Sunlight Reflection Management Primer

Walker Raymond Lee, Douglas MacMartin, Daniele Visioni, Spencer Adler

Table of Contents

1. Overview

1.1. Introduction

In a perfect world, this document would not need to exist. Unfortunately, society's continued reliance on fossil fuels is causing the planet to warm at a potentially catastrophic rate. The 2015 Paris Climate Accords established the goal of limiting global warming above pre-industrial temperatures to well below 2°C, and ideally to below 1.5°C. However, it is a near certainty that this goal will not be achieved. The international community's pledged emission cuts are well below what is needed to meet these targets; even if they are met, the current concentration of atmospheric CO_2 is 412 ppm, and the current NOAA Annual Greenhouse Gas Index, which factors in other greenhouse gases (GHGs), measures over 500 ppm of CO_2 equivalent. Even if emissions were zeroed tomorrow - an impossible feat - previous emissions could still cause temperatures to rise above the 1.5C target unless CO_2 is artificially removed from the atmosphere. While carbon dioxide removal (CDR) technologies do exist, no CDR technology is currently scientifically mature enough to be economically viable at scale, and inherent physical limitations mean that, while these efforts are vital, they may not be sufficient to avert the worst consequences of climate change.

One possible supplement to emission cuts and carbon dioxide removal is to reflect a small amount of incoming sunlight back to space in order to deliberately cool the planet (e.g., NASEM 2021); in this document, we refer to this idea using the abbreviation SRM, which could interchangeably stand for for sunlight reflection management, solar radiation management, or either of those terms using the words "methods" or "modification" in place of "management." All variations have been used before. The idea is also known variously as climate engineering, solar geoengineering, or solar climate intervention. Currently, the climate is approximately 1.2°C warmer than pre-industrial levels; using SRM alone, offsetting all of this warming would require reflecting a little less than 1% of incoming sunlight back into space, cooling the climate for as long as the SRM effort is sustained. This could be done in several ways, including stratospheric aerosol injection (SAI), marine cloud brightening (MCB), and cirrus cloud thinning (CCT); these are described briefly in the textbox below. In the remainder of this document, we provide an overview of what is known about these approaches and the context for considering them, as well as answers to common questions.

Main approaches:

There are three main SRM approaches frequently discussed; we will very briefly introduce them here for context. Each will be revisited in more detail in future sections.

Stratospheric aerosol injection (SAI) is the best understood approach. By spraying very small liquid droplets or solid particles - "aerosols" - into the upper atmosphere, we could reflect sunlight back into space. Nature does this periodically through large volcanic eruptions; the eruption of Mt. Pinatubo in the Philippines in 1991, for example, led to roughly 0.4°C cooling the following year. In principle, the same effect could be achieved by using aircraft to loft material to the stratosphere, constantly replenishing the aerosol layer.

Marine cloud brightening (MCB) is less well understood. The idea behind MCB is that spraying sea salt aerosols into the right type of clouds over the ocean can increase their reflectivity or "brightness". The same effect can be seen in satellite images of ship tracks where the aerosol pollution from the ship results in a cloud that can persist for a week.

Cirrus cloud thinning (CCT) is the least understood method. CCT aims to thin or remove cirrus clouds, which have a net warming effect on the planet. By seeding particles into the upper troposphere at high latitudes, the optical thickness and lifetime of cirrus clouds could be reduced, cooling the planet.

Note that CCT does not technically reflect sunlight in the way SAI and MCB do, but because all three involve the artificial introduction of particles into the atmosphere to change radiative forcing and cool the planet, we group them together under the same umbrella for brevity. In this document, we use the term SRM to refer to all three methods.

1.2. Background, context, and the role of SRM

In equilibrium, the Earth receives energy from the sun, and radiates the same amount of energy back into space as infrared radiation. If these two are not equal, then the climate must warm or cool until it reaches a new equilibrium temperature. Over the past hundred years or so, human activity - mainly, the burning of fossil fuels - has increased the concentration of *greenhouse gases* in the atmosphere. These gases, mainly carbon dioxide (CO_2) but also methane (CH_4) and others, absorb some of the infrared radiation trying to leave the Earth. As a result, the Earth is less able to cool itself, and the planet has warmed in response. Consequences of this temperature rise, commonly referred to as *global warming* or *climate change*, are expected to include sea level rise due to melting ice, threats to agricultural productivity, increased frequency and severity of tropical storms and wildfires, and mass migrations from the warmest regions of the world as those areas become less hospitable. It is widely accepted by the global scientific community that many of these consequences represent unacceptable threats to human lives, health, and livelihoods. According to a 2020 report by the McKinsey Global Institute, natural disasters made more likely or more powerful by global warming have already led to hundreds of billions of dollars in damages and tens of thousands of deaths, and both the costs and deaths will grow as climate change worsens[1].

