, J. Glaciol., 68, 1–3, https://doi.org/10.1017/jog.2022.25, 2022.

Morlighem, M.: MEaSUREs BedMachine Antarctica, Version 3, NASA National Snow and Ice Data Center Distributed Active Archive Center [data set], https://doi.org/10.5067/FPSU0V1MWUB6, 2022.

Morlighem, M., Goldberg, D., Barnes, J. M., Bassis, J. N., Benn, D. I., Crawford, A. J., Gudmundsson, G. H., and Seroussi, H.: The West Antarctic Ice Sheet may not be vulnerable to marine ice cliff instability during the 21st century, Science Advances, 10, eado7794, https://doi.org/10.1126/sciadv.ado7794, 2024.

Mouginot, J., Rignot, E., and Scheuchl, B.: Mapping of ice motion in Antarctica using interferometric synthetic-aperture radar data, Remote Sens.-Basel, 4, 2753–2767, https://doi.org/10.3390/rs4092753, 2012.

Mouginot, J., Scheuchl, B., and Rignot, E.: MEaSUREs Annual Antarctic Ice Velocity Maps 2005–2017, Version 1, [Version 1.1] Boulder, Colorado USA, NASA National Snow and Ice Data Center Distributed Active Archive Center [data set], https://doi.org/10.5067/9T4EPQXTJYW9 (updated 2022), 2017a.

Mouginot, J., Rignot, E., Scheuchl, B., and Millan, R.: Comprehensive annual ice sheet velocity mapping using Landsat-8, Sentinel-1, and RADARSAT-2 data, Remote Sens.-Basel, 9, 364, https://doi.org/10.3390/rs9040364, 2017b.

Naughten, K. A., Holland, P., and De Rydt, J.: Unavoidable future increase in West Antarctic ice-shelf melting over the twenty-first century, Nat. Clim. Chang., 14, 1222–1228, https://doi.org/10.1038/s41558-023-01818-x, 2023.

Nias, I. J., Cornford, S. L., and Payne, A. J.: Contrasting the modelled sensitivity of the Amundsen Sea Embayment ice streams, J. Glaciol., 62, 552–562, https://doi.org/10.1017/jog.2016.40, 2016.

Nias, I. J., Cornford, S. L., Edwards, T. L., Gourmelen, N., and Payne, A. J.: Assessing uncertainty in the dynamical ice response to ocean warming in the Amundsen Sea Embayment, West Antarctica, Geophys. Res. Lett., 46, 11253–11260, https://doi.org/10.1029/2019GL084941, 2019.

Nicola, L., Notz, D., and Winkelmann, R.: Revisiting temperature sensitivity: how does Antarctic precipitation change with temperature?, The Cryosphere, 17, 2563–2583, https://doi.org/10.5194/tc-17-2563-2023, 2023.

Nilsson, J., Gardner, A. S., and Paolo, F. S.: Elevation change of the Antarctic Ice Sheet: 1985 to 2020, Earth Syst. Sci. Data, 14, 3573–3598, https://doi.org/10.5194/essd-14-3573-2022, 2022.

Nilsson, J., Gardner, A., and Paolo, F. S.: MEaSUREs ITSLIVE Antarctic Grounded Ice Sheet Elevation Change, Version 1, NASA National Snow and Ice Data

Center Distributed Active Archive Center [data set], https://doi.org/10.5067/L3LSVDZS15ZV, 2023.

Nowicki, S. and Team, I.: ISMIP6 23rd Century Forcing Datasets, Zenodo [code], https://doi.org/10.5281/zenodo.13135571, 2024.

Paolo, F. S., Fricker, H. A., and Padman, L.: Volume loss from Antarctic ice shelves is accelerating, Science, 348, 327–331, https://doi.org/10.1126/science.aaa0940, 2015.

Pollard, D., DeConto, R. M., and Alley, R. B.: Potential Antarctic Ice Sheet retreat driven by hydrofracturing and ice cliff failure, Earth Planet. Sc. Lett., 412, 112–121, https://doi.org/10.1016/j.epsl.2014.12.035, 2015.

