Simulating the impacts of SAI deployment (November 2024)

Even in a best-case scenario where drastic emissions cuts and rapid scaling of carbon dioxide removal limit warming to less than 2°C this century, we are facing a world with significant near-and mid term climate impacts, including extreme weather events, biodiversity loss, and increasing risk of catastrophic “tipping points.” The Institute for Economics & Peace (IEP) predicts that around 1.2 billion people could be displaced by 2050 due to natural disasters and climate change (Institute for Economics & Peace 2020), posing significant risk to domestic and global stability.

Current evidence suggests that stratospheric aerosol injection (SAI) may be the only option to hold temperatures closer to the Paris 1.5°C target, thereby reducing climate-related harms to vulnerable populations around the world, protecting the natural environment, and easing the burden of adaptation.

But *how* SAI is deployed really matters, on both physical and geopolitical levels. A coordinated, international effort that deploys aerosols in both hemispheres could, according to current research, produce relatively uniform cooling around the planet. Unilateral deployment in a single hemisphere, however, could drastically change precipitation patterns around the globe, endanger the food supplies of billions, and lead to international conflicts.

While we understand, in broad strokes, that SAI could cool the planet, the science remains uncertain, and we lack sufficient understanding of the various impacts of different SAI deployment scenarios for policymakers to make well-informed decisions regarding its potential deployment.

There are a considerable number of potential choices when studying possible SAI deployment scenarios, including the latitude and altitude of injection, the aerosol or particle used, the cooling target/quantity of aerosols injected, and the starting date and ramp-up rate.

The vast majority of SAI research, currently, relies on supercomputers running incredibly complex Earth Systems Models (ESMs). Because of this, it’s expensive and difficult to generate SAI forecasts and scenarios: the research community doesn’t have access to the computing capacity required to analyze each permutation of these choices and study the local impacts of each.

Additionally, researchers in many parts of the Global South lack access to the required computational resources and cannot identify and run scenarios themselves. They’re limited to studying the downstream impact of scenarios defined and modeled by researchers with ready access to best-in-class modeling centers. This means that the majority of Global South researchers cannot explore how different scenarios might yield better or worse impacts for their own country and policymakers have to trust that scenarios primarily defined by researchers in the Global North have been fairly and adequately designed to address national priorities other than those of their own countries.

And there’s a staggering lack of access to the scientific output that has been generated. It’s hidden behind paywalls, written in scientific jargon, and often explores the effects of only one SAI deployment scenario (among millions of potential scenarios), leading to misleading statements like “SAI would do X.”

The result is that the debate around SAI is polarized, but mostly unscientific. Policymakers and their staffers who are interested in SAI, or who are fielding valid questions from constituents about the topic, don't have the tools they need to respond to questions and concerns or to assess if and how SAI could be deployed safely. And researchers cannot ask the questions germane to them.

We've designed this simulator to address this bottleneck.

We want to equip policymakers and the public, regardless of geography or access to best-in-class computational tools, to navigate this complicated science, participate in scenario design, and explore the local effects of different deployment scenarios. We've intentionally designed the user interface to help visualize how the effects of a potential SAI deployment would compare to the effects of continued warming (aka "risk-risk analysis").

We're building on top of a long history of the IPCC using similar emulators to accelerate and extend what can be done with ESMs and on top of a rich foundation of academic research. We've aimed to merge the strongest elements of several existing methods and models to create information that is accessible, actionable and (reasonably) intuitive. Below we describe how we developed the simulator and provide more detail on assumptions, methods, and uncertainties.

But first, two important notes:

1. Any model will always be a limited representation of reality. While climate models, including the one on which this simulator is based, are important tools for understanding past, current and future global climate variability, they are not able to provide perfect predictions of the future. Models don't entirely agree with one another about the effects of future warming, there's no way to represent all of the underlying physics and interactions even in the most complex models, and we lack adequate observational data to calibrate and evaluate models with accuracy.

   For now, we consider this to be a starting point for understanding the potential impacts of a deployment, but also want to use this as a tool to advocate for further, more robust research to reduce scientific uncertainties and to better quantify the broad range of effects of a potential deployment. And we plan to add data from multiple models, allowing us to better represent inter-model uncertainty, soon.

2. What we're releasing today is just a first step. Further scientific research may alter what is shown here, there are known limitations to this version of the simulator (see the Limitations section below), and there are also many yet-to-be-identified gaps in our approach. So what we'd ask of you: poke holes, tell us what you'd prioritize next, tell us what you'd adapt or change. We're eager to build on this with you.

