



Supplement of

An approach for projecting the timing of abrupt winter Arctic sea ice loss

Camille Hankel and Eli Tziperman

Correspondence to: Camille Hankel (camille_hankel@g.harvard.edu)

The copyright of individual parts of the supplement might differ from the article licence.

Introduction

This document contains the equations of the sea ice model used (Text S1), results from a modified version of the cubic ODE that has a periodic forcing term (Text S2), specific details of the calculations of the diagnostics on the time-changing forcing trajectories (Text S3), a heuristic derivation that explains the result that $\max(dx/d\beta) \approx 1/\mu$ (Text S4), and the calculations of the uncertainties on the predictions of bi-stability width (Text S5). Text S6 provides a method for estimating the likelihood of the existence of bi-stability using results from the main text. Text S7 provides a similar, but alternate approach for estimating the equilibrium hysteresis width from transient simulations alone. In Table S1 we provide the exact parameter values used to configure the sea ice model in the three scenarios described in the main text. Figures S1–S2, we provide extra information on model setup and experimental design, and Figures S3–S5 show March average quantities for all four state variables in the sea ice model from the experiments performed in the main text. Figure S6 shows results from the periodically-forced ODE described in Text S2, and Figures S7–S8 relate to understanding the convergence behavior of $\max(dx/d\beta)$ as a function of ramping rate. Figure S9 helps visualize our method for calculating CO_2^i and CO_2^d , the two edges of rate-dependent hysteresis described in the main text. Figure S10 demonstrates the skill of the alternative method for predicting the equilibrium hysteresis described in Text S7, and finally, Figure S11 demonstrates that the method proposed in the main text also works for predicting the equilibrium structure of a generic ODE (Eqn. 1).

Text S1: Sea ice model equations

The Eisenman model contains four state variables: sea ice volume (V), sea ice area (A), sea ice surface temperature (T_i), and mixed layer temperature (T_{ml}) for a single box representing the entire Arctic. The atmosphere is assumed to be in radiative equilibrium with the surface, and the model is forced by a seasonal cycle of insolation, of poleward heat transport, and of local optical thickness of the atmosphere, which represents cloudiness. The addition of CO_2 is represented by increasing the optical thickness and the midlatitude temperature, which increases poleward heat flux. Melt ponds are parameterized by allowing the ice to melt when the surface temperature reaches 0°C and by modifying the ice albedo when this condition is met. The equations for the model are written below, and

can also be found in the original paper (Eisenman, 2007) that used them. The surface longwave radiation imbalance at the surface is

$$\epsilon(T, T_s) = \frac{2a}{2+N} - \frac{D(T_s)}{2} + b \left(T - T_s + \frac{2T_s}{2+N} \right), \quad (1)$$

where N is the optical depth of the atmosphere and $T_s = AT_i + (1-A)T_{ml}$ is the surface temperature of the box. $D(T_s)$ is the atmospheric poleward heat transport given by

$$D(T_s) = k_D(T_{\text{midlat}} - T_s), \quad (2)$$

where T_{midlat} is a prescribed near-surface atmospheric temperature of the midlatitudes. The net heat flux into the mixed layer is given by

$$F_{ml} = (1-A)(-\epsilon(T_{ml}, T_s) + (1-\alpha_o)F_{sw}) - A\gamma T_{ml} + F_{entr}, \quad (3)$$

where F_{sw} is the shortwave radiation reaching the surface. The mixed layer temperature normally evolves according to this net heat flux, except for when it is at 0° C and cooling, at which point the negative heat flux goes entirely into new ice production and T_{ml} stays at 0° C. This is expressed as:

$$c_{ml}H_{ml} \frac{dT_{ml}}{dt} = \begin{cases} 0 & \text{if } T_{ml} = 0 \text{ and } F_{ml} < 0, \\ F_{ml} & \text{otherwise,} \end{cases} \quad (4)$$

$$F_{ni} = \begin{cases} -F_{ml} & \text{if } T_{ml} = 0 \text{ and } F_{ml} < 0, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where F_{ni} is the new ice production. The ice volume and surface temperature evolution are conditioned on whether or not the surface is melting. With the net ice surface heat flux when $T_i = 0$ written as $F_{net} = -\epsilon(0, T_s) + (1-\alpha_i)F_{sw}$, the equations for ice volume and surface temperature are:

$$L \frac{dV}{dt} = \begin{cases} A(\epsilon(T_i, T_s) - (1-\alpha_{mp})F_{sw} - \gamma T_{ml}) - v_0A & T_i = 0 \text{ and } F_{net} > 0, \\ A(-\frac{kT_i}{h} - \gamma T_{ml}) + F_{ni} - v_0LV & \text{otherwise,} \end{cases} \quad (6)$$

$$\frac{ch}{2} \frac{dT_i}{dt} = \begin{cases} 0 & \text{if } T_i = 0 \text{ and } F_{net} > 0, \\ -\epsilon(T_i, T_s) + (1-\alpha_i)F_{sw} - \frac{kT_i}{h} & \text{otherwise.} \end{cases} \quad (7)$$

