Emulating Inconsistencies in Stratospheric Aerosol Injection

Jared Farley[1], Douglas G. MacMartin[1], Daniele Visioni[2], Ben Kravitz[3,4]

[1] Department of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY, USA
[2] Department of Earth and Atmospheric Sciences, Cornell University, Ithaca, NY, USA
[3] Department of Earth and Atmospheric Sciences, Indiana University, Bloomington, IN, USA
[4] Atmospheric Sciences and Global Change Division, Pacific Northwest National Laboratory, Richland, WA, USA

Keywords

Stratospheric Aerosol Injection, Emulator, Scenario, Climate Engineering, Climate Intervention

Abstract

Stratospheric Aerosol Injection (SAI) would involve the addition of sulfate aerosols in the stratosphere in order to reflect part of the incoming solar radiation, thereby cooling the climate. Studies trying to explore the impacts of SAI have often focused on idealized scenarios without explicitly introducing what we call “inconsistencies” in a deployment. A concern often discussed is what would happen to the climate system after an abrupt termination of its deployment, whether inadvertent or deliberate. However, there is a much wider range of plausible inconsistencies in deployment than termination that should be evaluated to better understand associated risks. In this work, we simulate a few representative inconsistencies in a pre-existing SAI scenario: an abrupt termination, a decade-long gradual phase-out, and 1-year and 2-year temporary interruptions of deployment. After examining their climate impacts, we use these simulations to train an emulator, and use this to project global mean temperature response for a broader set of inconsistencies in deployment. Our work highlights the capacity of a finite set of explicitly-simulated scenarios that include inconsistencies to inform an emulator that is capable of expanding the space of scenarios that one might want to explore far more quickly and efficiently.

Section 1: Introduction

Stratospheric aerosol injection (SAI) would cool the surface climate through the injection of aerosols or precursor gasses into the stratosphere (Budyko 1977, Crutzen 2006). These aerosols would reflect a small fraction of incoming sunlight, and could be used to offset some fraction of the warming from anthropogenic greenhouse gasses (GHGs). The lifetime of aerosols injected

into the stratosphere is of order of one year (Visioni 2018, Quaglia 2023), so for consistent radiative cooling, consistent injection strategies would be needed.

As there are several forms of lag between the precursor injections and the surface climate response, we define a consistent injection strategy to be one in which the magnitude and frequency of injections allows for a relatively consistent aerosol load to exist in the stratosphere. An inconsistency in SAI, therefore, is any event or change in deployment with sufficient magnitude and duration to cause the aerosol load to significantly stray from the amount needed for a given cooling strategy, whether that modification is intentional or not.

A serious concern regarding any deployment of SAI (and more generally, any form of radiation management) is the complete and sudden stop of all injection, permanently, usually termed a "termination" of SAI deployment. In such a scenario, most of the SAI cooling would disappear within one to two decades (see for example Jones et al 2013). The rate of warming that arises from this inconsistency is linked to the amount of cooling reached in the first place, but it always exceeds the rate of the background GHG-induced warming at the time of termination. This is commonly named "termination shock" (Parker and Irvine 2018). Previous studies have considered the effects of an enhanced rate of warming with respect to economic damages caused by climate change (Goes et al 2011), risks of societal collapse (Baum et al 2013), changes to the Indian monsoon (Bhowmick et al 2021), triggering of tipping points (Ritchie et al 2023), and threat to the biosphere (Trisos et al 2018).

Other inconsistencies beyond this complete and final termination of SAI could occur during a deployment, and it is important to better understand what consequences these might have in order to better assess the risks associated with a deployment of SAI. Further inconsistencies might include 1) a deliberate gradual phase-out of SAI (Irvine et al 2012, MacMartin et al 2014, Keith and MacMartin 2015, Parker and Irvine 2018, and MacMartin et al 2022); 2) a temporary interruption of SAI (Parker and Irvine 2018, MacMartin et al 2022); 3) a volcanic eruption mid-deployment (Quaglia et al 2024); and 4) a sudden phase-in (Hueholt et al 2023, Trisos et al 2018, Irvine et al 2012, Pflüger et al 2024). The difficulty in assessing the effects of inconsistencies lies in the wide range of possible scenarios. Additionally, to date there exist few options between speculation of what an inconsistency could look like and fully simulating it in a computationally expensive climate model.

One option to help explore this expanded scenario space is to use a climate model emulator. A climate model emulator is a reduced-order model that captures some desired features of a climate model's output. These approximations can be used for scenario development and analysis (Kikstra et al 2022, Hansen et al 1984, Caldeira and Myhrvold 2013, MacMartin and Kravitz 2016). Climate emulators have been developed for SRM before; for example, MacMartin and Kravitz (2016) developed and validated emulators for simulating the response to a time-varying

solar reduction – see also MacMartin et al (2014). However, the focus previously was on scenarios where the forcing only varied slowly, and thus it did not need to include considerations over the actual lifetime of stratospheric aerosols. In order to emulate inconsistencies where the timescales of interest are commensurate with the lifetime of aerosols in the stratosphere, the dynamics of the stratospheric aerosols themselves are important, and it is therefore crucial to bridge the gap between the climate response to a forcing and the evolution of the stratospheric load from SAI.

In the following sections, we first describe the climate model we use and the scenarios used herein that were simulated in the full climate model, which include a consistent deployment, a termination, a 10-year phase-out, and 1-year and 2-year interruptions of SAI in Section 2. We discuss these simulation results in Section 3. In Section 4, we use these simulations to train and validate a climate model emulator. We then use this emulator in Section 5 to explore a wider range of inconsistencies. In Section 6, we discuss the impact of inconsistencies and how the scenarios modeled and emulated in this paper fit into the broader scope of scenario assessment. Finally, in Section 7 we conclude with how the emulator can be used to expand and explore the range of SAI scenarios.

Section 2: Climate Models and Scenarios

We use the fully-coupled Community Earth System Model version 2 with the Whole Atmosphere Community Climate Model version 6 (CESM2-WACCM6) (Danabasoglu et al 2020) with middle atmosphere chemistry (Davis et al 2023) and a high top (~140 km). The meridional resolution is 0.95°, the zonal resolution is 1.25°, and the model uses 70 vertical layers extending up to 10^{-5} hPa.

To look at inconsistencies directly, we must first choose a background scenario. We choose the SSP2-4.5-1.5 simulation from MacMartin et al (2022), which considers SAI deployed in 2035 in order to maintain a global mean temperature of 1.5°C above pre-industrial, defined by the mean over 2020-2039. We use the same strategy originally developed in Kravitz et al (2017) and used in GLENS (Tilmes et al 2018), in MacMartin et al (2022), and in ARISE (Richter et al 2022). This strategy uses SO_2 injected at 30°N, 15°N, 15°S, and 30°S. It uses feedback control to maintain global mean temperature at 1.5°C, interhemispheric temperature gradient, and equator-to-pole gradients; all are defined by their 2020-2039 means. SO_2 is injected at every timestep of the model into the gridbox centered at 21.5 km, and the injection rate for a given year is updated at the start of the year.

MacMartin et al (2022) outlines four inconsistencies, illustrated in Figure 1, with the goal of risk evaluation, named SSP2-4.5-1.5-T for a termination of SAI in 2055, SSP2-4.5-1.5-P for a 10-year, linear phase-out of injection beginning in 2055, and SSP2-4.5-1.5-I1 and

SSP2-4.5-1.5-I2 for 1-year and 2-year interruptions beginning in 2055. Three ensemble members are available from MacMartin et al (2022) for the SSP2-4.5 simulation without SAI and for the baseline simulation herein without inconsistencies. Nine new model runs (three ensemble members each for phase-out, 1-year interruption, and 2-year interruption) were performed herein, and two members of the termination model runs from MacMartin et al (2022) were extended. There exist multiple strategies for re-initiation of injection after an interruption. One could inject what one would have injected if the interruption had never happened, inject more than that to meet the temperature target quicker and reject the excess heat accumulated during the interruption, or inject less than that to avoid abrupt changes in forcing. In this paper, the first of these options was chosen for ease of comparison to the consistent deployment, but with an emulator one can easily scale up or down to explore the other options. For the 1-year and 2-year interruptions, the injection rate for each of the run's ensemble members is set to 0 for 1 year and 2 years, respectively, and then restored thereafter to the corresponding year of the SSP2-4.5-1.5 ensemble member from which it branched (without any subsequent feedback to adjust rates). For the phase-out, the injection rate after 2054 is reduced by 10% of 2054 levels per year, such that 0 injection is reached in 2064.

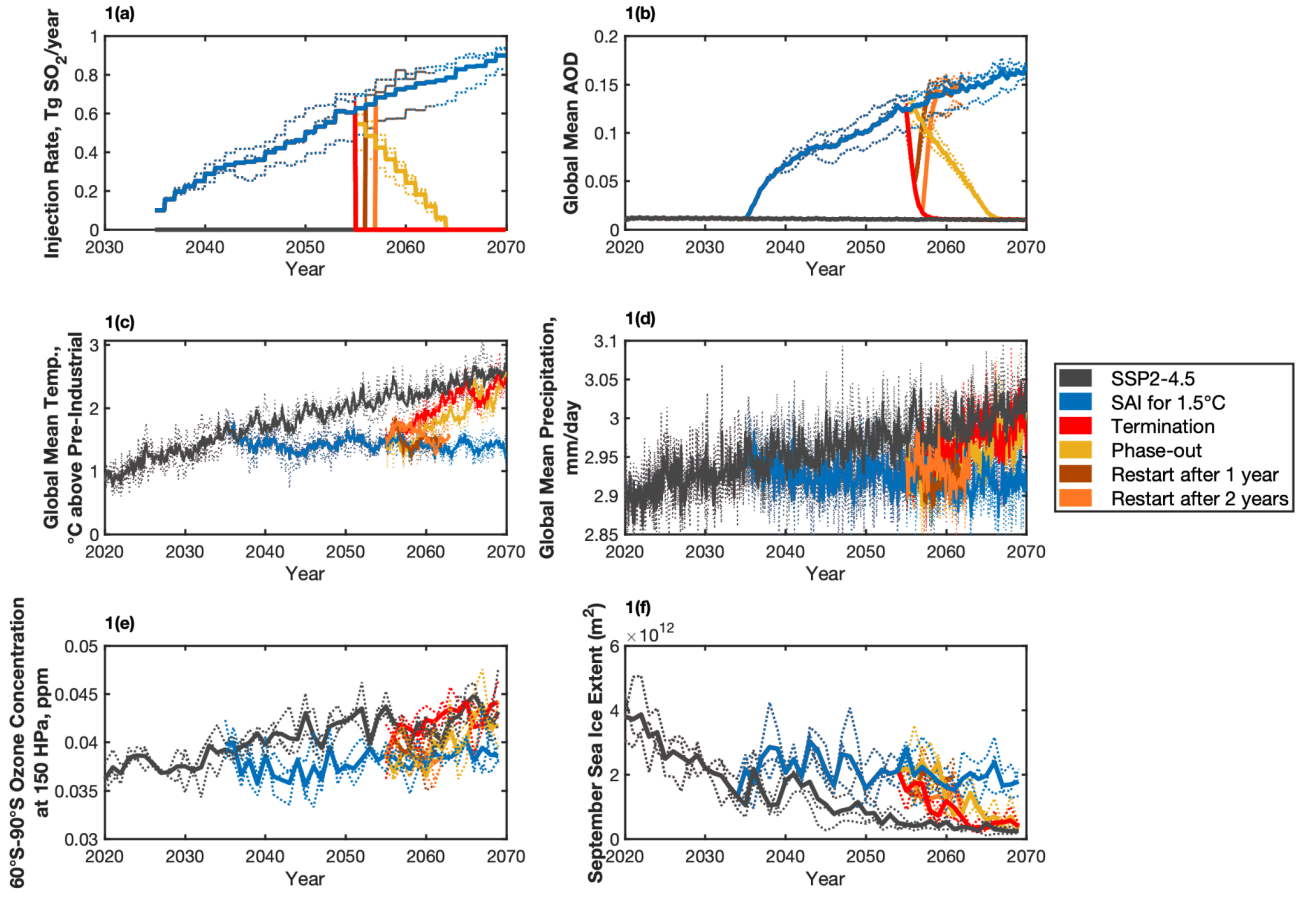

Figure 1: a) Overall SO_2 injection rate (summed across all 4 injection latitudes), b) Global-mean aerosol optical depth (AOD), c) Deseasonalized global-mean surface air temperature, d) Deseasonalized global-mean precipitation,

e) Southern polar ozone concentration, averaged across March, April, and May f) September Sea Ice Extent. Individual ensemble members are shown in dashed lines, the solid line shows the average over the 3-member ensemble.