How we developed the simulator

We've designed the tool to be modular: users select inputs, the SAI emulator module estimates global and regional temperature and precipitation, and then the impact analytics module uses projected temperature and precipitation as inputs to derive projections for other climate impacts. Temperature, precipitation and derived impacts are visualized for users to explore.

Choosing input variables

The tool is designed to help compare a world with SAI and a world without SAI. In order to help visualize this "risk-risk framework" and accommodate different potential scenarios about future warming, we wanted users to have the ability to select different scenarios to use as the comparative baseline. To do that, we've used the **Shared Socioeconomic Pathways**, which are used throughout the IPCC Sixth Assessment Report to derive greenhouse gas emissions under different climate policies.

Users then can determine how much to cool the planet by selecting a cooling target. We've chosen to express this as a target global mean temperature above pre-industrial, which mirrors international agreements, including the Paris Agreement which set the goals to "hold global temperature increase to well below 2°C above pre-industrial levels and pursue efforts to limit it to 1.5°C above pre-industrial levels, recognizing that this would significantly reduce the risks and impacts of climate change.

Users then can select the deployment start date based on their sense of when an SAI deployment could, or should, start. The start date affects the ramp-up rate and the level of cooling possible. Based on current estimates of the fastest-possible R&D pathway, we've selected 2035 as the earliest possible start date – this would imply significant investments in R&D now, such that we have the capability to deploy at the altitude and scale required to hit the user-selected policy target. We've flagged user-defined scenarios that would ramp up cooling faster than the current rate of warming: research into what would be a "safe" rate of cooling is an important, but nascent, research avenue (Hueholt et al. 2024). For now, we are simply highlighting global rates of cooling that are larger than the current rate of warming.

In this initial version, users do not have more granular control over the SAI deployment scenario, which follows the protocol defined in Richter et al. (2022). This protocol injects at four latitudes (15°S, 15°N, 30°S, 30°N) at an altitude of 21.5 km. A "controller" algorithm adjusts how much SO_2 is injected at each of these sites in order to meet the cooling target, while also aiming to maintain temperature gradients between hemispheres and from the equator to poles—essentially ensuring the most even cooling possible. SO_2 injection rates ramp up linearly from the start year to however much SO_2 is necessary to achieve the temperature target over a period of 10 years.

We anticipate future versions will allow users to change the latitude(s) of injection, adjust the ramp-up rate, and adjust the continuity of injection (as might be affected by global politics).

What we don't include and why

We don't include the option for users to select the altitude of deployment. Most simulations to date have injected material at 22 km – this altitude is high enough above the height of the tropopause in the tropics that aerosols will remain in the stratosphere for a sustained period of time, while also being low enough to be in the practical range of potential aircraft (Smith et al. 2020). Introducing the option to select injection altitude would introduce significant complexity for little return, as it has relatively little impact on surface climate (Usha et al. 2024).

While we plan to allow for the selection of injection latitude in the future, there is no need to add the ability to select for longitude of injection. The stratosphere mixes well east-west, therefore it is largely irrelevant at what longitude SO_2 is injected — not just here, but in research simulations as well.

SAI emulator module

The emulator module is split into two submodules: one which calculates regional values in a scenario with no SAI and one which calculates regional differences due to SAI. The methodologies and assumptions for each are described below.

Regional values with no SAI

We use the Finite-amplitude Impulse-Response (FaIR) (Leach et al. 2021) emulator to predict global mean temperature in each of the baseline warming scenarios. FaIR is an open source, reduced complexity climate model that allows for modeling of different emissions scenarios and uncertainties. We selected FaIR to emulate global temperatures in order to enable more flexible specifications of CO_2 pathways, which importantly allows us to extend the approach to eight different SSPs and eventually allows us to handle those SSPs across multiple climate models. We then ”pattern scale” the global mean temperature to obtain regional values of each variable under the selected SSP scenario without SAI.

Regional differences due to SAI

Based on the user-specified cooling target, we derive the temperature difference between the baseline SSP and the cooling target each year. If the target temperature is already at or below the temperature without SAI, we do not emulate any injection that year.

In order to determine injection amounts, we use CESM2-WACCM simulations described in Visioni et al. (2023) that model SAI scenarios to keep temperatures under 1.5, 1.0 and 0.5°C above pre-industrial. These simulations maintain global mean temperatures, as we've aimed to do with this simulator, while also maintaining temperature gradients between hemispheres and from the equator to poles—essentially ensuring the most even cooling possible.