Finally, ice area evolution occurs according to:

$$\frac{dA}{dt} = \frac{F_{ni}}{Lh_0} - \frac{A}{2V} \mathcal{R} \left(-\frac{dV}{dt} \right) - v_0A. \quad (8)$$

As mentioned in the main text, in our implementation of the model we also allow the CO₂ concentration to vary inter-annually, by allowing the optical depth (N) and the

midlatitude temperature to be functions of time. They can be written as:

$$N(t) = N_0 + A \sin\left(\frac{2\pi}{1 \text{ yr}}t\right) + \Delta N \times \log_2(\text{CO}_2(t)/280 \text{ ppm}), \quad (9)$$

$$T_{\text{midlat}}(t) = T_0 + B \sin\left(\frac{2\pi}{1 \text{ yr}}t\right) + 3^\circ\text{C} \times \log_2(\text{CO}_2(t)/280 \text{ ppm}), \quad (10)$$

where the time dependence of the final term in each equation is a modification from (Eisenman, 2007). The sin terms in each equation represent the seasonal cycles of the atmospheric optical depth and the midlatitude temperature, respectively, at 280 ppm of CO_2 . In our hysteresis experiments, $\text{CO}_2(t)$ is an exponentially increasing and then exponentially decreasing function of time (see Fig. S2), leading to CO_2 forcings that change linearly in time.

Test S2: Cubic ODE with periodic forcing

In this section, we analyze an ODE that is similar to the one presented in the main text but includes a periodic forcing term, which makes it more analogous to the seasonally forced model of sea ice. The equations for this system are:

$$\frac{dx}{dt} = -x^3 + \delta x + 50 \sin(2\pi t) + \beta(t), \quad \beta(t) = \beta_0 + \mu t, \quad (11)$$

The magnitude of 50 on the sin term is chosen such that the magnitude of the changes in β compared to the amplitude of the periodic forcing is roughly similar to the magnitude of CO_2 changes compared to the amplitude of the seasonal cycle of insolation in the sea ice model. The values of δ needed to configure the three scenarios are slightly different than those for the non-periodic ODE and are as follows: $\delta = 6$ for Scenario 1 (wide bi-stability), $\delta = 4$ for Scenario 2 (narrow bi-stability) and $\delta = 3.4$ for Scenario 3 (no bi-stability). Since the solution x is now oscillatory during time-changing and fixed forcing scenarios, we plot the maximum values of x during each oscillation; this is meant to parallel the plotting of March sea ice (which is approximately the maximum amount of ice during the annual cycle). We only range β from -5 to 5 as this is the range that is needed to sweep across the bifurcations. We see in Figure S6 that the qualitative characteristics of the rate-dependent hysteresis found in an ODE without periodic forcing (main text) also are found in this ODE. There is rate-dependent hysteresis for all three scenarios (panels a-c) and the width of this hysteresis gets wider as we move to faster ramping rates. The addition of the periodic forcing combined with the choice to plot the maximum value of x during each oscillation also generates asymmetry in the increasing and decreasing forcing trajectories (blue vs. red lines). In panel d we plot the maximum gradient of x with respect to β during transient forcing simulations versus the ramping rate, μ . We see that, similar to the result in the main text, the maximum gradient follows a negative power law as a function of μ , with the slope of the power law becoming steeper and approaching a value of -1 as we move from Scenario 3 to Scenario 1. Thus we conclude that

the comparisons we made in the main text between the sea ice model and the simple cubic ODE would also apply if we had chosen to include a periodic forcing in the cubic ODE.

Text S3: Calculating diagnostics on time-changing forcing trajectories

We calculate the maximum rate of change of sea ice volume with respect to CO₂ concentration by taking the maximum change in monthly-averaged March sea ice volume between any two consecutive years during the rate-dependent hysteresis period of the simulations (i.e., ignoring the initial fast decline of sea ice at low CO₂) divided by the change in yearly average CO₂ concentration between those two years. The maximum rate of change of sea ice in time is calculated analogously but divided by the time interval rather than the change in yearly mean CO₂.

To calculate the maximum rate of change of the solution to the cubic ODE with respect to the forcing parameter ($\beta(t)$) or time, we calculate the smoothed absolute change around the two time steps that show the greatest absolute change in x , divided by the change in β or by time. The “smoothed” absolute change is simply the difference between the mean value of x over the five time steps before the largest jump in x and the mean during the five time steps after the jump.