The only permanent, sustainable solution to global warming is to tackle the root of the problem by reducing the amount of CO_2 in the atmosphere. Unfortunately, this is easier said than done. The process of transitioning from a fossil-fuel-based economy to a renewable-based economy is challenging and expensive. Though many nations have taken steps to reduce their greenhouse gas emissions, progress in this area is well below the levels needed to meet the goals set out by the Paris Agreement. While no analogy is perfect, cutting emissions is like taking your foot off the gas when the car in front of you suddenly screeches to a halt: it would be silly not to, but it likely won't solve the problem by itself. In this case, the car is already going fast enough that taking one's foot off the accelerator won't prevent a crash; CO_2 levels are high enough that our emissions may well have already committed us to warming above 1.5°C and possibly above 2°C. The next step would be to hit the brakes, analogous to artificially removing carbon dioxide from the atmosphere. These technologies (called carbon dioxide removal, or CDR; see the box below) do exist; however, CDR needs to be scalable to gigatons per year, not have severe local impacts (like competition for food and clean water), and be cheap enough to be viable, and right now there is no single approach that meets all of these requirements. The challenge is that, while there is hope, we cannot know today how rapidly these ideas will be developed, and whether they will be developed in time to prevent the worst impacts of climate change. In other words, using the car analogy,

---

[1] *Climate risk and response: Physical hazards and socioeconomic impacts. McKinsey Global Institute, January 16, 2020.*

it is uncertain whether taking our foot off the gas and hitting the brakes will be sufficient to avoid an accident.

Carbon Dioxide Removal (CDR)

Currently, CO_2 levels in the atmosphere measure around 420 parts per million (ppm), and reducing or eliminating global warming will ultimately require reducing this concentration. In addition to reducing emissions, CO_2 already emitted can be artificially removed and stored someplace where it cannot easily escape back into the atmosphere; technologies that do this are known as Carbon Dioxide Removal (CDR) or Negative Emissions Technologies (NETs). The US National Academies' 2019 report on CDR, *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda,* calculates that the realistic global capacity for all combined NETs based on the current state of the technology is about 10 billion tons of CO_2 per year. Plausible pathways to keeping global warming under 2°C by the end of the century require on the order of 10 billion tons of CDR per year by mid-century, and 20/year by 2100; as such, these pathways require scaling all available means of CDR up to global capacity by mid-century, and technological breakthroughs and political commitments enabling a doubling of this capacity by the end of the century. Ultimately, this would require building an infrastructure commensurate in scale with our entire global energy infrastructure today, but without the direct economic incentives that created that infrastructure over the last century.

Here we give a brief overview of the different CDR methods, along with estimated costs (in \$/ton of CO_2) and storage capacities (in billions of tons CO_2/year). Quantities cited in this box come from the sources used by www.cdrprimer.org and the National Academies' 2019 report.

Afforestation, Reforestation, and other forest/agricultural management practices: increasing carbon uptake and storage in plants and soils through the planting of trees and management of forest and agricultural practices. The NAS Report quotes a practically achievable global capacity of 5.5 Gt/yr at "low" costs (\$0-20 per ton), and the CDR Primer cites a global capacity of up to 12.5 Gt/yr at a price of \$100/ton. The lower-cost estimates would be relatively inexpensive, but limited by total available land area, competition for land use with food production, and inability to completely implement changes to practices worldwide, in addition to concerns over permanence of storage.

Coastal Blue Carbon: management of mangroves, marshes, and other wetland regions to improve carbon uptake and storage. The NAS Report cites a current maximum global capacity (estimated based on analysis of the U.S.) of 0.13 Gt/yr at \$0-20 per ton, and the CDR Primer cites an estimate of 2.4-4.5 Gt/yr at \$100/ton or less; inexpensive, but limited by total wetland area.

Bio Energy with Carbon Capture and Storage (BECCS): the use of plant-based fuels, combined with the capture and storage of carbon released by their combustion. The NAS report cites a "safe" potential global capacity of 3.5-5.2 Gt/yr at \$20-100/ton without significant social/environmental impacts, and the CDR primer cites a range of 1-77 Gt/yr at a range of \$15-400/ton. Limited by cost and competition for use of land and biomass.

Direct Air Capture (DAC): the "scrubbing" of CO_2 from the open atmosphere via chemical reaction. Theoretically, capacity is bounded by infrastructure and energy draw, as well as storage capacity for captured CO_2. The NAS report cites current costs as ranging from \$100-600/ton, and that scaling up to 5 Gt/yr at \$100/ton would require significant global investment; the CDR primer estimates current costs at

$600/ton, with the possibility of realizing lower costs ($100/ton) with mass deployment. Limited primarily by high cost.