Section 3: Simulation Results

Figure 1 shows the overall SO_2 injection rate (across all 4 injection latitudes), global mean aerosol optical depth (AOD), global-mean surface air temperature, global-mean precipitation, ozone concentration from 60°S to 90°S, and September Arctic sea ice extent. In panels 1a, 1b, and 1c, one can see that SO_2 injection, AOD, and Temperature changes operate on different timescales. In a termination, SO_2 injection halts immediately, half of the AOD above background is lost within one year, and half of the cooling is lost within five years. One can clearly see the effect of this timescale difference with the 1-year and 2-year interruptions, which lose at maximum approximately 60% and 90% of their AOD but slightly less than 20% and 40% of their cooling, respectively.

The mean rate of temperature rise in SSP2-4.5 (10-year smoothing by convention, see Hueholt et al 2023) in this model from 2035 to 2065 is 0.31°C/decade, with an interannual standard deviation of 0.10°C/decade. Termination peaks at 0.88°C/decade and the 10-year phase-out peaks at 1.0°C/decade (statistically equivalent). Both meet a common definition for “termination shock” – more than 0.2°C/decade higher than under SSP2-4.5 (Parker and Irvine 2018). In contrast, the 1-year and 2-year interruptions peak at 0.26°C/decade and 0.34°C/decade, respectively. The variability in 10-year warming rates arising from natural variability can be estimated as having a standard deviation of 0.11°C/decade by using the consistent-SAI case, where the rate of change is near zero.

Other outputs are correlated with either the rapid response of aerosol concentrations or the slower climate response. In a termination or interruption, Southern Hemisphere polar ozone returns to the baseline level within a year, as it is highly related to aerosol concentrations. In contrast, global mean precipitation and Arctic sea ice minimum extent follow the changes in global mean temperature. Understanding the evolution of global mean aerosol concentrations and temperature during a given inconsistency thus means one can estimate ozone, precipitation, and September sea ice extent as well.

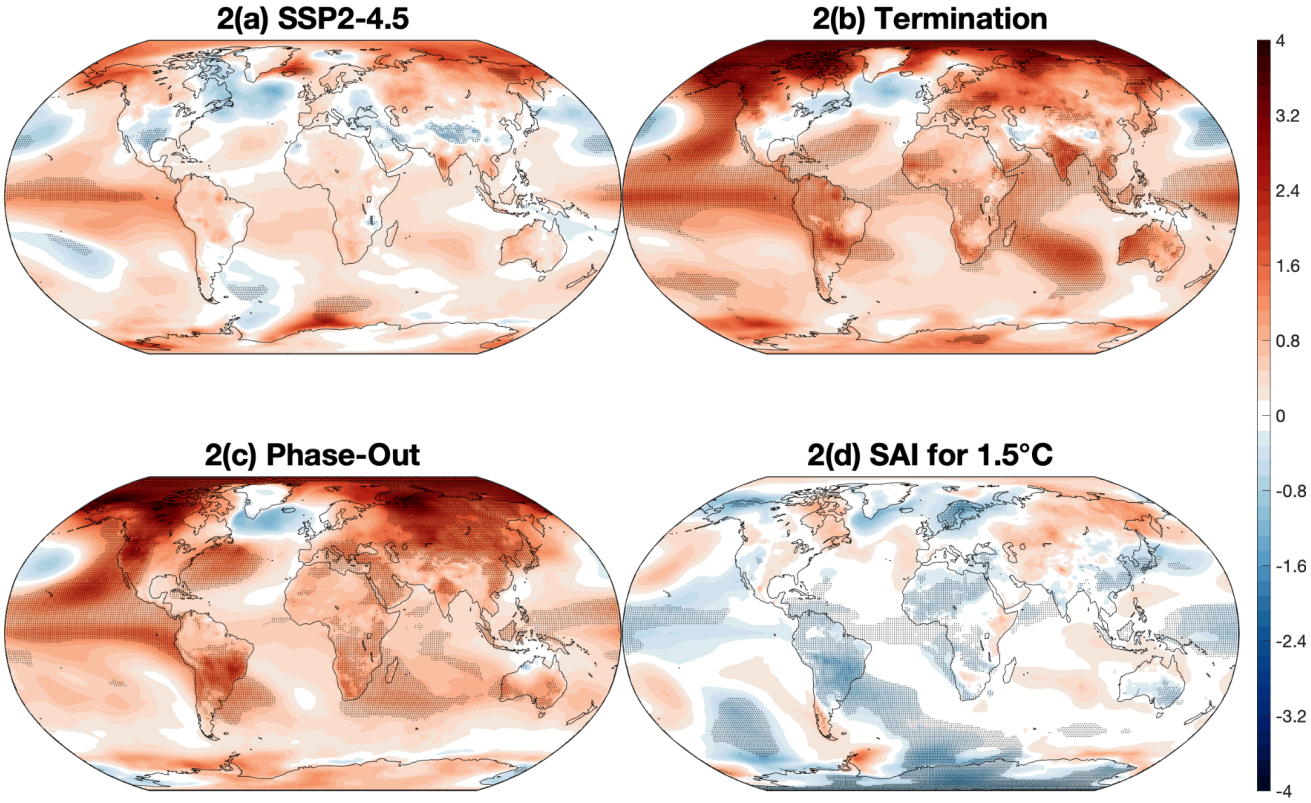

Figure 2: Maps of Heating Rates, in °C per decade. Stippling indicates a grid cell whose slope is more than 2 standard errors away (above or below) from the mean rate of temperature rise for SSP2-4.5 (taken over range of 2020 to 2098). a) SSP2-4.5, b) Termination, c) 10-year Phase-Out, and d) Consistent SAI for 1.5°C

Figure 2 shows that the spatial pattern of warming is similar for the termination, phase-out, and SSP2-4.5 (with the magnitude of the last being significantly smaller). The greatest amount of warming occurs above 60° in latitude, heating up at 1.8°C/decade and matching the pattern of Arctic amplification found in the SSP2-4.5 simulations. While the absolute warming rates are higher at high latitudes, variability is lower in the tropics, so to better understand the ability to detect changes, it is useful to normalize by the standard deviation of interannual variability.

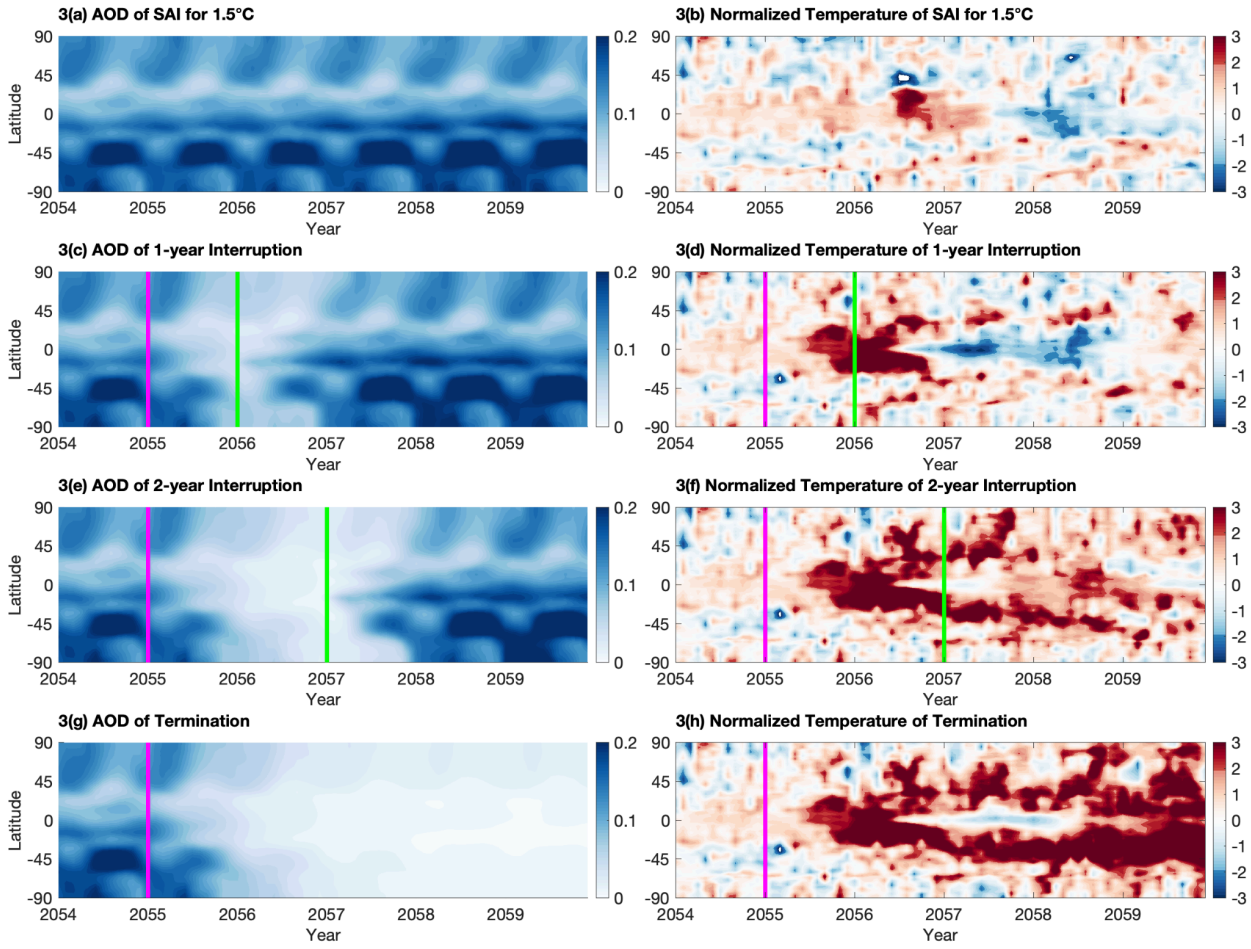

Figure 3: AOD changes and Temperature changes normalized by standard error from a-b) SAI for 1.5°C, c-d) 1-year Interruption, e-f) 2-year Interruption, and g-h) Termination. The magenta line indicates when injection ceases, and the green line indicates when injection is restarted.

Figure 3 shows the evolution of zonal-mean AOD and normalized temperature. When an inconsistency diverges from the SAI induced steady temperatures, the deviation appears first in the tropics, where temperature variability is low. We can see this warming begin to appear beginning at 6 months but fully by the end of the first year (Figures 3d, 3f, and 3h, but also see Figure 4b for the last three months of the first year). Notably, over this timeframe even the 1-year interruption has regions of tropical warming where the signal exceeds three standard deviations of variability. Moreover, it is clear to see from Figures 3c-f that restarting injection does not lead to immediate restoration of either the AOD or the temperature. The disruption persists beyond the restart for an amount of time roughly equivalent to the disruption itself. To a degree, this is due to the strategy of injection chosen for these simulations (injection rate was not increased to compensate for aerosol losses during injection), but due to the multiple forms of lag between injection of SO2 and the climate system response, restoration will never be immediate.