Text S4: Deriving $\max(dx/d\beta) \propto \mu^{-1}$

To understand why $\max(dx/d\beta) \propto \mu^{-1}$ and thus why the maximum rate of change of sea ice also follows a similar negative power law as a function of ramping rate, we first note that,

$$\max\left(\frac{dx}{d\beta}\right) = \max\left(\frac{dx}{dt} \frac{dt}{d\beta}\right) = \frac{1}{\mu} \max\left(\frac{dx}{dt}\right). \quad (12)$$

Thus in Fig. S7 we plot the maximum March dV/dt and the maximum dx/dt respectively as a function of the ramping rate. We can see that as predicted by eqn. 12, the slopes of the power laws in main text Figs. 2a and 2b are those found in Figs. S7 minus 1. In particular, the μ^{-1} rate of convergence in Scenario 1 is recovered when noting that $\max(dx/dt)$ appears to be a constant value as a function of the ramping rate for small enough rates; an unintuitive result that is explained further below. The slopes of the power laws that characterize the convergence of the transient simulations to their equilibrium behavior may also prove a useful tool for inferring the equilibrium, in addition to the method we proposed in the main text.

Next, we provide a heuristic derivation for why $\max(dx/dt)$ approaches a constant in Scenarios 1 and 2 when μ is small. Using Figure S8 for reference, we can see that as x moves from one equilibrium, x_a^* , to the next, x_b^* , during the bifurcation, its maximum rate of change is given by the local minimum of the dx/dt curve, given by c . However, because we are considering a non-autonomous equation with the time-changing forcing $\beta(t)$ that shifts the dx/dt curve down in time, the local minimum c is a function of time,

according to the ramp rate, μ . In other words, the maximum rate of change (dx/dt) during the transition from x_a^* to x_b^* (the value of which is also changing in time) under a given forcing $\beta(t) = \beta_0 + \mu t$ is greater than or equal to c_0 and less than or equal to $c = c_0 + \mu t$, where c_0 is the local minimum of dx/dt exactly at the point of bifurcation and t is the time it takes to complete the transition from one steady-state to the other. In the case where c_0 is large compared to μt , we can make the approximation that $c \approx c_0$ as $\mu \rightarrow 0$. Thus, the $\max(dx/dt) \rightarrow c_0$ as $\mu \rightarrow 0$. Finally, returning to equation 11, we get that $\max(dx/d\beta) = c_0/\mu$ for small μ , recovering the μ^{-1} rate of convergence we estimated empirically.

We do not expect this derivation to hold in cases where c_0 is not large compared to μ . Indeed, in the ODE without a bifurcation where $c_0 = 0$, we see that $\max(\frac{dx}{dt})$ is a positive power of μ , which, when divided by μ according to equation 11, causes $\max(\frac{dx}{d\beta})$ to be a negative power of μ with a magnitude less than 1. We argue that this derivation from a simple ODE provides intuition for the more gradual slopes for the physics-based sea ice model seen in Fig. 2a as we move from a scenario (1) with a wide region of bi-stability to a scenario (3) with no bi-stability or bifurcation. Specifically, in a cubic ODE, the value c_0 exactly corresponds to the width of parameter forcing for which there is bi-stability; while this may not hold exactly for the physical sea ice model, we expect Scenario 1 in the sea ice model to be associated with a large c_0 (a fast maximum rate of change of sea ice in time), Scenario 2 to be associated with a smaller c_0 , and Scenario 3 to be associated with small or zero c_0 . As discussed previously, the larger the magnitude of c_0 the closer the slope of the maximum rate of change in time versus the ramp rate is to zero, which in turn sets the slope of the maximum rate of change in CO_2 versus the ramp rate. Thus, the derivation that $\max(dx/d\beta) \approx c_0\mu^{-1}$ for the cubic ODE with a bifurcation provides insight into the convergence behavior of the transient sea ice simulations for all three scenarios.

Text S5: Calculating uncertainty on predictions of the CO_2 value of tipping points

When fitting curves to CO_2^i and CO_2^d in order to predict CO_2^i and CO_2^d at infinitely slow ramping rates, we noticed that there was some auto-correlation in the residuals (which are the difference between the fitted curves and the actual values of CO_2^i and CO_2^d at all 18 ramping rates). This means that using the covariance matrix of the fitted parameters underestimates the uncertainty on the prediction of CO_2^i and CO_2^d at infinitely slow ramping rates (especially in the case of Scenario 3). To address this issue, we instead use a block-bootstrapping method to calculate the uncertainty on our predictions. We sample with a block size of three, and bootstrap 1000 times, giving us a distribution of estimates of the equilibrium values of CO_2^i and CO_2^d . From these distributions, we can calculate the standard deviation of the predictions, and we use these standard deviations

to plot 95% confidence intervals around the predictions in main text Fig. 3. We perform this block-bootstrapping procedure for each of the predictions that use fewer and fewer simulations to produce all the confidence intervals plotted in Fig. 3b.