Carbon Mineralization: accelerating the natural process of "weathering," by which CO_2 bonds to certain minerals and is stored in rocks. The NAS report is reluctant to estimate storage potential due to lack of understanding, citing costs ranging from $20/ton to over $100/ton. The CDR primer cites a range of estimates for a dozen potential methods varying from less than 1 Mt/yr to 10 Gt/yr or more, with costs varying from $10-500/ton.

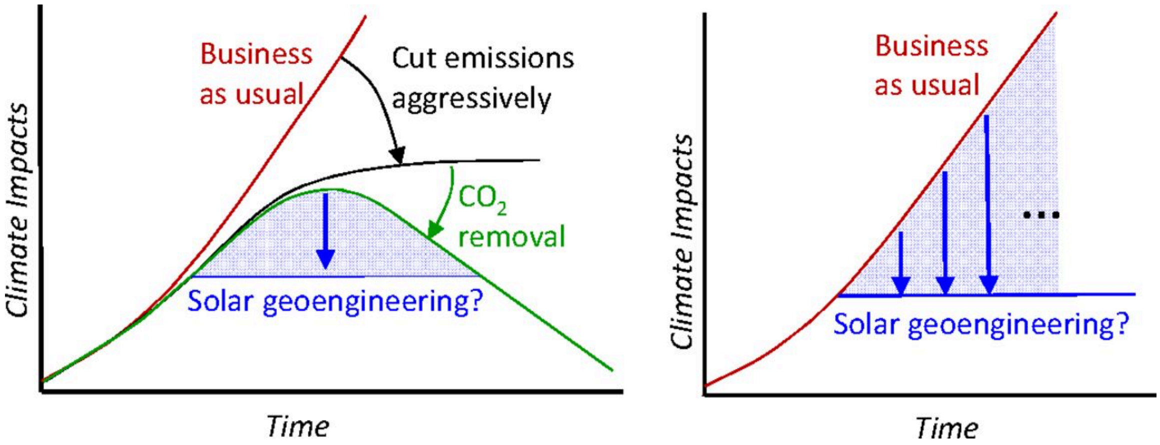

Figure 1 describes a "contextual framing" for the use of SRM - in other words, the role of SRM as part of a multi-faceted response to climate change. There are many variations on this sketch, which is often referred to as the "napkin diagram" (as it was originally sketched on a napkin by J. Shepherd in 2009).

Solar geoengineering - the artificial creation of mechanisms to reflect sunlight away from the earth - would reduce global warming, but it does not deal with the underlying problem of increased CO_2 levels. As such, following the car accident analogy, SRM is analogous to seatbelts and airbags: it doesn't slow the car down, but it may reduce the loss of lives and livelihoods in the event of a crash. As such, SRM should be considered only as a possible supplement to aggressive mitigation, and not as a substitute or excuse to delay it. For the same reason, SRM could only ever be a temporary endeavor; as demonstrated by the "napkin diagram" above, the amount of SRM required to reduce impacts scales with the amount of global warming we wish to offset, and as shown in the panel on the right, without cutting emissions and/or CO_2 removal, the amount of SRM needed to stave off climate impacts would keep increasing. However, as shown in the left panel, when used as a supplement to mitigation and CDR, solar geoengineering can act as a temporary salve to lessen the worst impacts of climate change. As discussed earlier, nobody knows how rapidly we can cut our emissions to zero, or how severe the climate impacts will be by the time we reach that point. SRM could lessen or even eliminate some of those impacts until we stabilize and reduce CO_2 levels.

1.3. The need for SRM research

While we know that SRM would cool the climate, these methods would not perfectly cancel the effects of increased greenhouse gases; some climate impacts would be reduced, others might not be, and there would be new risks introduced, including both purely physical effects as well as ethical, political, and governance concerns. Is it less risky to consider some limited amount of SRM, or to not consider it? This is not a straightforward question, as we need to consider not only physical climate impacts and uncertainties, but also the human dimension. The evidence to date from climate modeling suggests that if only the physical impacts were considered, it is plausible that SRM would reduce risk for most (and possibly all) people on the planet; however, “plausible” is an inadequate level of certainty for risks of this magnitude, and further research is needed to better understand the physical impacts of SRM. These physical uncertainties are compounded by complications from the human dimension, such as the potential for conflict over whether and how to deploy SRM, or the potential for “moral hazard”: will SRM lessen our resolve to cut emissions and invest in mitigation and CDR technology? These complications increase the need for research.