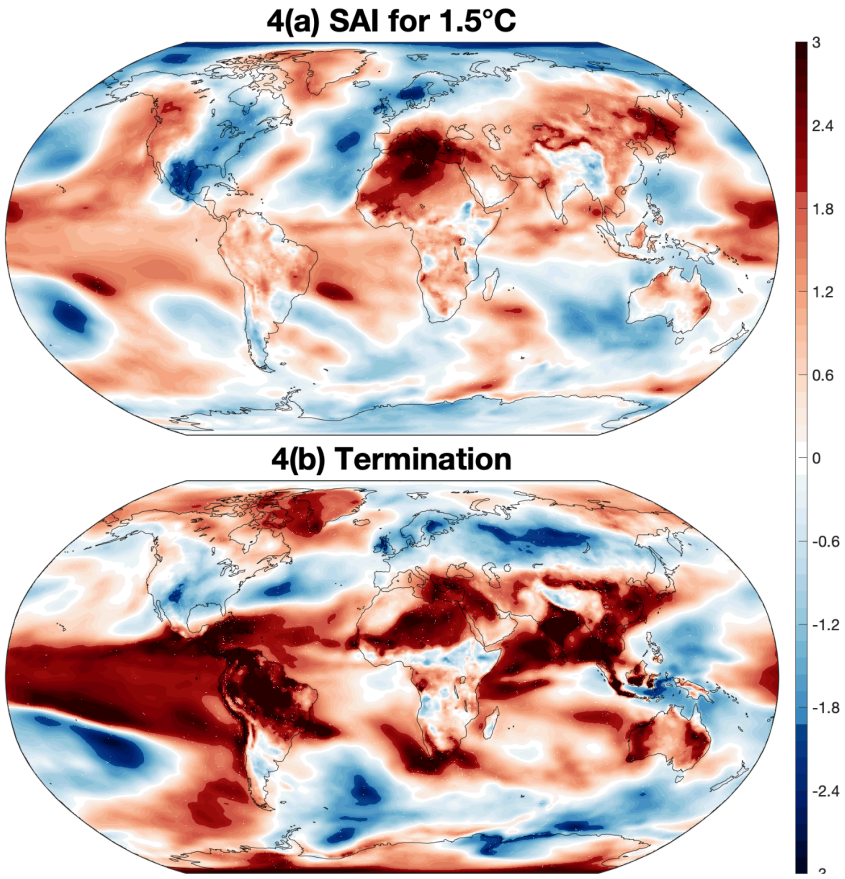

Figure 4: Standard errors of a 3-member ensemble away from the steady-state (2050-2069 consistent SAI) mean. Values are averaged across ensemble members and across months 10-12 of a terminated injection (4b) as compared to the same months of the consistent case (4a), and all are calculated grid cell to grid cell. Within one year, an El-Niño-like pattern of greater than 2 standard deviations appears, with some regions seeing a Z-score above 3. Since the termination, 1-year interruption, and 2-year interruption are identical during the first 12 months, this pattern appears for those interruption cases as well.

Figure 4 shows the number of standard errors away from the mean of the steady state, of temperature averaged across months 10-12 in the consistent (4a) and terminated SAI (4b) (termination, 1-year interruption, and 2-year interruption are identical during the first 12 months). The 3-member ensemble is used, and each grid cell's standard error is calculated by using 3-month deviations away from the steady-state (2050-2069 consistent SAI) mean. There is a strong El-Niño-like pattern of warming, indicating that the temperature rise of a termination or interruption may project onto this mode of variability. It is unclear if the precise timing of when injection halted (January 1 of the shown year) or the state of the El Niño-Southern Oscillation during the year of the inconsistency has an effect on the pattern of warming; regardless, Figure 4 shows that even a 1-year interruption has the potential for strong regional warming (and cooling), even if the signal is masked in the global mean.

As an aside, a La-Niña-like pattern of cooling appears after this warming in many of the ensemble members following an inconsistency (See Supplement 1). It is plausible that a sudden aerosol-removal-induced warming may lead to a La-Niña-like pattern in a similar fashion that a sudden aerosol-induced cooling may lead to an El-Niño-like pattern (Dogar et al 2022), but this relationship should be examined more closely in a dedicated study. In particular, global mean temperatures during the phase-out's La-Niña-like pattern have not greatly diverged from the consistent SAI yet, so these La-Niña-like patterns could simply be expressions of natural variability.

Section 4: Dynamic Emulator

There is a much wider range of possible inconsistencies than can be reasonably simulated in a full climate model; even for the same baseline consistent-SAI scenario considered here there are a range of factors that matter including the timing of any of the events (and hence the amount of cooling at the time), the speed of a phase-out, and the magnitude and duration of an interruption. The effects of these inconsistencies will also depend on the specific baseline SAI scenario from which they deviate.

The goal of the emulator herein is to estimate the forced response of certain features of the Earth system to a wide range of SAI deployment scenarios. Emulators in previous works have estimated the AOD evolution from volcanic SO_2 injection (Aubry et al 2020) and the temperature response from a simple radiative forcing (Hansen et al 1984, Caldeira and Myhrvold 2013, MacMartin et al 2014, MacMartin and Kravitz 2016). The emulator we aim to develop in this paper consists of these same two steps - estimating AOD evolution from SO_2 injections (building off of prior emulators for volcanic forcing), and AOD to radiative forcing in order to obtain the temperature response (building off of earlier SRM emulators). Both use discrete time steps of 1 month, so monthly resolution (for both inconsistency input and emulation output) is possible.

The first step is a modified version of EVA_H from Aubry et al (2020). EVA_H uses eight boxes (three for the middle stratosphere, three for the lower stratosphere, and two for the lowermost extra-tropical stratosphere) with transfer between them to model stratospheric aerosol creation and transport after a volcanic eruption. SO_2 is injected and converted to SO_4. The resultant AOD is then calculated, with a nonlinearity at 10 Tg of SO_4. AOD is calculated using a linear fit below this threshold, and a power law fit above it (see Table 1).

To calibrate the aerosol emulator towards a continuous injection of SO_2 and representation in CESM2, we use the simulations of SSP2-4.5 and SAI for 1.5°C. We subtract the SSP2-4.5 total SO_2 burden, total SO_4 burden, and AOD from their respective SAI cases to get the difference SAI injection rates make on each. This leaves us with three data sets (SO_2, SO_4, AOD for 1.5°C).

The first parameter to set was the timescale for production of SO_4 from SO_2. In this emulator, we set it to be 1 month, as opposed to Aubry et al (2020)'s 7.8 months – although the total sulfur amounts between a volcanic eruption and SAI are comparable, all of the SO_2 from a volcanic eruption is injected at the same time, while for SAI the SO_2 is injected at a much lower rate but continuously. The longer timescale used in Aubry et al (2020) is appropriate when the oxidation of SO_2 is rate-limited by availability of OH; for our SAI simulations, the lower injection rate means that the SO_2 from SAI is able to convert to SO_4 on the order of one month (Quaglia et al 2023).

The second parameter to set was the deposition time scale. Instead of changing the balance between the boxes of EVA_H, we opted for simplicity to scale all deposition time scales by a single parameter, which we found to change with SO_4 burden. At low SO_4 burden, the deposition time constants are increased by 25% from their nominal EVA_H values. Above 7 Tg, this extra scaling is reduced by 1.5% for every Tg above 7 Tg until it reaches 1. This combination was found to best match the evolution of SO_4 in our simulations. Finally, we calibrated the SO_4-to-AOD conversion through minimization of root-mean-square error to be described by the nonlinear power law in rows 4 and 5 of Table 1, rather than the power law described in Aubry et al (2020).

The amount of sulfate aerosols deployed by SAI can change stratospheric dynamics or coagulate in ways that increase or decrease their lifetime (Visioni et al 2018), so being able to capture the nonlinear dynamics associated with the full range of injection magnitudes of a consistent deployment is necessary to have accurate forcing throughout.

All of these adjustments are listed in Table 1.

Row	Adjustment Type	Adjustment Amount
1	Production time scale*	1 month
2	Deposition time scale*	Increased by 25% from nominal EVA_H values. Decreases by 1.5% per Tg for every Tg above 7 Tg until nominal values are reached again. Tropical middle stratosphere: 11.9 -> 9.5 months Middle stratosphere: 3.0 months -> 2.3 months Lower stratosphere: 3.5 months -> 2.7 months Lowermost stratosphere: 4.9 months -> 3.8 months Tropical lower stratosphere: 18.1 -> 14.5 months
3	Mixing time scale	10.7 months

4	SO_4 to AOD Conversion (Linear <= 7 Tg)*	$0.013*(SO_4)$
5	SO_4 to AOD Conversion (Nonlinear > 7 Tg)*	$0.013*(7\ Tg)^{(0.2)}*(SO_4)^{(0.8)}$

Table 1: All parameters of the modified EVA_H (Aubry et al 2020) for usage in this paper's emulator. Parameters marked with a * are modified from the original values or are new.

The second step uses a semi-infinite diffusion model used previously in climate emulators (Hansen et al 1984, Caldeira and Myhrvold 2013, MacMartin et al 2014, MacMartin and Kravitz 2016) to predict global mean temperature response to AOD. The impulse response of the semi-infinite diffusion model takes the form of

$$h_x(t) = \mu_x \left(\frac{1}{\sqrt{\pi\, t\, \tau^*}} - e^{\frac{t}{\tau^*}} erfc\left(\sqrt{\frac{t}{\tau^*}}\right)\right)$$

In this form, t is months since the emulation's start, τ^* is a time scale, and μ_x is a scaling factor capturing the equilibrium temperature response for a particular forcing input x. μ_{ghg} has units °C/ Wm^{-2}, and μ_{sai} has units °C / unit AOD. One can then get the temperature response by convolving the impulse response with the forcing:

$$T(t) = \sum_{j=1}^{t} [h_{ghg}(j) * f(t-j) + h_{sai}(j) * g(t-j)]$$

$$f(t) = 5.35\ \log\left(\frac{CO_2(t)}{CO_2(reference)}\right)\ Wm^{-2}$$

Where t is months since the emulation's start, f is the forcing from CO_2 concentrations (Myhre et al 1998), and g is the AOD above baseline from either simulation output directly (used here in training the temperature-response part of the emulator) or the first step (used in the full emulator).

In order to estimate the parameters, each step of the emulator is trained with three simulations - one with no SAI (SSP2-4.5) to capture the baseline temperature evolution due to GHG, one with consistent SAI (SAI for 1.5°C above pre-industrial) to capture the full magnitude of cooling out to the end of the simulation (and hence estimate μ_{sai}), and one with a large and sudden inconsistency (termination) to capture the response to a step-like input, needed to accurately estimate all of the time-scale parameters. The remaining simulations described in Section 2 are used for validation. The emulator parameters are optimized to minimize RMS error between the emulator output and the simulation output. In the training of the forcing-to-temperature step,

simulation AOD was used (not emulated AOD), such that any errors in the emulation of AOD in the first step were not compensated for in the optimization of the second.

In the second step, we train three parameters: an equilibrium response of temperature to GHG forcing, an equilibrium response of temperature to SAI forcing, and a timescale.