Text S6: Calculating the probability of the existence of bi-stability

The block-bootstrapping process described in Text S5 gives us distributions of estimated values of CO_2^i and CO_2^d at infinitely slow ramping rates. Since the width of the “true” hysteresis is the difference between CO_2^i and CO_2^d , we can take the difference of these distributions to estimate the likelihood that the hysteresis width is greater than zero and thus that bi-stability and a tipping point exist. Using only experiments with a ramping rate of 75 years/doubling or faster, in both Scenarios 1 or 2 we find that >95% of the difference distribution (hysteresis width distribution) is greater than zero; in other words, we can say that there is less than a 5% chance that bi-stability does *not* exist. Excluding even more of the ramping experiments (which would be computationally expensive in a GCM) to use only experiments with a ramping rate of 47.5 years/doubling or less, we find that there is an >80% chance that bi-stability exists for both Scenarios 1 and 2. For Scenario 3 (which we know does not have bi-stability), no matter how many experiments we exclude there is never more than 55% chance that bi-stability exists; in fact, when using few experiments the distribution skews towards predicting a *negative* width of hysteresis, an unphysical result that may itself suggest the lack of bi-stability. These statistical tests can be used to estimate the binary existence or non-existence of tipping points, in addition to the method for predicting of the CO_2 location of tipping points presented in the main text.

Text S7: An alternate approach for predicting the equilibrium structure of sea ice

Another approach for inferring the equilibrium structure of sea ice from transient runs only would be to analyze only the difference between CO_2^i and CO_2^d (i.e., the hysteresis width) as a function of the ramping rate instead of the two values separately and fit a curve to see if this width approaches zero (no bi-stability) at infinitely slow ramping rates. We perform this analysis in Figure S10 and find that it successfully identifies bi-stability in the Scenarios. However, unlike the method in Fig. 4, this method does not provide any prediction of the CO_2 value of the tipping point (in Scenarios 1 and 2), so we suggest it is used in addition to, but not in place of, the method in the main text.

References

- Eisenman, I. (2007). Arctic catastrophes in an idealized sea ice model. *2006 Program of Studies: Ice (Geophysical Fluid Dynamics Program)*, 133–161.

	α_i	α_{mp}	α_o	Equilibrium behavior
Scenario 1	.75	.45	.1	wide bi-stability
Scenario 2	.5	.4	.2	narrow bi-stability
Scenario 3	.4	.4	.4	no bi-stability

Table S1. Model configurations leading to wide, narrow, and no bi-stability regimes. The symbols α_i , α_{mp} , and α_o , refer to the albedo of bare ice, melt ponds, and open ocean respectively.




Figure S1. A comparison of our seasonal cycle of insolation to that of Eisenman (2007).




Figure S2. A subset our CO₂ hysteresis experiments (ramp up, hold CO₂ fixed, ramp down) compared to the ramping rate of the RCP8.5 Scenario in CMIP5 (which is nearly the same as the ramping rate of SSP585 in CMIP6). Exponential increases in the concentration of CO₂ in time lead to linear increases of the CO₂ radiative forcing.