To model other regional output values, such as precipitation or number of days above 35 C, we assume that the changes in these variables are proportional to the change in global mean temperature (Tye et al. 2022), which allows us to emulate many of the outputs for cooling targets between 0.5°C and 2.5°C by linearly interpolating the simulated values from the three experiments above. We separately add the predicted contributions from greenhouse gas emissions and SAI, and with the proportionality for each variable calibrated using CESM2-WACCM6-MA simulations. This “pattern scaling” is a standard, well-validated approach in emulating the response to climate change (Tebaldi et al. 2014).

Output data from the emulator module includes temperature variables (mean/max/min annual temperature, number of hot and extreme heat days per year, number of freezing days per year, number of person-days exposed to hot and extreme heat per year, number of person-days exposed to hot and extreme heat, number of person-days exposed to extreme cold) and precipitation variables (mean water availability, number of heavy/extreme precipitation days per year, number of person-days exposed to heavy/extreme precipitation per year).

Impact analytics module

The impact analytics module is designed to take these direct model outputs and use them as inputs for “downstream” impact modeling. We are currently developing support for a limited set of impacts that, based on current literature, can be robustly derived from climate model outputs (primarily temperature and precipitation). Specifically, we are in the process of modeling deposition, air quality (PM2.5), sea ice extent, and sea level rise. As the academic literature grows, and as spatial and temporal downscaling can be built into the tool, we will add additional, well-studied climate impacts.

Choosing baselines for comparison

We define the pre-industrial baseline as the climate models' historical mean (from 1850-1900), and plot values of the variables relative to this baseline.

When viewing the emulator, you will see either one or two lines for values during past dates. The red line seen for every variable is based on the same CESM2-WACCM simulations used for the baseline values as described above. The black line represents the observed mean temperature, specifically the annual land and ocean anomaly using air temperature above sea ice from Berkeley Earth. Climate models are not fit to historical temperatures, which is why the black and red lines diverge in places. However, as you can see, climate models have historically been quite successful at predicting future warming.

Derived statistics

There are a number of useful statistics we calculate based on the SAI scenario selected. Here we explain how those calculations are made.

Estimation of the amount of SO_2

We calculate this figure based on the amount of cooling required as prescribed by the scenario selected. We use a linear regression to interpolate between the amount of SO_2 injected to produce the amount of cooling in each of the scenarios simulated in Visioni et al. (2023).

How much SO2 is injected relative to Pinatubo

We calculate this by comparing the annual amount of SO_2 injected (averaged over the decade) to the ~15 Tg of SO_2 lofted into the stratosphere by the Mt. Pinatubo volcanic eruption (Quaglia et al. 2023). This eruption caused ~0.5°C of global cooling which peaked at 18 months after the eruption (Soden et al. 2022 and, as such, may also provide a useful comparison point for the impacts of SAI.

Number of airplanes required to deploy

We use the approach from Smith et al. 2020 to calculate the number of planes required to deploy and direct costs. In the paper, Mr. Smith projects the cost of a variety of potential deployment scenarios, including three different warming scenarios and three potential radiative forcing targets (halving future warming, halting warming, and reversing temperatures to 2020 levels). His forecasts assume development of three successive generations of specialized aircraft able to reach ~20 km and use of sulfate aerosols.

We calculate the number of planes based on his assumption that each plane can carry 20 tons of SO_2 per flight. While no airplane can currently reach the ~20 km altitude necessary for deployment while also having a large payload, this is a reasonable estimate for a purpose-built high-altitude delivery aircraft. We assume each plane can make 10 flights per day, and will fly daily. Based on these values, we can calculate the amount of planes necessary in order to carry the annual SO_2 amount determined above.

How much will deployment cost

Notably, while the aggregate cost of the scenarios studied in Smith et al. 2020 differs by an order of magnitude, cost-per-ton of deployed aerosol varies little among scenarios (ranging from $2223 - $2987 per ton of SO_2 as depicted in Table 6 of the above paper). We take the average amount of SO_2 injected (averaged over the decade) and multiply it by this range to arrive at the range of total costs per year. This paper accounts for many key drivers of direct cost, including the cost of aircraft R&D, marginal build costs for new aircraft, operating costs (i.e. maintenance, ground operations, fuel, staffing), and the cost of the aerosol injected. We do not attempt to include indirect economic costs such as environmental externalities yet.