The final emulated global mean temperature is the sum of the contribution from GHG and that from SAI (T = a*GHG + b*SAI). In a similar way, we can emulate global mean precipitation as the as the sum of the contribution from GHG and SAI (P = c*GHG + d*SAI), though the ratio of the coefficients will be different for precipitation than for temperature (the amount of SAI that compensates global mean temperature will overcompensate global mean precipitation; see e.g. Bala et al 2008). The emulator can be trained to model global mean precipitation using the same type of optimization as it had been trained to temperature, simply by holding the timescale constant and estimating equilibrium precipitation responses to GHG and SAI forcing instead. (One does not need to actually hold the timescale constant, but for best comparison with the temperature emulator, we did). In a similar fashion, other climate variables that depend linearly on the GHG forcing and SAI forcing could be emulated even if the ratio of the influence due to GHG and SAI forcing is different.

	SO_2 Standard Deviation (Tg)	SO_4 Standard Deviation (Tg)	AOD Standard Deviation	Temperature Standard Deviation (°C)	Precipitation Standard Deviation (mm/day)
	0.10	0.26	0.0134	0.13°C	0.031
Scenario	SO_2 RMSE (Tg)	SO_4 RMSE (Tg)	AOD RMSE	Temperature RMSE (°C)	Precipitation RMSE (mm/day)
SSP2-4.5	0	0	7.7e-04	0.12	0.020
SAI for 1.5°C	0.08	0.20	0.0029	0.087	0.018
Termination	0.04	0.20	0.0032	0.12	0.019
10-year Phase-Out	0.06	0.39	0.0027	0.12	0.020
1-Year Interruption	0.07	0.28	0.0038	0.12	0.018
2-Year Interruption	0.07	0.29	0.0040	0.10	0.019

Table 2: Root mean square error (RMSE) of monthly values of emulations to their respective monthly ensemble means values from simulation; the cases in the first 3 rows are used in training the emulator while the remaining cases are only used for validation. RMSE for SSP2-4.5 and SAI for 1.5°C calculated from 2035 to 2069. RMSE for inconsistencies calculated from 2055 to the simulation's end date (2069 for termination and phase-out; 2061 and 2062 for 1-year and 2-year interruptions, respectively). Interannual standard deviation calculated from the 2050-2069 period of the SAI for 1.5°C simulation ensemble. Temperature is relatively constant during this period, and other values hold a relatively constant slope.

Figure 5 and Table 2 show the comparison between the emulator and the simulation data for both the training datasets and validation datasets. Atmospheric SO_2 load, atmospheric SO_4 load, global mean AOD, global mean temperature, and global mean precipitation are shown. For comparison, the standard error for temperature is 0.078°C and standard deviation of 0.13°C. If the RMSE of the emulator is small compared to natural variability (deseasonalized), this indicates that the emulator provides a good estimate of the forced response. Scenarios that are members of the emulator's training set (SSP2-4.5, Default, and Termination) can be expected to have a low RMS error, since the optimizer's goal is to minimize it. If the form of the emulator is sufficiently representative of the underlying physics of the climate model, the emulator should be able to estimate the output of runs it was not trained on. To validate that the emulator can estimate the forced response of inconsistencies, we emulate and examine the Phase-Out, 1-Year, and 2-Year simulations and emulations. The RMS error in these cases still remains smaller than interannual natural variability over the span of 2055 to the the end of the simulation (2069 for Phase-Out, 2061 for 1-year interruption, and 2062 for 2-year interruption). Final temperature emulator parameters are as follows: $\tau^* = 260$ months, $\mu_{ghg} = 0.0078°C * / Wm^{-2}$, and $\mu_{sai} = -0.056°C *$ month per unit AOD.

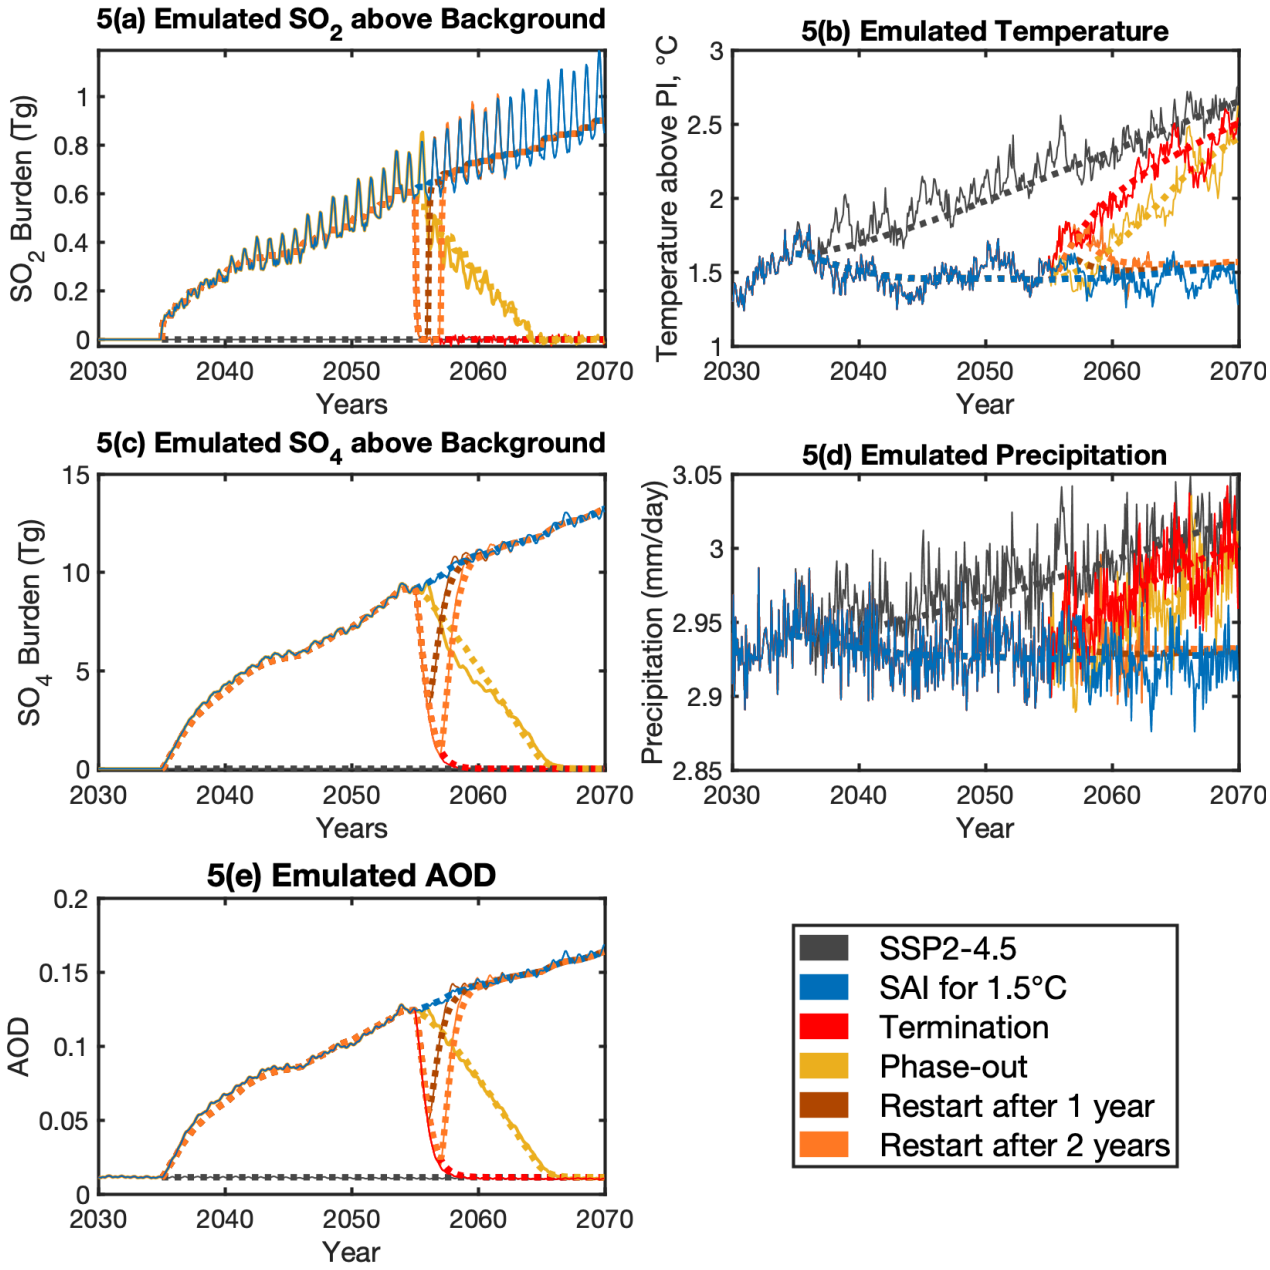

Figure 5: Emulation results compared to simulation data. Solid, thin lines denote simulation monthly ensemble averages, deseasonalized; bold, dotted lines denote emulations. a) Atmospheric SO_2 load (Tg), b) Global Mean Temperature (°C above pre-industrial), c) Atmospheric SO_4 Load (Tg), d) Global Mean Precipitation (mm/day), e) Global Mean AOD

It is also valuable to consider different trajectories for the amount of cooling provided by SAI. In this case we also have simulation data for two cases from MacMartin et al (2022): cooling to the same 1.5°C temperature target but only starting in 2045, or starting in 2035 as in the 1.5°C case here, but cooling down to 0.5°C over 10 years. These can test the ability of the emulator both to capture scenarios where the desired target varies with time and ones with more cooling than used

in the "training" cases. Supplement 2 shows the result of using only the training scenarios of SAI for 1.5°C and SSP2-4.5 for AOD (and use these two as validation) – the AOD below 0.25 is well-captured, but high levels of AOD see the emulation and the simulation diverge. It is possible to include the lower temperature target in the training of the aerosol portion of the emulator: Figure 6 shows the result of this expanded training set. A nonlinearity in cooling rate resulting from high AOD can still be observed: a delayed SAI deployment and a lower temperature target both result in the emulator's estimate being slightly too low. From these results we learn that there still exist some aspects of the climate response that it does not estimate as well as it estimates the inconsistencies modeled for this paper, namely scenarios which involve large and rapid amounts of cooling. (A note: the emulator can be trained on these lower temperature target SAI scenarios, but this results in the cooling of the consistent 1.5°C scenario being underestimated. The timescales of responses are still broadly correct, but one should aim to emulate scenarios near what one trains the emulator on.)

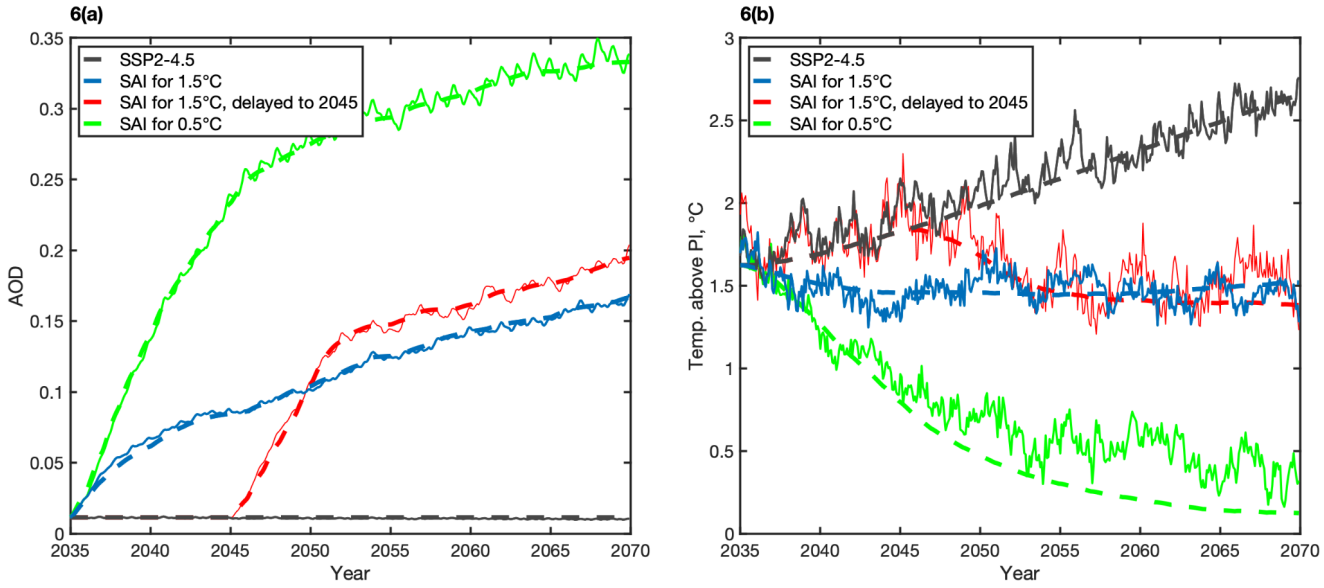

Figure 6: (a) AOD emulation and (b) temperature emulation of other consistent scenarios – a lower temperature target of 0.5°C and a temperature target of 1.5°C, delayed to 2045.