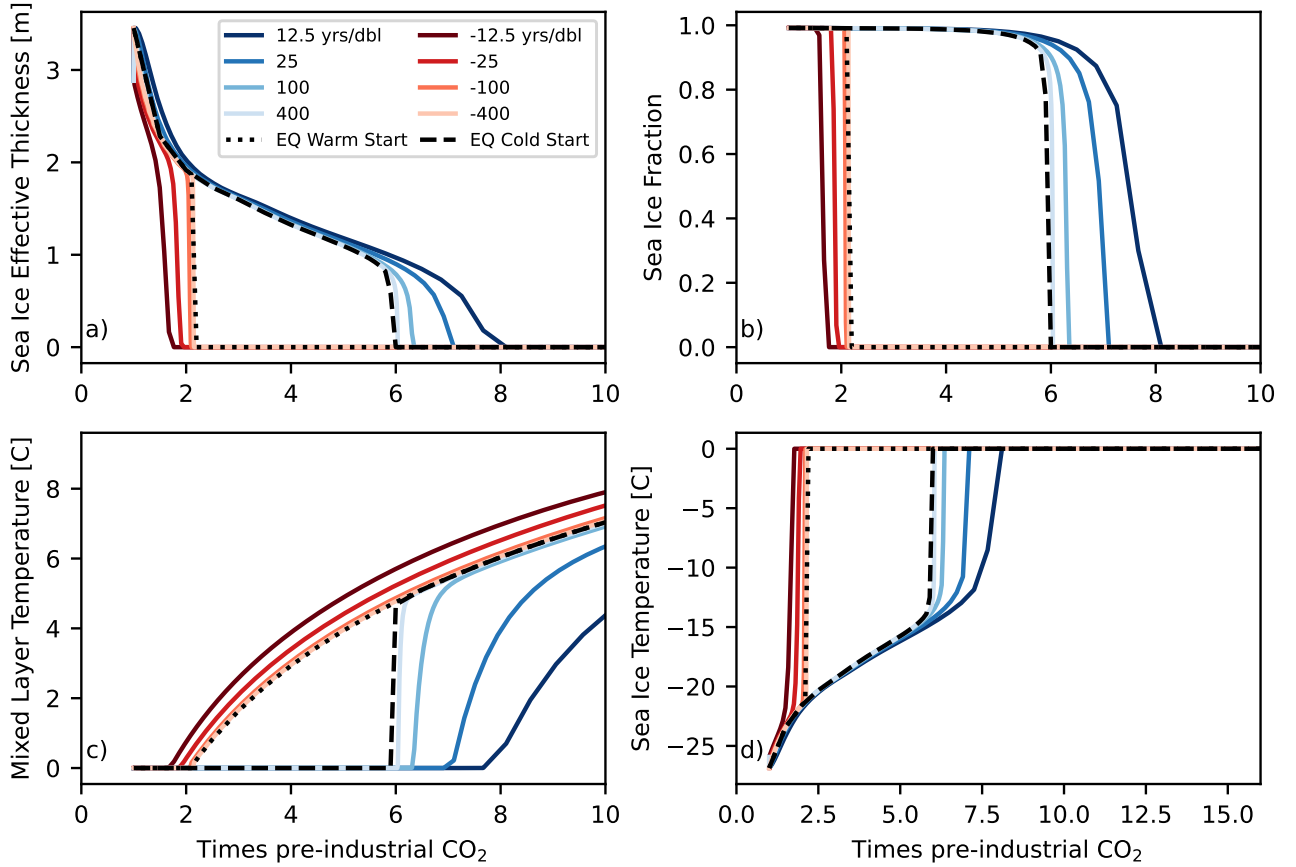


Figure S3. All four state variables of the sea ice model from Scenario 1 (wide bi-stability) runs, shown as their March monthly averages: sea ice effective thickness (a), sea ice fraction (b), mixed layer temperature (c), and sea ice temperature (d). Legend and coloring are the same as in Fig. 1 in the main text.




Figure S4. Same as Figure S3, but for Scenario 2 (narrow bi-stability).




Figure S5. Same as Figure S3, but for Scenario 3 (no bi-stability).




Figure S6. Results from a version of the cubic ODE that is periodically forced (eq. 11). Panels a-c show transient (red and blue colored lines) and fixed-forcing (black lines) simulations for Scenarios 1, 2, and 3 respectively. Panel d shows the maximum gradient of x with respect to the forcing parameter β during transient simulations versus the ramping rate of each simulation for all three scenarios.