Just as one would never talk about chemotherapy without talking about cancer, or air bags without talking about car accidents, the risks of SRM must be considered not against the world as it is today, but against the risk of future climate change. The litany of concerns about climate change are well known, including droughts, forest fires, heat waves, stronger hurricanes, and increased sea level rise. As the Earth’s climate is pushed away from the regime that we are familiar with, in addition to the impacts getting worse, the uncertainty also gets worse; for example, Arctic permafrost contains vast amounts of stored organic material; as this permafrost thaws, the organic material is released as carbon dioxide and methane, and we don’t know how quickly this would accelerate climate change. We also don’t know, for example, how much global warming might trigger the collapse of the ice shelves that hold back Antarctic ice and subsequent catastrophic sea level rise. In total, our ability as an international community to limit global warming and its consequences to “acceptable” levels through mitigation and CDR alone requires four things, none of which are certain right now:

- *First,* it requires sufficient political will in every major GHG-emitting country in the world, sustained for multiple decades. Popular discourse often assumes both that this is the only requirement, and that if we want this to be true hard enough, then it will happen (as an aside, our use of the word “we” here glosses over several significant assumptions. The question of who “we” represents, and who has agency, is a recurring theme here; there are many different actors whose decisions matter).
- *Second,* it requires technology that, while plausible, doesn’t yet exist. A truly sustainable economy requires net-zero emission (either zero emission of CO_2, or reduced CO_2 emission balanced by CDR); while we know how to cut our emissions dramatically with existing technology (and we should be doing that today), we don’t yet have the scalable energy storage options required for 100% of our electricity to come from renewables, and it is difficult to remove the dependency of some sectors, like aviation, on fossil fuels. As discussed above, we also don’t have demonstrated scalable CDR approaches required to reach net-zero emissions. Net-zero solutions also need to be cheap enough to enable deployment not just in the developed world, but to allow the rest of the world to develop economically.
- *Third,* we need reasonable luck on climate sensitivity - in other words, there is uncertainty in the remaining carbon budget - how much additional CO2 can be emitted - in order to limit warming to 1.5°C or 2°C above pre-industrial conditions. According to the IPCC’s 6th Climate Assessment Report, typical emissions scenarios that are described as being sufficient to stabilize below the Paris

Agreement upper limit of 2°C actually only give a 2/3 probability of meeting that. No-one would ever accept those odds with these consequences in any other realm of life.

- *Fourth,* we need to gamble that the climate impacts from a given temperature are acceptable; there is no certainty that staying below 2°C or even the 1.5°C aspirational goal of the Paris Agreement will be sufficient to avoid significant risks. For example, there are already increases in extreme heat or flooding, and it is unclear how much warming will lead to even more severe impacts such as the collapse of Antarctic ice shelves and subsequent catastrophic sea level rise, or when, how much, and how quickly carbon will be released from melting permafrost.

While we can certainly hold out hope that all four of these conditions will be met, they are clearly not guaranteed, especially given that two of these four aren't even under human agency. This leaves us facing the reality that there is a significant possibility of a substantial and sustained overshoot of global temperature targets - in layman's terms, there is a real chance that we will miss the goal of limiting global warming to 1.5°C or even 2°C, that we will miss the goal by a lot, and that temperatures will remain over the target for a long time before technology evolves to the point where we can reduce them. Furthermore, even if we make the 1.5°C target, it is not guaranteed that the resulting impacts will be "acceptable". Different people can reasonably disagree about the probability of avoiding significant damages, but given the uncertainties involved, it is impossible to pretend that the probability is small.

Currently, not enough is known about any of the SRM methods described above to make informed decisions about their use, and further research is needed. The questions on the table today are not the hard ones of whether (and how, and where, and when) to deploy, nor the equally challenging question of how to make that decision, or even who gets to decide. Rather, the current question is whether or not to conduct additional research into SRM, and how that research should be governed. The US National Academies of Science, Engineering and Medicine (NASEM) recently released a report that reached the conclusion that further research was warranted, that the US should create a research program, and made recommendations on research governance. Note that as laid out in that report, while an international research program and international governance framework would be preferable, a US body can only legitimately make recommendations for actions taken by the US. The recommended research agenda includes not only research into the physical climate effects, but into the human dimensions of the challenge as well. If some form of SRM were ever implemented, the ideal scenario would be one in which every country agreed and cooperated. This aspiration seems unlikely today (though the 2015 Paris Agreement does provide some precedent). While the ideal world will likely continue to be just that – if we lived in an ideal world we would have already addressed our CO_2 emissions and wouldn't be faced with this conversation – it is appropriate that in parallel with the research, we must think about how decisions involving these ideas might get made. One aspect of this is simply capacity-building: making sure that people in every country understand the issues and how these decisions might affect them.