Pörtner, H.-O., Roberts, D., Masson-Delmotte, V., Zhai, P., Tignor, M., Poloczanska, E., Mintenbeck, K., Alegría, A., Nicolai, M., Okem, A., Petzold, J., Rama, B., and Weyer, N. (Eds.): IPCC Special Report on the Ocean and Cryosphere in a Changing Climate, Cambridge University Press, Cambridge, UK and New York, NY, USA, https://doi.org/10.1017/9781009157964, 2019.

Pritchard, H. D., Ligtenberg, S. R. M., Fricker, H. A., Vaughan, D. G., van den Broeke, M. R., and Padman, L.: Antarctic ice-sheet loss driven by basal melting of ice shelves, Nature, 484, 502–505, https://doi.org/10.1038/nature10968, 2012.

Qi, C. and Goldsby, D. L.: An experimental investigation of the effect of grain size on "dislocation creep" of ice, J. Geophys. Res.-Sol. Ea., 126, e2021JB021824, https://doi.org/10.1029/2021JB021824, 2021.

Reese, R., Winkelmann, R., and Gudmundsson, G. H.: Grounding-line flux formula applied as a flux condition in numerical simulations fails for buttressed Antarctic ice streams, The Cryosphere, 12, 3229–3242, https://doi.org/10.5194/tc-12-3229-2018, 2018.

Reese, R., Garbe, J., Hill, E. A., Urruty, B., Naughten, K. A., Gagliardini, O., Durand, G., Gillet-Chaulet, F., Gudmundsson, G. H., Chandler, D., Langebroek, P. M., and Winkelmann, R.: The stability of present-day Antarctic grounding lines – Part 2: Onset of irreversible retreat of Amundsen Sea glaciers under current climate on centennial timescales cannot be excluded, The Cryosphere, 17, 3761–3783, https://doi.org/10.5194/tc-17-3761-2023, 2023.

Rignot, E., Jacobs, S., Mouginot, J., and Scheuchl, B.: Ice-shelf melting around Antarctica, Science, 341, 266–270, https://doi.org/10.1126/science.1235798, 2013.

Rignot, E., Mouginot, J., Morlighem, M., Seroussi, H., and Scheuchl, B.: Widespread, rapid grounding line retreat of Pine Island, Thwaites, Smith, and Kohler glaciers, West Antarctica, from 1992 to 2011, Geophys. Res. Lett., 41, 3502–3509, https://doi.org/10.1002/2014GL060140, 2014a.

Rignot, E., Mouginot, J., and Scheuchl, B.: MEaSUREs InSAR-Based Ice Velocity of the Amundsen Sea Embayment, Antarctica, Version 1, Boulder, Colorado USA. NASA National Snow and Ice Data Center Distributed Active Archive Center [data set], https://doi.org/10.5067/MEASURES/CRYOSPHERE/nsidc-0545.001, 2014b.

Rignot, E., Ciracì, E., Scheuchl, B., Tolpekin, V., Wollersheim, M., and Dow, C.: Widespread seawater intrusions beneath the grounded ice of Thwaites Glacier, West Antarctica, P. Natl. Acad. Sci. USA, 121, e2404766121, https://doi.org/10.1073/pnas.2404766121, 2024.

Robel, A. A., Seroussi, H., and Roe, G. H.: Marine ice sheet instability amplifies and skews uncertainty in projections of future sea-level rise, P. Natl. Acad. Sci. USA, 116, 14887–14892, https://doi.org/10.1073/pnas.1904822116, 2019.

Rosier, S. H. R.: Uncertainty quantification code for ASE ice sheet model projections, Zenodo [code], https://doi.org/10.5281/zenodo.11922614, 2024.

Rosier, S.: Amundsen Sea Embayment ice sheet model simulations 2021–2300 forced with Paris2C and RCP8.5 climate scenarios, Zenodo [data set], https://doi.org/10.5281/zenodo.14712131, 2025.

Rosier, S. H. R., Reese, R., Donges, J. F., De Rydt, J., Gudmundsson, G. H., and Winkelmann, R.: The tipping points and early warning indicators for Pine Island Glacier, West Antarctica, The Cryosphere, 15, 1501–1516, https://doi.org/10.5194/tc-15-1501-2021, 2021.