Section 5: Using the Emulator to Expand the Inconsistencies Space

Having developed an emulator, estimated the parameters, and validated that it reasonably approximates the behavior of the full climate model for several additional scenarios, we now apply the emulator to explore a broader range of scenarios, focusing on inconsistencies in deployment.

Alternate versions of the inconsistencies modeled in this paper are a reasonable place to start – terminations in different years, phase-outs of different rates, and interruptions of different lengths. Figure 7a-6f shows emulated temperatures and corresponding temperature rates of change for a range of these inconsistencies, under the same baseline scenario (SSP2-4.5 with SAI to maintain 1.5°C). By emulating these alternate scenarios, we can both glean more information about how the system behaves and ascertain at what value of parameter a certain threshold is reached.

All temperature rates of change are calculated by finding the slope of the least-squares line of best fit through a 10-year span of emulated points. From the alternate terminations (Figure 7a, 6b) we can see that the rate of temperature rise increases as the amount of SAI does, and that the maximum rate of temperature rise is at the start of termination, delayed about a month due to the lag of the last consistent injection converting to aerosol. A termination after 30 years of SAI causes temperature rates to surpass 0.2°C/decade over SSP2-4.5 for about a decade, but even the termination after a 10-year deployment causes the 10-year temperature rate to momentarily surpass this threshold proposed by Parker and Irvine (2018). All terminations trend to asymptotic convergence with the background, as expected.

The alternate phase-outs (Figure 7c, 6d) show something similar. As one might expect, phasing out gradually does lower the rate of temperature rise. The peak of the 10-year phase-out is only slightly below that of a termination in the same year, and the peak of the 50-year phase-out stays below 0.4°C/decade. All surpass the background rate of temperature rise by 2065, one decade after the phase-out began. Examining Figure 7c and comparing the inconsistency's start at 2055 to its eventual trajectory reveals an underlying feature of phasing out during rising GHG concentrations – much of the phase-out's trajectory is in pursuit of a climbing background temperature. Even the slowest phase-out arrives at the background temperature of its initiation within 35 years. The fastest phase-out reaches this temperature by 10 years. The rest of the phase-out's duration beyond this point consists of catching up to where the SSP2-4.5 temperature profile has moved since the phase-out began.

The series of emulated interruptions (Figure 7e, 6f) are in a way snapshots of the termination in 2055; the warming sections of each interruption are the same as termination and diverge when injection is restarted. However, because the rates of temperature change are taken over the span of 10 years, they capture both the rise in temperature and the return to the target. As a result, interruptions of 2 years or less do not surpass the background rate of temperature rise. A 1-month interruption, similar to the "technological failure" briefly explored in Jackson et al (2015), is predicted here to lead to a warming of less than 0.03°C. The 1-year interruption leads to a peak warming of 0.17°C, and every year thereafter warms by a smaller magnitude (as one would expect from asymptotic behavior).

The key takeaway from Figure 7 is the importance of timescales - termination takes somewhere from 10 to 20 years to complete, depending on where one considers it finished. A gradual phase-out over a time less than or comparable to this duration has only a limited effect on warming rates. The beginning years of termination contain much of its warming, but restarting after two years keeps the 10-year-average rate of warming below that of SSP2-4.5, while a 1-month interruption only deviates from the consistent deployment by a very small amount.

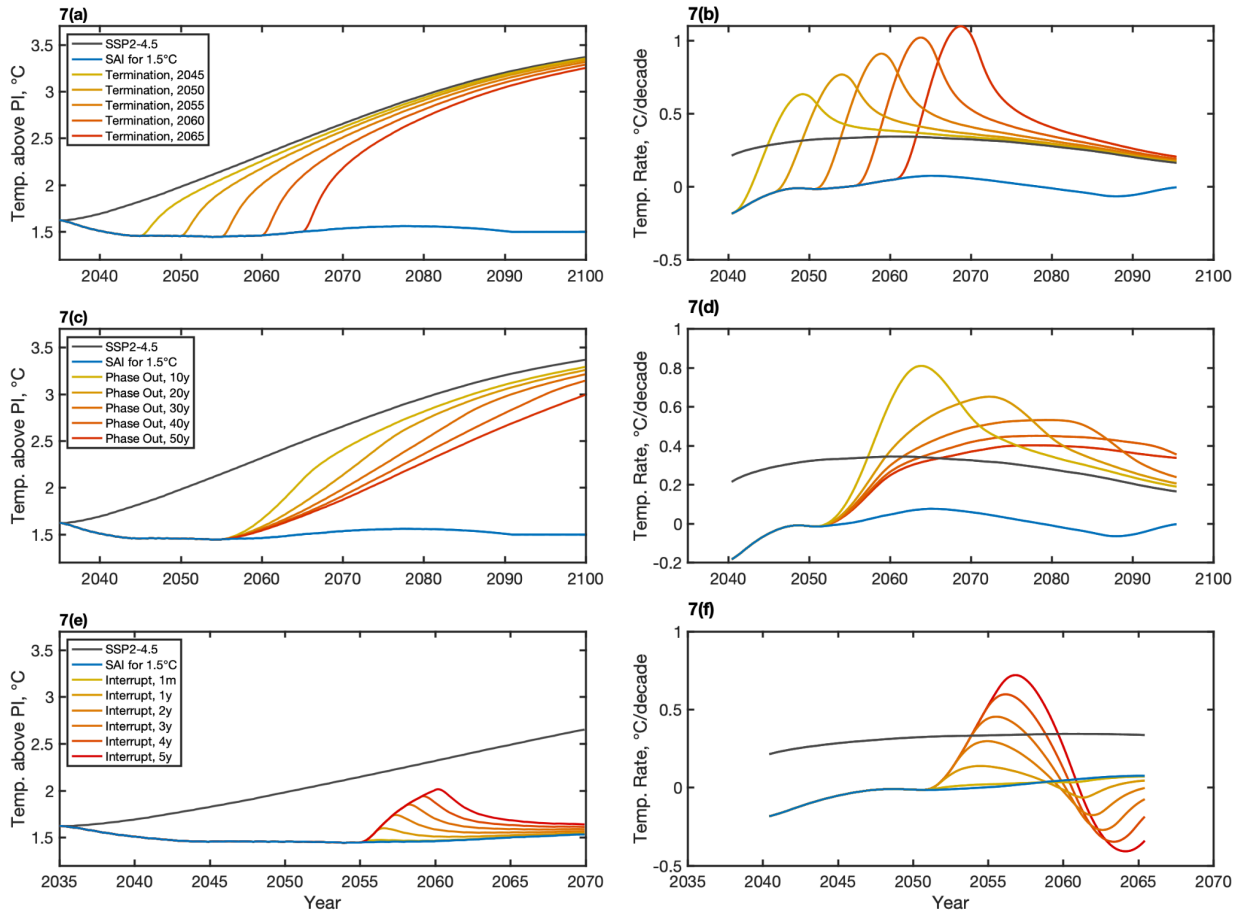

Figure 7: Emulation of alternate inconsistencies to the ones simulated. Left column of panels (a,c,e) shows the temperature above pre-industrial in °C, while the right column of panels (b,d,f) shows the 10-year temperature rates in °C/decade of the panel to the left. **a-b)** Terminations with a range of starting points, from 2045 to 2065 in 5-year intervals. The peak value of the 10-year warming rate occurs when the first month of that 10-year interval is the first month with no new aerosols. A slight delay occurs due to the lag between SO_2 injection and SO_4 generation. Notably, even terminating after only 10 years of SAI deployment sees a 10-year temperature rate above 0.5°C/decade, while terminating after 30 years sees a peak 10-year temperature rate near 1°C/decade. **c-d)** Phase-outs of alternate durations, from 10 years to 50 years in 10-year intervals. Phase-outs of 30 years or longer stay below 0.5°C/decade in rate. All contain higher rates than that of SSP2-4.5. **e-f)** Interruptions of alternate durations, 1-month and from 1 year to 5 years in 1-year intervals. The 1-month interruption sees a warming of less than 0.03°C. The 1-year interruption sees a warming of 0.17°C, and every year thereafter warms a smaller magnitude (as

one would expect from asymptotic behavior). Of note is that the 10-year warming rate for the 2-year interruption does not surpass that of SSP2-4.5, and that of the 4-year interruption does not surpass 0.5°C.

All of the simulations in Section 4 used SSP2-4.5 as the background; Figure 8a shows the inconsistencies simulated in this paper but under the SSP1-2.6 CO_2 concentrations instead. The termination plot at first glance may look different, but it can be helpful to think of termination as a loss of forced cooling, rather than strictly a rise in temperature. Another notable feature is that "peak shaving" is not as straightforward as terminating SAI when GHG-forced temperatures return to the temperature target – it takes time for the aerosol cooling to be lost, so the curves converge asymptotically (see also Figure 2 of MacMartin et al 2014 and Figures 3 and 4 of MacMartin et al 2018).

Figure 8c shows that the termination behavior depends not only on how much cooling there is but also on how long SAI has been deployed. It does this by prescribing no GHG forcing and a constant injection, and then terminating after different durations of SAI. The cooling from SAI is lost more quickly (the warming from a termination is more rapid) the shorter that SAI has been deployed; after a long consistent deployment then roughly ~50% of cooling is lost in a decade (as appears in Caldeira and Myhrvold 2013) while after a shorter duration that fraction is more like ~75% (as seen in Figure 1c as well). This difference in rate is due to the deep ocean having more time to equilibrate to SAI's cooler climate.

The emulator is also useful for illustrating how termination and interruption are affected by the ~1 year lifetime of aerosols in the stratosphere. Figure 8b shows interruption and termination with and without the dynamical effect of aerosols, corresponding either to zeroing the injection rate or instantaneously zeroing the AOD. The aerosols provide a damping effect to changes in injection, rounding out and delaying the peak temperature of an interruption. While the emulator is intended only to capture the response to SAI, removing the radiative forcing from the aerosols instantaneously is relevant to evaluating inconsistencies to other forms of SRM such as marine cloud brightening.

Finally, Figure 8d is of potential interest to understanding how realistically scenarios might need to be represented in climate modeling. In an actual deployment one might reasonably expect that injection rates would gradually change over time. In some modeling simulations, however, (notably the G6-sulfur simulations from the Geoengineering Modeling Intercomparison Project; Kravitz et al 2015, Visioni et al 2021), modelers have chosen to run the simulation for 10 years with fixed injection, look at the results, and update the injection rate for the next 10 years - a human-in-the-loop feedback algorithm. Comparing this modeling approach to one where the injection rates are smoothly varied (as in the simulations used herein) illustrates that the difference between the two cases is not necessarily large relative to natural variability.