2. Stratospheric Aerosol Injection

2.1. Background

Some volcanic eruptions are powerful enough that some of their payload is ejected into the stratosphere, a layer of the atmosphere sufficiently stable so that added material can remain aloft for a year or longer. Sulfur dioxide (SO_2), one of the products of volcanic eruptions, reacts with OH and combines with water vapor to form tiny droplets, called *sulfate aerosols*. Sulfate aerosols scatter sunlight, reflecting some of it back to space, producing a net cooling effect on the planet. The 1991 eruption of Mt. Pinatubo in the Philippines ejected an estimated 10-20 million tons of sulfur dioxide (SO_2) into the stratosphere; this was sufficient to cool the Northern Hemisphere by 0.5-0.6°C and the entire planet by about 0.4°C. The basic premise of SAI is that the cooling effects of volcanic eruptions could be replicated by artificially introducing aerosols or their precursors (such as SO_2) into the stratosphere to reflect sunlight. Humans emit significantly more SO_2 into the troposphere every year than the eruption of Mt. Pinatubo did, and this does have a cooling effect, masking some of the warming from increased concentrations of atmospheric greenhouse gases. However, because the lifetime of aerosols in the stratosphere is so much longer, a relatively small (but still massive) amount of material added to the stratosphere can have a much more significant cooling effect than all of the tropospheric pollution.

SAI would likely be deployed by aircraft specially designed to carry heavy payloads from the surface into the stratosphere. SAI would not be an exact replica of a volcanic eruption; rather than deploying in a single burst, SAI would be continuous, requiring repeated flights to replenish the aerosol layer. This difference has consequences both in the details of the aerosol physics, and in the climate response; for example, monsoons are driven by the difference in temperature between land and ocean, but since the land cools more rapidly than the ocean, an “impulsive” forcing from a volcanic eruption will have a greater impact on monsoons than a corresponding gradual forcing from SAI. As a result, while volcanic eruptions provide an analog, care is needed in directly comparing the effects. Furthermore, unlike with an eruption, for SAI the location and amount of injection is a choice.

SAI could also be conducted at multiple locations simultaneously (see a discussion of the “design space” below). Most analyses of SAI consider injecting in the tropics. In the stratosphere, the air is rapidly mixed in longitude (east to west), and the Brewer-Dobson circulation carries air slowly toward the poles. Injecting near the equator means the aerosols would spread over the globe and provide relatively uniform cooling (a “global” strategy). In the tropics, the tropopause (the boundary between troposphere and stratosphere) is around 17 km (55,000 feet). However, because a higher injection altitude would produce a longer aerosol lifetime, an injection altitude several kilometers above the tropopause (20 km or higher) is typically assumed to ensure the aerosols wouldn’t immediately re-enter the troposphere. While this is higher than any existing aircraft is capable of lofting a payload of sufficient size (commercial airliners, for example, typically cruise between 35,000 and 41,000 feet), designing such an aircraft does not appear to present any fundamental challenges, and if the need arose, the design and construction of such aircraft could possibly be done in as little as a few years. Other strategies consider injecting at higher latitudes, such as 60°N (an “Arctic” or “high-latitude” strategy) or at both 60°N and 60°S; because of the direction of the Brewer-Dobson circulation, such a strategy would preferentially cool the poles, with relatively less effect on lower latitudes. While the aerosol lifetime would be considerably shorter, this would likely provide a greater impact on the Arctic per unit of injection, and could be used to preserve sea ice, ice sheet mass, and permafrost. However, it is not possible to cool only the Arctic (or any single region of the planet) without affecting the rest of the world; since cooling high latitudes would draw heat from the tropics towards the preferentially cooled poles, there would still be some effect on the tropics. At higher latitudes, the tropopause is lower (9-10 km closer to the poles), and sufficient altitude could thus be reached by existing aircraft for injection poleward of about 45°N or 45°S.