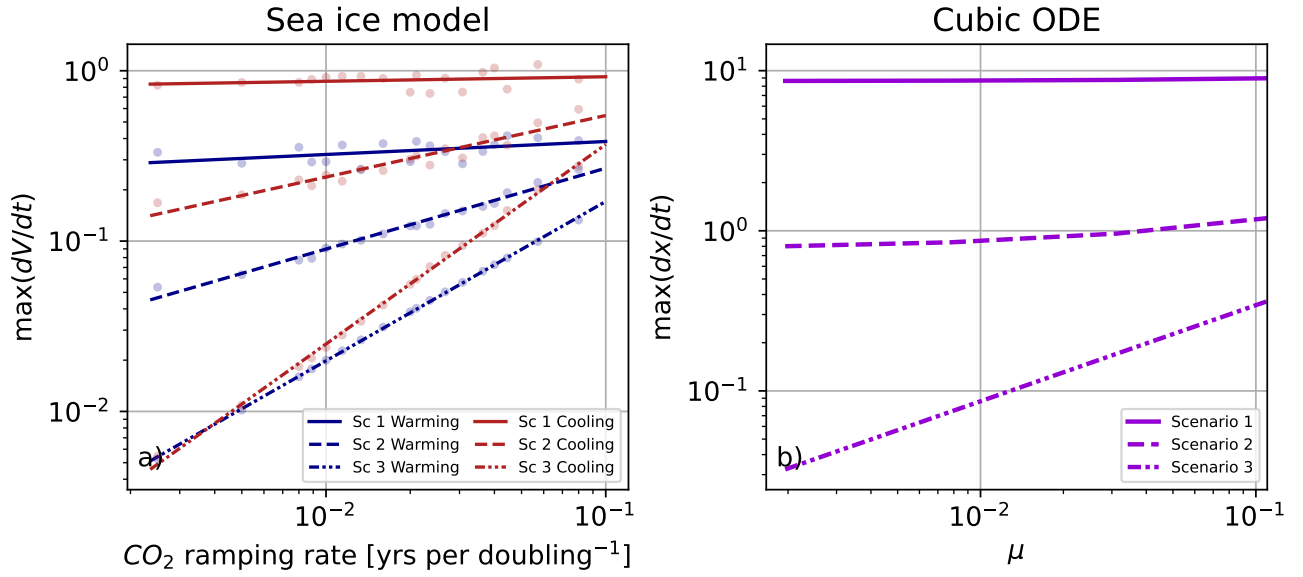


Figure S7. Maximum rate of change of March sea ice effective thickness (a) and maximum rate of change of the variable x from the cubic ODE (b) during time-dependent forcing simulations.




Figure S8. Schematic of the upper limit on $\max(dx/dt)$ during the bifurcation for the equation $dx/dt = -x^3 + 5x - \mu t$ (Scenario 1 for the cubic ODE in the main text). The points x_a^* and x_b^* represent the two stable equilibria before the bifurcation; when the bifurcation happens, x_a^* disappears and the solution must transition to x_b^* . The variable c , represents the upper limit on $\max(dx/dt)$, and c_0 is the value of c at the time of the bifurcation, which is also the width of the true bi-stability.

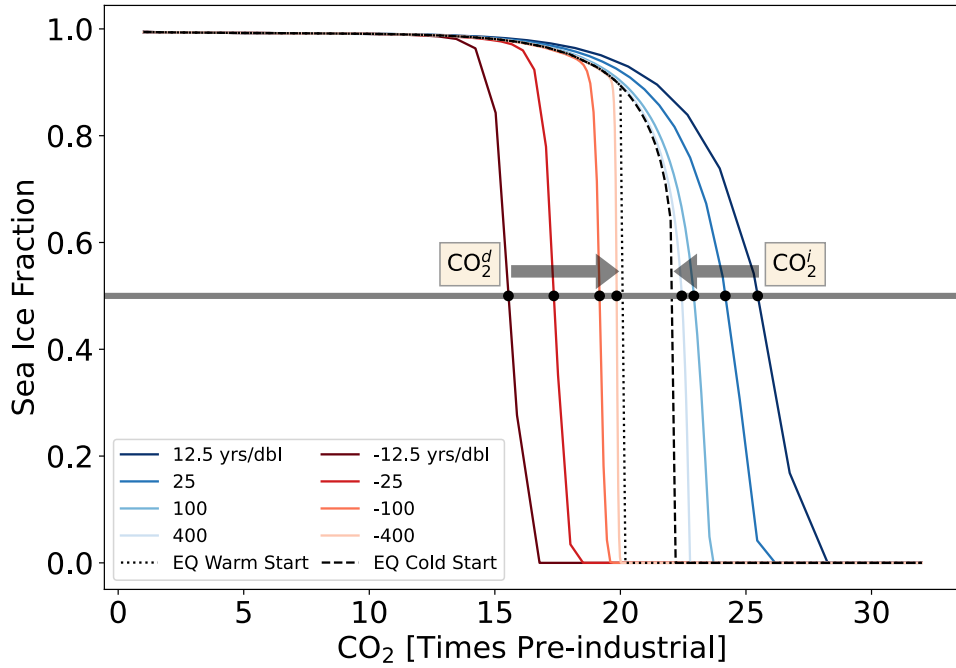


Figure S9. Visualization of how the two edges of the rate-dependent hysteresis (CO_2^i and CO_2^d) are calculated, demonstrated for Scenario 2 only. CO_2^i and CO_2^d are shown for a subset of the ramping rates as the block dots, and are the CO_2 values at which March ice fraction crosses a critical threshold (fraction of .5, shown in gray) along increasing and decreasing CO_2 forcing simulations respectively (see Methods in main text). As indicated by the gray arrows, we can see that for slower and slower ramping rates, CO_2^i and CO_2^d are converging to the width of bi-stability in the equilibrium.