The scenarios shown in Figure 8 represent an illustrative but small portion of the vast number of scenarios that someone may be interested in and that could be emulated. Additionally, having parameters of the climate system easily accessible means that one can use these emulators as tools for uncertainty assessment. A base example is shown in Figure 8b, answering the question of "What if the timescale of the aerosol dynamics was instantaneous?" One can imagine questions in the vein of "What if this parameter was larger/smaller?" that could be asked of every part of the emulator.

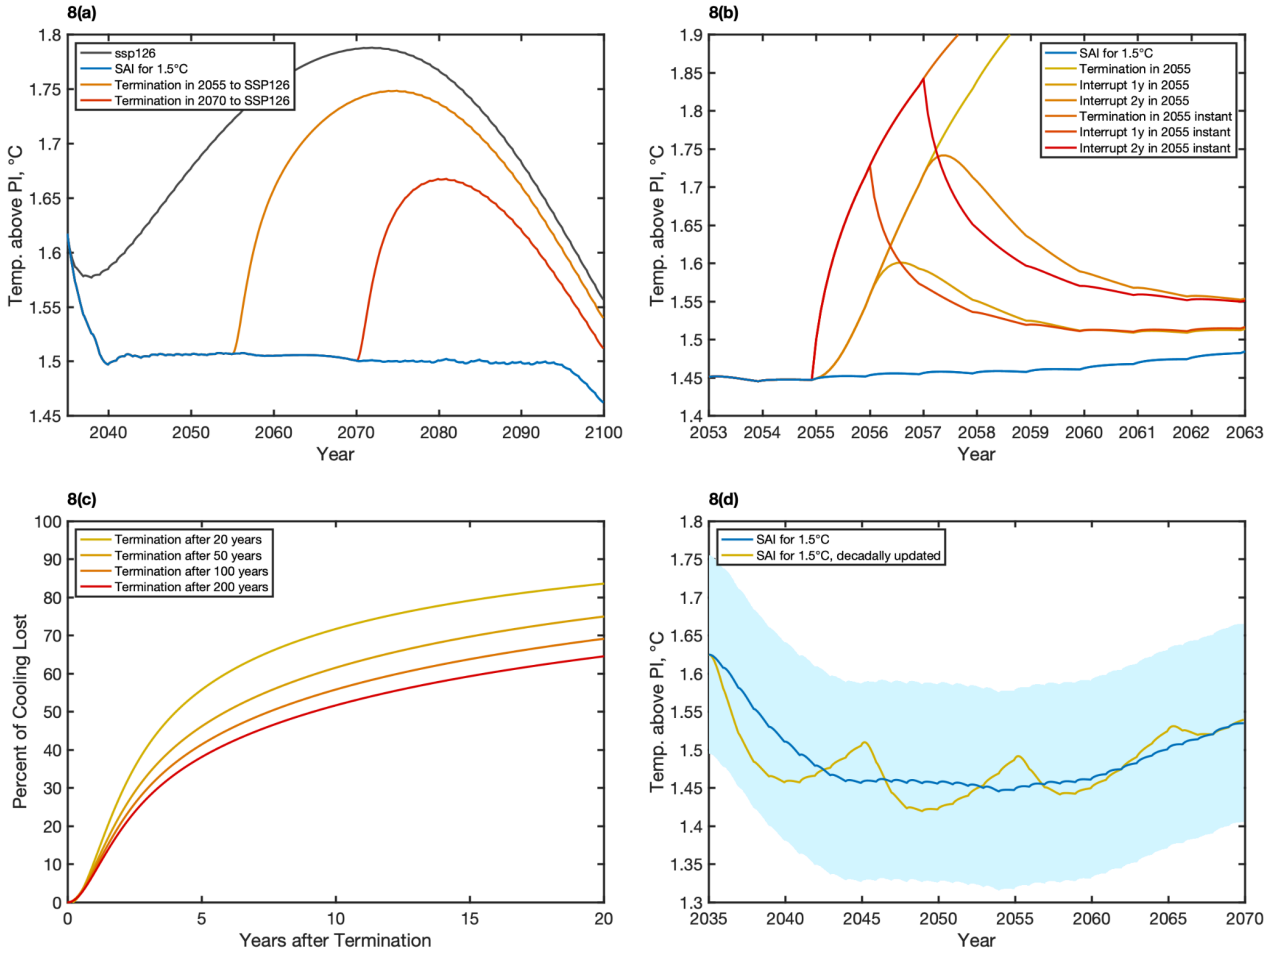

Figure 8: Alternate scenarios and physics exploration a) Inconsistencies under SSP1-2.6. Because termination is a loss of forced cooling, it will converge to the background temperature profile. If that background profile is cooling, terminated SAI will also cool. b) Loss of cooling (in percent) versus time after different lengths of consistent deployment. Loss of cooling slows as SAI is maintained. c) Inconsistency with and without aerosol lifetime, i.e. comparing stopping injection with zeroing AOD. When there are no aerosol effects, AOD is set to background levels at 2055 and returned to their consistent levels after an interruption ends. Aerosol dynamics create a delay in post-inconsistency warming of about one year - note the 1-year no-aerosol interruption as compared to the 2-year with-aerosol interruption, as well as when each interruption peaks. d) Ten-year-hold feedback injection as compared to annually-updated injection. Decadally-updated SAI is used in some simulations (G6-sulfur simulations from the Geoengineering Modeling Intercomparison Project; Kravitz et al 2015, Visioni et al 2021). Shading marks ± 1 standard deviation.

Section 6: Discussion

The previous sections explored "inconsistencies" in deployment such as termination or interruptions, both through direct simulation and through emulation; based on that a number of observations can be made.

First, as shown in Figure 1, the termination of SAI causes around 75% of its cooling to be lost within a decade. However, this percentage is not a fixed number. If SAI is deployed for a long time, as it is in Figure 5h, SAI is slower to lose its cooling in a termination. The most-sustained cooling emulated (320 years) loses only 50% of its cooling in a decade, rather than 75%. This 50% loss of cooling agrees with the equilibrium temperature response found in Caldeira and Myhrvold (2013). However, it only reaches this number if the termination occurs after a long, consistent deployment. Thus, safety criteria based around not exceeding some rate of warming in the event of termination should be built with the notion that this rate slows with time. One could use the emulator herein to find the allowable amount of cooling as it changes over time. For example, if one wanted to use the 0.2°C/decade of Parker and Irvine (2018), the maximum cooling allowed would start at 0.26°C and increase to 0.4°C over the course of the deployment.

It is worth pointing out that an interruption, until one begins re-injecting, is identical to a termination. If one wants to assess the first segment of an interruption (before re-initiation of injection), one can assess the respective segment of termination simulation output. However, the segment of an interruption after re-initiation of injection is affected by the strategy chosen. Herein we chose to restart injection simply by returning to using the injection amounts for the consistent deployment; with this restart strategy AOD returns at roughly the same rate it is depleted, implying the temporary disruption to cooling is roughly twice the length of the interruption to injection and reaches its maximum roughly when deployment restarts. In other words, a 1-year interruption creates a 2-year deviation in forcing.

Since it takes a full decade or longer for SAI to lose 75% of its cooling, a temporary interruption where SAI is restarted (potentially by a different actor) before the full effects of termination occur is plausible, and thus an important scenario to evaluate. The shorter the interruption, the less the deviation from the temperature target (see Figure 7e). Indeed, examining the frequency response of the climate system (MacMynowski et al 2011, Geoffroy et al 2013) can provide a general overview of what to expect regarding an interruption – the low-pass nature of global mean temperature means that the shorter the period of a particular forcing, more it is damped. Figure 5e follows this rule – while the five-year interruption warms 0.56°C by the peak of the interruption, the one-month interruption warms 0.03°C (in comparison to the standard deviation of 0.13°C). As such, robustness of deployment and design towards the quick re-institution of

injection in the event of losing it, on the order of months, can greatly limit the impact of this inconsistency.

Short-term inconsistencies to an otherwise consistent climate already exist, in the form of occasional large volcanic eruptions. Examining these may give a broad sense of what an interruption to injection during SAI might look like, with cooling instead of warming. During 2055 and 2056, the simulated strategies would inject an average of 7.5 and 7.7 Tg of SO_2 respectively. The maximum deviations from target temperatures caused by these 1-year and 2-year interruptions are 0.2°C and 0.3°C respectively in our simulations (the emulated values, which do not have natural variability, are more proportional to the duration of interruption at 0.15°C and 0.3°C respectively). This temporary change in temperature is roughly comparable to that from a volcanic eruption the size of Pinatubo, which cooled the globe by 0.4°C in 2 years (McCormick et al 1995).

Thus far, this discussion has centered around a sudden inconsistency. Such an abrupt change in deployment as termination has been shown (in this paper and many others) to have large and rapid effects on the climate. Nevertheless, an actor deploying SAI might still want to stop, and they might want to do it safely. How long must such a phase-out be, then? Clearly, a 10-year phase-out does not appreciably reduce the rate of warming in this climate model relative to an abrupt termination for the scenario simulated (where the termination or phase-out occurs after 20 years of deployment with 0.7°C cooling). It is possible to design a phase-out that would meet some specified safety criteria, but it could be a long enough operation that it might be more apt to describe it as a change in strategy than an inconsistency, akin to switching to a rate-limiting strategy as in MacMartin et al (2014). If one assumes a background emissions scenario with a roughly constant rate of temperature rise (e.g. SSP2-4.5), a constant temperature target (e.g.1.5°C above pre-industrial), and a safety limit on rate (e.g. 0.2°C/decade higher than SSP2-4.5, the limit proposed in Parker and Irvine (2018), one can easily find the minimum duration of a safe SAI phase-out. This simple linear approximation (see Supplement 3) shows that for the strategy used, a phase-out of length 1.5 times the amount of time deployed would be appropriate. For a phase-out beginning in 2055 under our scenario to meet these safety criteria, it would need to be approximately 30 years. The emulator confirms this approximation in Figure 5b.

Much of this phase-out's duration is due to the constantly-rising greenhouse gas emissions. Indeed, the story for each inconsistency in this paper changes as warming is limited. Each figure (save figures of consistent deployment and Fig. 5i) carries a GHG-induced increasing warming superimposed over the inconsistency's loss of cooling. This has a few consequences. Namely, any safety limit on rate has the background rate of warming already partway to the limit as a baseline. Implied within the limit of 0.2°C/decade of additional warming (Parker and Irvine 2018) is that there exists a rate of warming that should not be surpassed (0.5°C/decade for

CESM2). With steady background temperatures, new options are revealed. One can either phase out at a shorter duration, using all 0.5°C/decade, or one can keep the same 0.2°C/decade one would have otherwise been forced to use with rising emissions. This second option allows one to maintain an overall lower rate of warming on par with SSP2-4.5 and potentially fulfill lower rate limits like that found in Kosugi (2013).

The risks of potential inconsistencies in SAI deployment provide further reasons why mitigation of greenhouse gas emissions is necessary. The magnitude of any inconsistency's temperature deviation is determined by the background temperature profile – the lower the background temperature, the smaller the deviation. Additionally, the background temperature rate is added as a baseline to any inconsistency, meaning that inconsistencies like interruption and termination see a quicker emergence of impacts and faster climate velocities and phase-outs take longer.