The aircraft payload is often assumed to consist of sulfur dioxide (SO_2) gas, the same substance ejected by volcanoes; over the next month or so, the SO_2 would oxidize into sulfate aerosol particles, just like after the eruption of Pinatubo (SO_2 is thus described as an aerosol precursor). Alternatively, it might also be possible to inject H_2SO_4 directly into the wake of an aircraft to produce aerosols right away and skip the 1-month oxidation time. The aerosols do not have to be sulfur, but at present, most climate model simulations of SAI use sulfur; the primary reason for this is risk - even though sulfate aerosols are a pollutant at ground level, because they are ejected by volcanoes, we can bound the uncertainty. Additionally, the natural occurrence of sulfur in the stratosphere enables model validation, as the behavior of sulfur in climate models can be compared with the real-life effects of volcanic eruptions (to an extent - as discussed above, volcanoes are not a perfect comparison to SAI; because of the different means of deployment, there are differences in the aerosol microphysics, an ongoing uncertainty in SAI research). Alternate aerosols such as calcite could provide advantages relative to sulfur, but because they do not naturally exist in the stratosphere, there is less observational evidence either to calibrate models or to bound “unknown unknowns”. The amount of sulfur that would need to be lifted would depend on how much cooling we wanted to achieve, as well as other goals such as sea ice (see below). In simulated SAI experiments, the amount of sulfur injected is often measured in teragrams (Tg), also called megatons (Mt); 1 Tg equals approximately 2.2 billion pounds. Roughly speaking, an injection of 10 Tg of SO_2 would produce about 1°C of cooling, and this can be scaled appropriately, with somewhat less efficiency for larger injections because the aerosols are more likely to clump together.

2.2 The “design space” of SAI

One of the first questions in considering the effects of SAI is the amount of cooling desired. More cooling would result in greater reduction in some climate impacts (such as greater recovery in Arctic sea ice, for example), but more cooling would require more injection and would increase risks from side effects. The amount of cooling is not the only question, though; the when and where are also important. As discussed above, preferentially cooling one portion of the planet (such as the Arctic) will still affect other parts of the planet because the cooling will influence heat transport. As such, any SAI deployment would be a global endeavor, affecting every person on earth. However, injecting different amounts in different places or in different seasons can still produce different effects on the climate. Therefore, SAI is not simply a “yes or no” problem, but rather a design problem - the choice is not simply whether to deploy, but rather at which location(s), how much to inject at each, and when to carry out the injection. This “design space” of SAI - the range of AOD patterns and climate effects we can achieve - is ultimately defined by the earth’s circulation. Here we discuss the limits of the design space.

Of the possible considerations for injection location(s) - latitudes, longitude, and altitude - only the choice of latitude is particularly important. The choice of longitude is unimportant for climate reasons, and the choice of altitude is largely constrained by technical limitations. East-west transport in the stratosphere happens on timescales of two weeks to a month, which is much shorter than the aerosol lifetime. As such, if we wish to inject over the equator (for example), then it doesn’t matter whether we inject over Brazil, the Congo, or Malaysia, or a little bit over all of them - the results would be statistically identical after a relatively short period of time as the aerosols are distributed around the globe. While a higher injection altitude might increase aerosol lifetime relative to a lower altitude, designing and building aircraft capable of carrying loads to increasingly high altitudes would be an increasingly difficult and expensive undertaking, and the increased “bang for your buck” may not be worth the investment. As such, longitude is arbitrary, and the ideal injection altitude is at least a few kilometers above the tropopause, but ultimately a trade-off with cost; this leaves injection latitude as the most influential location choice.

The choice of latitude is governed by the Brewer-Dobson circulation, which tends poleward in the stratosphere, with a strong seasonal dependence at lower latitudes. This has two major implications: firstly, aerosols in one hemisphere tend to stay in that hemisphere, meaning that injecting in one hemisphere will preferentially cool that hemisphere. This is important because of the implications for heat transport; injecting too much in one hemisphere would shift the Intertropical Convergence Zone (or *ITCZ;* the band of high precipitation near the equator where the trade winds converge) towards the cooler hemisphere, which could have adverse effects on tropical precipitation. To avoid this shift, any well-intentioned strategy would likely require either injecting close to the equator or injecting into both hemispheres. Secondly, because stratospheric aerosols tend to be transported poleward, it is possible to inject at high latitudes to preferentially cool the Arctic or Antarctic with relatively smaller impacts on the tropics. This is significant because, combined with the lower tropopause at high latitudes (which, as discussed above, results in reduced technical barriers), high-latitude SAI and low-latitude (or "global") SAI become somewhat more distinct topics. Injection at 60°N in one climate model was found to restore approximately twice as much summer sea ice per unit injection relative to a global strategy, and the preferential cooling at high latitudes would likely also result in greater preservation of permafrost and ice sheet mass. We comment more on the topic of Arctic SAI in section 5.4 of this document, "Arctic Impacts and High-Latitude SRM". However, even if a global strategy is being considered, injecting further away from the equator would still have benefits; injecting only near the equator could bring the *average* global temperature down to a desired goal, but doing this would over-cool the tropics and under-cool the poles (especially if too close to the tropics, as much of the aerosols could remain trapped in the tropical pipe). This could be mitigated by spreading the injection out over multiple latitudes, perhaps splitting it between 30°N, 15°N, 15°S, and 30°S (a combination used in several studies of "global" strategies to date), which would cool the globe more evenly.