Sanderson, B. M., Xu, Y., Tebaldi, C., Wehner, M., O'Neill, B., Jahn, A., Pendergrass, A. G., Lehner, F., Strand, W. G., Lin, L., Knutti, R., and Lamarque, J. F.: Community climate simulations to assess avoided impacts in 1.5 and 2 °C futures, Earth Syst. Dynam., 8, 827–847, https://doi.org/10.5194/esd-8-827-2017, 2017.

Sanderson, B. M., Oleson, K. W., Strand, W. G., Lehner, F., and O'Neill, B. C.: A new ensemble of GCM simulations to assess avoided impacts in a climate mitigation scenario, Climatic Change, 146, 303–318, 2018.

Seroussi, H., Pelle, T., Lipscomb, W. H., Abe-Ouchi, A., Albrecht, T., Alvarez-Solas, J., Asay-Davis, X., Barre, J.-B., Berends, C. J., Bernales, J., Blasco, J., Caillet, J., Chandler, D. M., Coulon, V., Cullather, R., Dumas, C., Galton-Fenzi, B. K., Garbe, J., Gillet-Chaulet, F., Gladstone, R., Goelzer, H., Golledge, N., Greve, R., Gudmundsson, G. H., Han, H. K., Hillebrand, T. R., Hoffman, M. J., Huybrechts, P., Jourdain, N. C., Klose, A. K., Langebroek, P. M., Leguy, G. R., Lowry, D. P., Mathiot, P., Montoya, M., Morlighem, M., Nowicki, S., Pattyn, F., Payne, A. J., Quiquet, A., Reese, R., Robinson, A., Saraste, L., Simon, E. G., Sun, S., Twarog, J. P., Trusel, L. D., Urruty, B., Van Breedam, J., van de Wal, R. S. W., Wang, Y., Zhao, C., and Zwinger, T.: Evolution of the Antarctic ice sheet over the next three centuries from an ISMIP6 model ensemble, Earths Future, 12, e2024EF004561, https://doi.org/10.1029/2024EF004561, 2024.

Sun, S., Cornford, S. L., Liu, Y., and Moore, J. C.: Dynamic response of Antarctic ice shelves to bedrock uncertainty, The Cryosphere, 8, 1561–1576, https://doi.org/10.5194/tc-8-1561-2014, 2014.

Wake, L. and Marshall, S.: Assessment of current methods of positive degree-day calculation using in situ observations from glaciated regions, J. Glaciol., 61, 329–344, https://doi.org/10.3189/2015JoG14J116, 2015.

Wang, Y., Lai, C.-Y., Prior, D. J., and Cowen-Breen, C.: Deep learning the flow law of Antarctic ice shelves, Science, 387, 1219–1224, https://doi.org/10.1126/science.adp3300, 2025.

Wang, Y., Zhang, X., Ning, W., Lazzara, M. A., Ding, M., Reijmer, C. H., Smeets, P. C. J. P., Grigioni, P., Heil, P., Thomas, E. R., Mikolajczyk, D., Welhouse, L. J., Keller, L. M., Zhai, Z., Sun, Y., and Hou, S.: The AntAWS dataset: a compilation of Antarctic automatic weather station observations, Earth Syst. Sci. Data, 15, 411–429, https://doi.org/10.5194/essd-15-411-2023, 2023.

Wernecke, A., Edwards, T. L., Nias, I. J., Holden, P. B., and Edwards, N. R.: Spatial probabilistic calibration of a high-resolution Amundsen Sea Embayment ice sheet model

with satellite altimeter data, The Cryosphere, 14, 1459–1474, https://doi.org/10.5194/tc-14-1459-2020, 2020.

Wernecke, A., Edwards, T. L., Holden, P. B., Edwards, N. R., and Cornford, S. L.: Quantifying the impact of bedrock topography uncertainty in Pine Island Glacier projections for this century, Geophys. Res. Lett., 49, e2021GL096589, https://doi.org/10.1029/2021GL096589, 2022.