Section 7: Conclusion

Overall, inconsistencies are critical to keep in mind when assessing SAI impacts. However, they also add exponentially to the already expansive scenarios of SAI. Within this study we provide tools for successfully exploring this expanded scenario space. For long-term and intentional inconsistencies, treating them as a change of strategy means that one can design them as one would design any other strategy, with safety criteria in mind. For even broader assessment, one can use an emulator like the one herein to survey an infinite range of scenarios, consistent and inconsistent) using a finite amount of simulations. These relatively simple dynamic models are able to capture key features of scenario space without needing to run any additional simulations.

There are limitations to the emulator presented in this study, however. This emulator only calculates global mean values, so spatial patterns are lost. Within this study, we also demonstrated that significant regional warming can appear within even a 1-year interruption, and one would not notice this if one looked at only the global mean temperature. A complete version of an emulator with regional resolution should also be able to model the effect of multiple injection locations. Such an emulator would also be able to present the effects of a partial inconsistency (e.g. a hemispherical termination), as well as a framework by which one might assess the physical impacts of a multi-actor, uncoordinated deployment of SAI. Development of such an emulator is planned as future work.

Additionally, this study and its emulator is constrained to only one climate model – any limitations or faults of the climate model are inherited by the emulator. An emulator of this form can be developed for other climate models without much adjustment necessary. In the same line of thought, forms of radiation management other than SAI (like marine cloud brightening) have the potential for inconsistency, and some have even more degrees of freedom than SAI.

Emulators developed for these other forms of geoengineering would assist in exploring their scenario spaces.

As a final note about inconsistencies in SAI is that many are inherently tied into the planning, decision-making, and execution of a deployment. This is obviously true for any intentional inconsistencies (as they can be viewed as a change in strategy), but it is also true for unintentional ones. This paper has shown that there are major differences between one-month and one-year interruptions. Proper preparation can shorten an interruption (see e.g. Parker & Irvine 2018), while poor logistics might itself be the cause of one. Discussion around the development of SAI should integrate the impacts of inconsistencies beyond their influence on whether or not to deploy – they should also shape how it deploys.

Acknowledgements

Support for J. Farley, D. G. MacMartin and D. Visioni was provided in part by the Cornell Atkinson Center for a Sustainable Future and in part by the National Science Foundation through agreement CBET-2038246. Support for B. Kravitz was provided in part by the National Science Foundation through agreement SES-1754740, NOAA's Climate Program Office, Earth's Radiation Budget (ERB) (Grant NA22OAR4310479), and the Indiana University Environmental Resilience Institute. The Pacific Northwest National Laboratory is operated for the U.S. Department of Energy by Battelle Memorial Institute under contract DEAC05-76RL01830. The CESM project is supported primarily by the National Science Foundation. High-performance computing was provided by NSF NCAR's Computational and Information Systems Laboratory (https://doi.org/10.5065/D6RX99HX), sponsored by the National Science Foundation. Code for emulator and data used for inconsistency figures are available at https://github.com/jf678-cornell/Inconsistency_Emulator.git. Other key simulation variables for consistent deployment are available through the Cornell eCommons platform (https://hdl.handle.net/1813/111357) (Visioni et al. 2022). All climate model simulation output is available at Globus (https://app.globus.org/file-manager?destination_id=dc94169c-30f3-4866-a305-f112408bcf8f&destination_path=%2Fdgm224_0001%2F); authorization is needed to access the dataset; please contact D. Visioni.

References

Aubry, T. J., Toohey, M., Marshall, L.,Schmidt, A., & Jellinek, A. M. (2020). A new volcanic stratospheric sulfate aerosol forcing emulator (EVA_H): Comparison with interactive

stratospheric aerosol models.Journal of Geophysical Research:Atmospheres,125, e2019JD031303. https://doi.org/10.1029/2019JD031303

Baum, S.D., Maher, T.M. & Haqq-Misra, J. (2013). Double catastrophe: intermittent stratospheric geoengineering induced by societal collapse. *Environ Syst Decis* **33**, 168–180 https://doi.org/10.1007/s10669-012-9429-y

Bhowmick, M., Mishra, S.K., Kravitz, B. *et al* (2021). Response of the Indian summer monsoon to global warming, solar geoengineering and its termination. *Sci Rep* **11**, 9791 https://doi.org/10.1038/s41598-021-89249-6

Caldeira, K. & Myhrvold, N. P. Res. Lett. 8 034039 (2013). https://doi.org/10.1088/1748-9326/8/3/034039

Danabasoglu, G., Lamarque, J.‐F.,Bacmeister, J., Bailey, D. A., DuVivier,A. K., Edwards, J., et al (2020). TheCommunity Earth System ModelVersion 2 (CESM2).Journal ofAdvances in Modeling Earth Systems,12, e2019MS001916. https://doi.org/10.1029/2019MS001916

Davis, N. A., Visioni, D., Garcia, R. R., Kinnison, D. E., Marsh, D. R., Mills, M., et al (2023). Climate, variability, and climate sensitivity of "Middle Atmosphere" chemistry configurations of the Community Earth System Model Version 2, Whole Atmosphere Community Climate Model Version 6 (CESM2(WACCM6)). Journal of Advances in Modeling Earth Systems, 15, e2022MS003579. https://doi.org/10.1029/2022MS003579

Dogar, M.M., Hermanson, L., Scaife, A.A. et al (2022). A Review of El Nino Southern Oscillation Linkage to Strong Volcanic Eruptions and Post-Volcanic Winter Warming. Earth Syst Environ https://doi.org/10.1007/s41748-022-00331-z

Geoffroy, Olivier & Saint-Martin, D. & Olivié, D. & Voldoire, A. & Bellon, Gilles & Tytéca, S. (2013). Transient Climate Response in a Two-Layer Energy-Balance Model. Part I: Analytical Solution and Parameter Calibration Using CMIP5 AOGCM Experiments. Journal of Climate. 26. 1841-1857 . https://doi.org/10.1175/JCLI-D-12-00195.1

Goes, M., Tuana, N. & Keller, K. (2011) The economics (or lack thereof) of aerosol geoengineering. *Climatic Change* **109**, 719–744 https://doi.org/10.1007/s10584-010-9961-z

Hansen, J. et al (1984). Climate sensitivity: Analysis of feedback mechanisms. in Geophysical Monograph Series (eds. Hansen, J. E. & Takahashi, T.) vol. 29 130–163 (American Geophysical Union, Washington, D. C.).

Hueholt, Daniel M., Elizabeth A. Barnes, James W. Hurrell, and Ariel L. Morrison, (2023). Climate speeds help frame relative ecological risk in future climate change and stratospheric aerosol injection scenarios. In review at Nature Communications. Preprint: https://doi.org/10.21203/rs.3.rs-3463352/v1

Irvine, P., Sriver, R. & Keller, K. (2012). Tension between reducing sea-level rise and global warming through solar-radiation management. *Nature Clim Change* **2**, 97–100. https://doi.org/10.1038/nclimate1351

Jackson, L. S., Crook, J. A., Jarvis, A., Leedal, D., Ridgwell, A., Vaughan, N. and Forster, P. M. (2015). Assessing the controllability of Arctic sea ice extent by sulfate aerosol geoengineering. *Geophys. Res. Lett.*, 42: 1223–1231. https://doi.org/10.1002/2014GL062240

Jones, A., et al (2013). The impact of abrupt suspension of solar radiation management (termination effect)in experiment G2 of the Geoengineering Model Intercomparison Project (GeoMIP),J. Geophys. Res. Atmos.,118,9743–9752. https://doi.org/10.1002/jgrd.50762

Keith, D., MacMartin, D. (2015). A temporary, moderate and responsive scenario for solar geoengineering. *Nature Clim Change* **5**, 201–206. https://doi.org/10.1038/nclimate2493

Kikstra, J. S. et al (2022). The IPCC Sixth Assessment Report WGIII climate assessment of mitigation pathways: from emissions to global temperatures. Geosci. Model Dev. 15, 9075–9109 https://doi.org/10.5194/gmd-15-9075-2022

Kosugi, T. (2013). Fail-safe solar radiation management geoengineering. Mitig Adapt Strateg Glob Change 18, 1141–1166. https://doi.org/10.1007/s11027-012-9414-2

Kravitz, B., MacMartin, D. G., Mills, M. J.,Richter, J. H., Tilmes, S., Lamarque, J.-F.,...Vitt, F. (2017). First simulations of designing stratospheric sulfate aerosol geoengineering to meet multiple simultaneous climate objectives. Journal of GeophysicalResearch: Atmospheres,122, 12,616–12,634.https://doi.org/10.1002/2017JD026874

Lee, W., MacMartin, D., Visioni, D., and Kravitz, B. (2020). Expanding the design space of stratospheric aerosol geoengineering to include precipitation-based objectives and explore trade-offs, Earth Syst. Dynam., 11, 1051–1072 https://doi.org/10.5194/esd-11-1051-2020

MacMartin, D. G., Caldeira, K. & Keith, D. W. (2014). Solar geoengineering to limit the rate of temperature change. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 372, 20140134. https://doi.org/10.1098/rsta.2014.0134

MacMartin, D. G. and Kravitz, B. (2016). Dynamic climate emulators for solar geoengineering, Atmos. Chem. Phys., 16, 15789–15799. https://doi.org/10.5194/acp-16-15789-2016

MacMartin, D. G. et al (2022). Scenarios for modeling solar radiation modification. Proc. Natl. Acad. Sci. U.S.A. 119, e2202230119. https://doi.org/10.1073/pnas.2202230119

MacMynowski, D. G., Shin, H.-J. & Caldeira, K. (2011). The frequency response of temperature and precipitation in a climate model: CLIMATE FREQUENCY RESPONSE. Geophys. Res. Lett. 38, n/a-n/a https://doi.org/10.1029/2011GL048623

McCormick, M., Thomason, L. & Trepte, C. (1995). Atmospheric effects of the Mt Pinatubo eruption. *Nature* **373**, 399–404. https://doi.org/10.1038/373399a0

Myhre, G., Highwood, E. J., Shine, K. P. & Stordal, F. (1998). New estimates of radiative forcing due to well mixed greenhouse gases. Geophysical Research Letters 25, 2715–2718. https://doi.org/10.1029/98GL01908

Parker, A. and Irvine, P. J. (2018). The Risk of Termination Shock From Solar Geoengineering, Earth's Future, 6, 456–467. https://doi.org/10.1002/2017EF000735

Daniel Pflüger, Claudia Elisabeth Wieners, Leo van Kampenhout, et al No Emergency Brake: Slow Ocean Response to Abrupt Stratospheric Aerosol Injection. ESS Open Archive . September 11, 2023. In review

Quaglia, I., Timmreck, C., Niemeier, U., Visioni, D., Pitari, G., Brodowsky, C., Brühl, C., Dhomse, S. S., Franke, H., Laakso, A., Mann, G. W., Rozanov, E., and Sukhodolov, T. (2023). Interactive stratospheric aerosol models' response to different amounts and altitudes of SO_2 injection during the 1991 Pinatubo eruption, Atmos. Chem. Phys., 23, 921–948. https://doi.org/10.5194/acp-23-921-2023

Quaglia, I., Visioni, D., Bednarz, E.M., MacMartin, D.G., Kravitz, B.: The potential of Stratospheric Aerosol Injection to reduce the climatic risks of explosive volcanic eruptions (2024), Geophysical Research Letters, in review

Richter, J. H., Visioni, D., MacMartin, D. G., Bailey, D. A., Rosenbloom, N., Dobbins, B., Lee, W. R., Tye, M., and Lamarque, J.-F. (2022). Assessing Responses and Impacts of Solar climate intervention on the Earth system with stratospheric aerosol injection (ARISE-SAI): protocol and initial results from the first simulations, Geosci. Model Dev., 15, 8221–8243. https://doi.org/10.5194/gmd-15-8221-2022

Ritchie, P. D. L., Alkhayuon, H., Cox, P. M., and Wieczorek, S. (2023). Rate-induced tipping in natural and human systems, Earth Syst. Dynam., 14, 669–683. https://doi.org/10.5194/esd-14-669-2023

Tilmes, Simone, Jadwiga H. Richter, Ben Kravitz, Douglas G. MacMartin, Michael J. Mills, Isla R. Simpson, Anne S. Glanville, John T. Fasullo, Adam S. Phillips, Jean-Francois Lamarque, Joseph Tribbia, Jim Edwards, Sheri Mickelson, and Siddhartha Ghosh. (2018). "CESM1(WACCM) Stratospheric Aerosol Geoengineering Large Ensemble Project". *Bulletin of the American Meteorological Society* 99.11: 2361-2371. https://doi.org/10.1175/BAMS-D-17-0267.1

Trisos, C.H., Amatulli, G., Gurevitch, J. et al (2018). Potentially dangerous consequences for biodiversity of solar geoengineering implementation and termination. *Nat Ecol Evol* **2**, 475–482 https://doi.org/10.1038/s41559-017-0431-0

Visioni, D., Pitari, G., Tuccella, P. & Curci, G. (2018). Sulfur deposition changes under sulfate geoengineering conditions: quasi-biennial oscillation effects on the transport and lifetime of stratospheric aerosols. Atmos. Chem. Phys. 18, 2787–2808. https://doi.org/10.5194/acp-18-2787-2018

Visioni, D. et al. (2022), Data from “Scenarios for modeling solar radiation modification.” https://hdl.handle.net/1813/111357. Deposited 7 July 2022.