Injecting at different times of the year will produce different effects on the climate, and therefore injection timing is also a factor to consider. The differences are smaller in the tropics, which receive similar amounts of sunlight year-round; "global" simulations injecting in the tropics commonly inject the same amount of SO_2 every month of the year. However, the poles receive much more sunlight in summer than in winter, and aerosols present in winter have little sunlight to reflect (their lifetime is also shorter). At high latitudes, SO_2 injected in spring will oxidize to sulfate aerosols before and during the summer, when the poles receive the most sunlight. As such, concentrating injection in the spring at high latitudes results in substantially more cooling per unit injection than injecting at a constant rate throughout the year, both near the poles and at lower latitudes. Therefore, high-latitude injections would likely be seasonal, while low-latitude or equatorial injections would be more likely to be year-round.

2.3. Climate impacts and uncertainties

We can say with certainty that SAI would cool the planet. However, the magnitude and distribution of temperature reduction corresponding to a given injection rate is less certain, as are the other effects of SAI. The only way to predict the effects of an intervention (without actually doing it) is to use a model, but models are of course not perfect. Sub-grid-cell processes need to be parameterized, often with limited observational data to constrain the parameter choices. Uncertainties in predicting the response can be conceptually divided into 1.) uncertainties in stratospheric and upper-tropospheric processes that determine the magnitude and spatial pattern of the "forcing", and 2.) uncertainties in how the climate system responds to that forcing.

Key stratospheric processes include aerosol microphysics and stratospheric transport. Aerosol droplets will occasionally collide to form larger droplets; larger droplets both have a shorter lifetime and are less efficient at scattering sunlight (because the scattering is a function of surface area, while the mass is proportional to volume; additionally, the extinction coefficient decreases with particle radius above a certain maximum). This coagulation process is thus important for understanding the efficacy - how much material would need to be added to achieve a given cooling. But because this happens at a spatial scale much smaller than the scale of a grid-cell in a climate model (typically of order 100km by 100km, and ~1km high in the stratosphere), these processes are parameterized: a global climate model does not - and never will - simulate the actual behavior of all of the molecules injected into the stratosphere. Additionally, as discussed above, volcanic eruptions (from which aerosol behavior in many climate models is calibrated) are not perfect analogues for SAI, which limits our ability to validate model behavior. The coagulation behaviour is likely nonlinear; doubling the rate at which material is added does not double the amount of solar energy reflected. Furthermore, observations after the eruption of Mt. Pinatubo in 1991, the most recent “large” eruption, are often insufficient - for example, the total amount of SO_2 injected into the stratosphere by the eruption is uncertain because the satellite instruments to measure it saturated. The stratospheric circulation is also both uncertain, and not perfectly represented in climate models; this leads to uncertainty in how the aerosols are distributed in the stratosphere and thus the spatial pattern of the resulting radiative forcing. Furthermore, sulfate aerosols would heat the lower stratosphere (because they absorb some of the infrared energy). This heating alters stratospheric circulation, altering the aerosol distribution; it also warms the tropopause, allowing more water vapour into the normally dry stratosphere. Because water vapour is a greenhouse gas, this would dampen the amount of cooling otherwise provided. The stratospheric heating might also affect upper tropospheric cirrus clouds, further altering the net forcing.

In addition to the uncertainty in the amount and spatial distribution of the forcing that results from a particular amount of injection at a particular location, how the climate responds to that forcing is also uncertain. Many of the effects of increased atmospheric GHG concentrations are directly related to the resultant warming, including the ability of a warmer world to hold more moisture (increasing both mean and extreme precipitation events), an increase in evaporative demand (affecting soil moisture and droughts), increased extreme heat, and melting of snow and ice. In addition to temperature reduction, we can thus confidently predict the sign of many of the other effects of SAI. However, while we know the sign of the effect, substantial uncertainty still remains in the magnitude, and in how the effects will interact with each other. For example, while we are confident that SAI will preserve Arctic sea ice extent and volume, the response of sea ice to SAI is extremely complex and would depend on many interacting mechanisms, including not only the direct forcing of the aerosols, but also humidity and temperature feedbacks, heat transport from lower latitudes, and cloud changes. As another example, aerosols scattering sunlight would heat the stratosphere, which would in turn slow down poleward transport by the Brewer-Dobson circulation. The implications of this change in circulation are not fully understood; stratospheric heating is likely related to changes in precipitation in some regions and to changes in the seasonal cycle of temperature in some places.