Current limitations of the simulator

There are a number of limitations to this version of the simulator. Some of these we plan to address in the short-term, while others we simply want to flag.

- We do not include depictions of interruption or termination, both of which are legitimate causes for concern about a potential SAI deployment.
- The emulator module has been developed using data from a limited set of model runs, all of which represent idealized deployment scenarios. This means that users, currently, cannot explore the effects of non-cooperative SAI deployment scenarios nor can users vary the latitude/altitude of deployment to try and tune the effects more specifically or explore, for example, a high-latitude but lower-altitude deployment scenario that could be initiated more rapidly using existing aircraft..

- We've started analyzing only a limited set of climate impacts which we felt could be robustly derived from temperature and precipitation. This means that we are not including many policy-relevant variables, some of which may reveal significant risks stemming from SAI deployment. We will keep adding additional climate impacts as the academic literature grows.
- We're only using one model. The simulator is currently based on a single global climate model, the Community Earth Systems Model Version 2, Whole Atmosphere Community Climate Model Version 6 (CESM2(WACCM6)), which is one of the most rigorously validated model for SAI simulations to date (Davis et al., 2023). We plan to include outputs from multiple climate models in future versions of the simulator to show the differences in projected SAI impacts. However, in previous studies it has been found that in general, climate models show smaller differences in temperature and precipitation outcomes (in most regions) due to SAI forcing relative to the much larger differences due to anthropogenic forcing and internal variability (Visioni et al., 2021). This means that the differences in temperature and precipitation between models due to SAI will contribute only partially to the uncertainty estimates shown in Reflective's simulator, and that climate model intercomparison, based on the few models that have performed the necessary SAI simulations, is in of itself a not so meaningful representation of the overall uncertainty present in SAI simulations, especially considering the potential for differences under different SAI injection strategies and targets. Hopefully, future efforts to expand research and modeling of SAI will provide us with clearer ways to expand our depiction of uncertainty (Visioni et al., 2024).

Moreover, there are currently no established diagnostics for evaluating model representations of past volcanic eruptions against observations, which makes it challenging to assess the skill of one model relative to another. Further research is also needed in this area to understand what constraints are provided by observations in terms of SAI projections.

In response to the dearth of data and tools needed for uncertainty assessment, Reflective team members are currently developing a set of observations-based diagnostics to support routine model evaluation and intercomparison. The diagnostics package will also include solutions to distinguish the response of the climate system to SAI from that of internal variability. Future versions of the simulator will feature uncertainty assessment metrics derived from the new diagnostics.

Broadly speaking, this simulator doesn't tell the full story about the potential impacts of SAI and how those impacts compare to the impacts of climate change. We're just beginning to characterize these impacts and, while we hope this tool can help inform future decision-making, we don't want to lose necessary nuance.

Looking forward

This is the first step. Moving forward, we initially plan to build on this work in four ways, some of which have been mentioned above:

- We want to give users the ability to run scenarios with varying latitudes of injection. This is important because higher-latitude deployment may provide an avenue to preferentially cool the poles and reduce the risks posed by polar ice melting. Additionally, the tropopause, the zone between the troposphere and stratosphere, is much lower at high-latitudes and hypothetically could be reached using modified existing aircraft, providing a potential pathway to earlier deployment and/or lower costs.
- In order to allow us to compute an expanded set of climate and socio-economic impacts, we need to build in spatial and temporal downscaling.
- We want to incorporate more training data, from more models, to broaden the scenario space that can be simulated and better represent uncertainty – or certainty.

- We plan on developing a benchmarking framework so that we can validate the simulator's performance. If this also spurs the development of even better simulators, all the better!

Also on the immediate horizon: we're following the lead of our friends at CarbonPlan and are making all of the input and output datasets public, alongside documented open-source code. As they said well, "as this kind of data increasingly becomes the basis for decisions — which could impact millions of people's lives and trillions of dollars — we need full transparency to enable public accountability. Alongside the dataset itself, we hope this work provides an example for how to model climate impacts in the open." See our GitHub repository for links to all datasets and documented code.

Also on the immediate horizon: we're following the lead of CarbonPlan and are making all of the input and output datasets public, alongside documented open-source code. See our GitHub repository for links to all datasets and documented code.

Finally, we plan to co-develop future versions of this simulator with users. If you have questions, feedback, or specific functions you'd like to see built into this tool, please get in touch here.