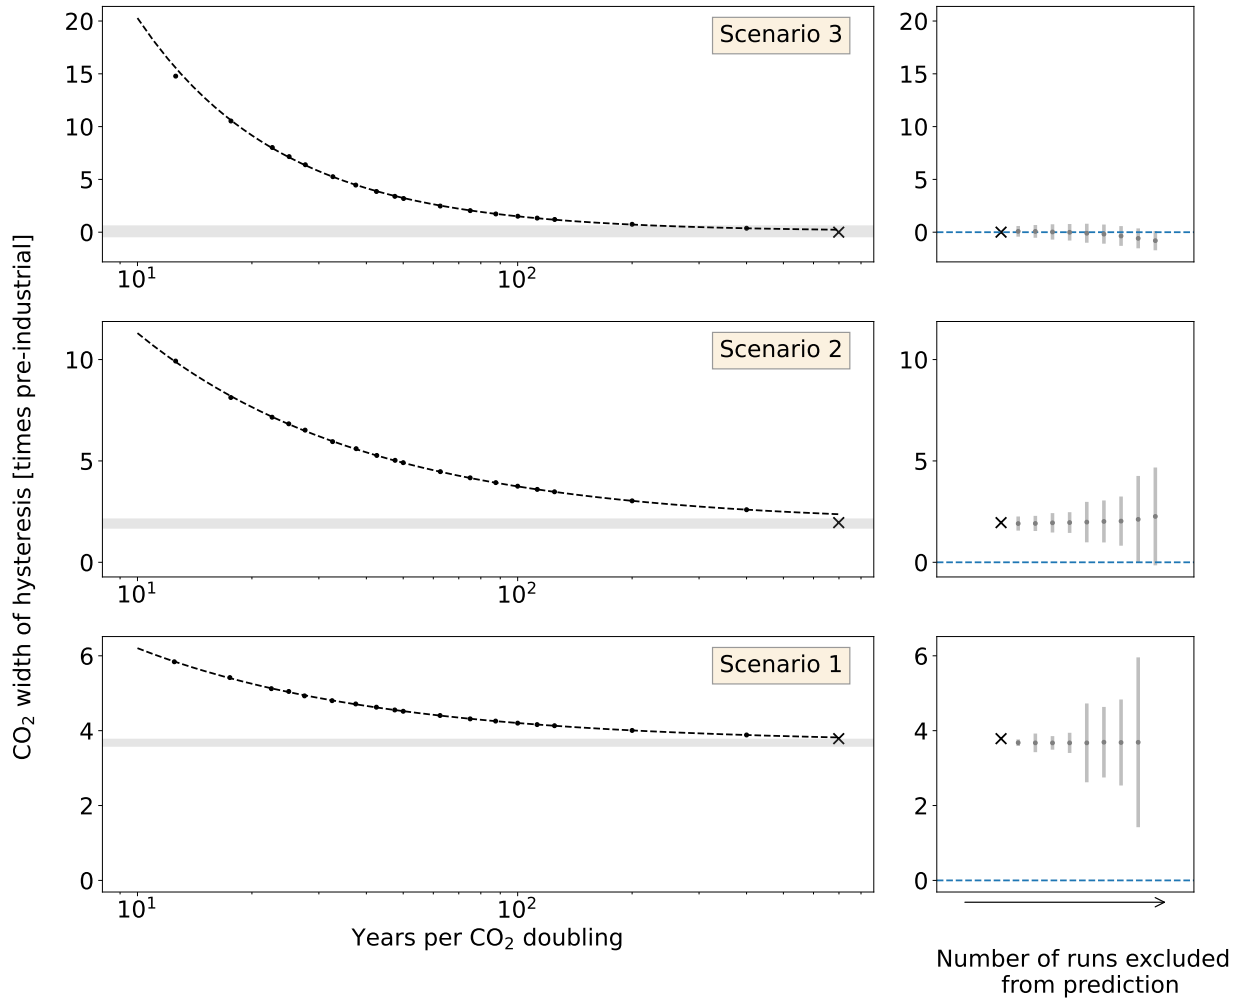


Figure S10. An alternate method for estimating the equilibrium hysteresis width. Here, we fit a single curve to the difference between CO_2^i and CO_2^d , which gives the rate-dependent hysteresis width. We use the same polynomial function fit and block-bootstrapping approach for estimating the uncertainty as done for the main text Figure 4. The left panels show the rate-dependent hysteresis width at different ramping rates as scatter points, the polynomial fit to this width as a dashed line, the true equilibrium hysteresis width as \times 's, and $\pm 2\sigma$ around the predicted equilibrium hysteresis width as gray shading. The right panels show the resulting predictions of equilibrium hysteresis width (gray dots) and the $\pm 2\sigma$ uncertainty around them (gray bars) as the most gradual ramping experiments are sequentially removed from the fitting process. We can see that this method successfully identifies a non-zero equilibrium hysteresis width for Scenarios 1 and 2, even as several runs are excluded from the prediction, and correctly identifies an equilibrium hysteresis width of zero for Scenario 3.