Supplement 1: La-Niña-like pattern one year after inconsistency

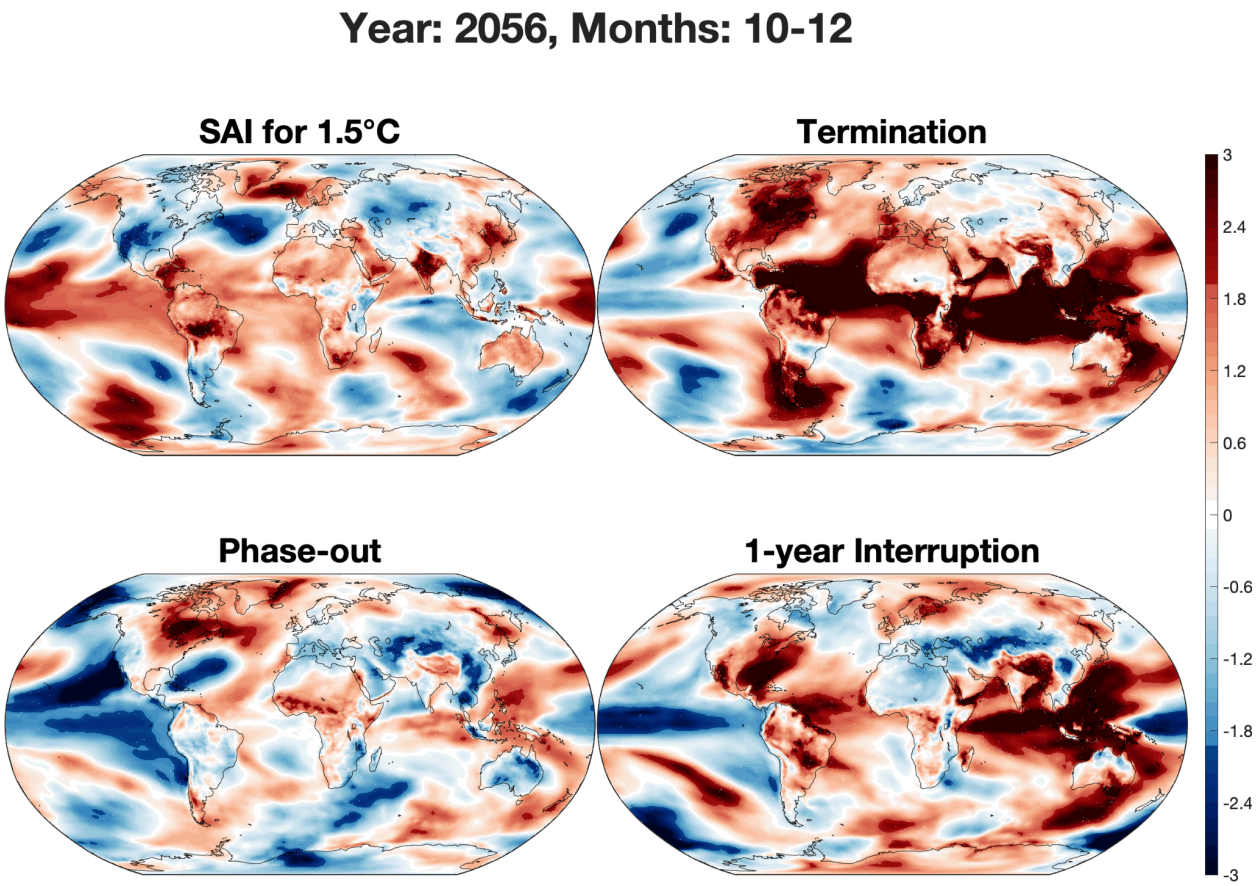

Supplement 1: Standard errors of a 3-member ensemble away from the steady-state (2050-2069 consistent SAI) mean. Values are averaged across ensemble members and across months 22-24 of the labeled inconsistency. All errors are calculated grid cell to grid cell.

Supplement 2: Emulator trained on SAI for 1.5°C, emulating higher-AOD deployments

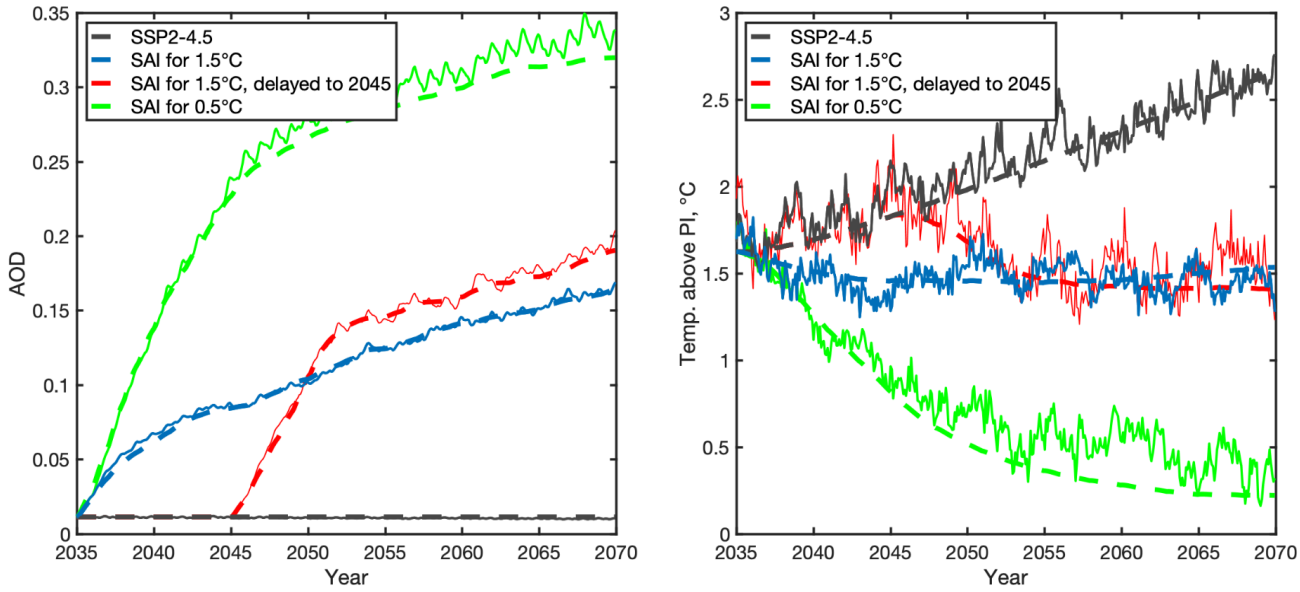

Supplement 2: (a) AOD emulation and (b) temperature emulation of other consistent scenarios – a lower temperature target of 0.5°C and a temperature target of 1.5°C, delayed to 2045. In these plots, the emulator is trained on only the SSP2-4.5 and SAI-for-1.5°C simulations, so high AOD is underestimated.

Supplement 3: A Phase-Out Approximation

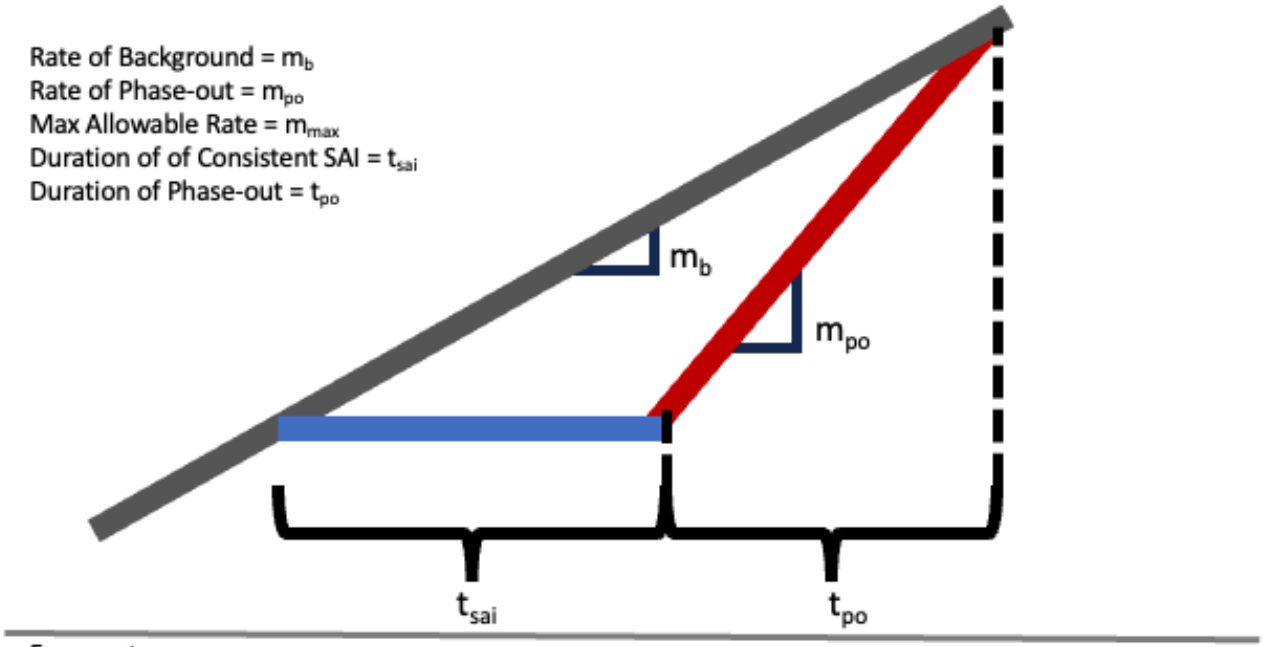

For $m_{po} \leq m_{max}$:

$$\frac{m_b\,(t_{sai} + t_{po})}{t_{po}} \leq m_{max}$$

$$t_{po} \geq \frac{m_b}{m_{max} - m_{sai}}\, t_{sai}$$

If m_b = 0.3°C/decade (CESM2, SSP2-4.5)
And m_{max} = m_b + 0.2°C/decade = 0.5°C/decade (Parker and Irvine 2018)

Then **t_{po} = 3/2 * t_{sai}**

For our t_{sai} = 20 years, **t_{po} = 30 years**