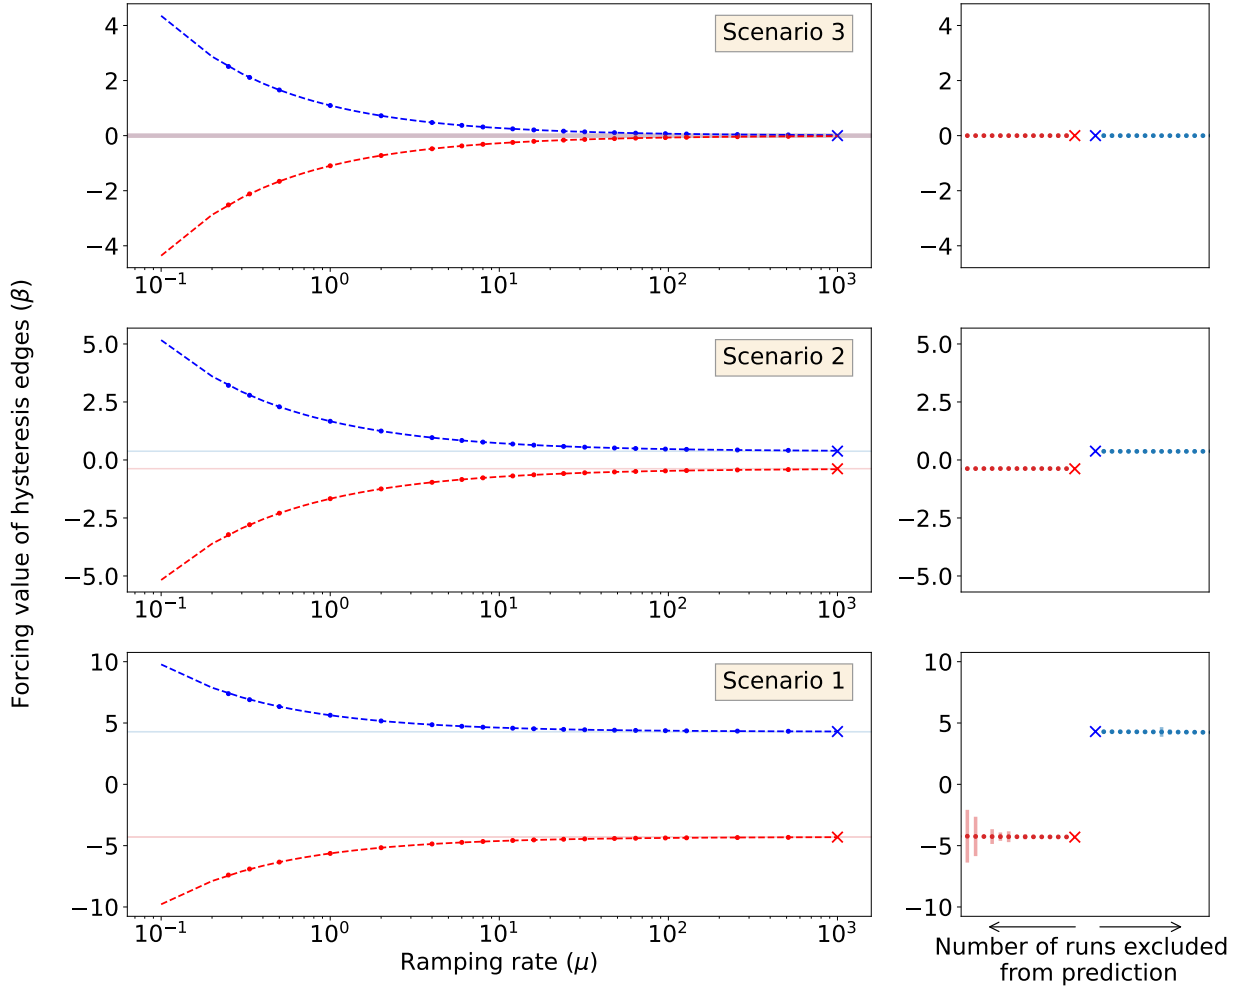


Figure S11. Predicting the equilibrium hysteresis width in a simple cubic ODE. We show that the same method described in the main text and applied in the context of sea ice can be used to predict the equilibrium structure of Eqn. 1 using only transient runs. In the left panels, the blue and red scatter points show the rate-dependent hysteresis width edges, the dashed lines show the curves that fit them, the \times 's show the edges of the true equilibrium hysteresis found from fixed forcing runs, and the shaded bars indicate $\pm 2\sigma$ around the predicted edges of the equilibrium hysteresis. The right panel shows the predictions of the equilibrium hysteresis edges (dots) as the most gradual runs are sequentially excluded from the function fitting, and the uncertainty around these predictions in the shaded bars. In this simple system, we predict the location of the edges of the equilibrium hysteresis with very low uncertainty even when very few transient runs are used to make the prediction.