Ultimately, all SAI simulations must be done by climate models, as at the point where a field experiment is sufficiently large to produce a detectable signal on the climate response itself, it becomes a deployment. As such, there will always be uncertainties inherent to SAI that will never be resolved prior to a decision to deploy. To some extent, the uncertainty can be reduced by the use of feedback; we may not know, for example, the exact amount and distribution of cooling that an injection of 5 Tg/yr at the equator would produce in the real world, but we could adjust the injection quantity and latitude over time to push the cooling toward the amount and distribution we desire. As a result, it is not clear today how important it is to reduce any of the uncertainties in stratospheric processes, for example; the use of feedback might be sufficient to ensure that projected outcomes are nearly independent of some uncertainties. There will of course always be unknowns, as well as “unknown unknowns” representing impacts or interactions we have not considered. However, there are also significant unknowns associated with *not* deploying SAI and potentially allowing the climate to warm to unprecedented levels. The goal of SAI research is to reduce the uncertainty to the point where we can reasonably compare the benefits and risks of SAI to the impacts and risks of *not* deploying SAI, and determine whether it should be seriously considered as part of a multi-faceted response to climate change.

3. Marine Cloud Brightening

Under the right conditions, smokestack pollution from ships at sea causes vessels to leave behind bright “tracks” of clouds that can persist for days. This happens because aerosols in the ship’s exhaust serve as cloud condensation nuclei: droplets of water coalesce around the aerosols, forming cloud particles. When aerosols contribute to cloud formation in this way, the clouds are “brighter” because the same amount of water vapor is distributed across more, smaller droplets, resulting in a higher cloud surface area and thus higher reflectivity. This mechanism, known as the Twomey Effect, is the principle behind marine cloud brightening (MCB): by deliberately seeding marine skies with cloud condensation nuclei, which could be done by spraying salt water into the atmosphere (rather than with pollution), we could brighten existing clouds, reflecting more sunlight back into space to cool the planet.

MCB is less well understood than SAI. The primary reason for this is the low level of understanding of how clouds and water droplets interact with aerosols; laboratory observations of these interactions at micro-scale cannot easily be translated to the complexity of the real world at climate-wide scales, and the IPCC describes cloud-aerosol interactions as “one of the current major challenges in climate modeling in general”. While we know there are areas and conditions where MCB will “work,” in the sense of increasing the amount of reflection, there are also meteorological conditions where adding aerosols might have the opposite effect (by causing a cloud to rain out earlier, for example), and the exact conditions where adding aerosols increases reflectivity vs. decreasing it are not well understood. Ultimately, it is unclear today where, when, and how much we can increase cloud albedo. As such, the maximum forcing achievable through MCB is not certain. Humans currently emit significant aerosol pollution, though most of this is not in the regions likely to be most susceptible to increases in cloud cover. Current estimates place existing forcing due to worldwide cloud-aerosol interactions at -0.9 W/m^2, but with considerable uncertainty.

MCB could be achieved using salt water, reducing concerns associated with the material itself (relative to sulfur, typically assumed for SAI). While MCB would be implemented only 1-3 km above the surface, which means that it could be deployed by ships, the technology needed to deploy would still involve challenges. Developing spray nozzles that generate the right size of salt aerosols is challenging, and the number of vessels needed, and the extent to which their spray strategy might need to be finely tuned based on local conditions, remains unclear. Unlike SAI, MCB could be applied locally; while the stratospheric aerosols of SAI would spread globally and can persist for a year, MCB could be applied over a smaller region, meaning MCB could be either a “regional” approach or a “global” approach. However, as heat is transported through the atmosphere, the resultant cooling would not be local even in the former case, and there would still be transboundary effects. Additionally, the right type of clouds only exist over perhaps 10% of the planet's surface, and so if the goal is global cooling, a stronger effect would be needed in those smaller areas. This could lead to more spatially heterogeneous outcomes, though the details are unclear.

4. Cirrus Cloud Thinning (CCT)

Most clouds have a net cooling effect on the atmosphere: low, warm clouds emit approximately the same amount of longwave radiation which they absorb, and therefore "trap" little upwelling longwave radiation, but they do reflect some of the incoming sunlight. However, high, cold cirrus (ice) clouds that form in the upper troposphere (6-13 km) emit much less longwave than they absorb, and therefore trap outgoing longwave radiation. These cirrus clouds, found in mid- to high-latitudes in particular, trap more longwave radiation than the amount of sunlight they reflect and are therefore a net warming influence on the earth. The goal of cirrus cloud thinning (CCT) is to artificially thin or remove cirrus clouds, thus lessening their warming influence.

Cirrus clouds can form in two ways: homogeneous freezing and heterogeneous freezing. Heterogeneous freezing occurs where there are sufficient ice nuclei, typically dust. In the absence of sufficient ice nuclei, homogeneous freezing can occur; this requires colder temperatures and higher relative humidity, and results in smaller ice crystals than heterogeneous freezing. The relative fraction of cirrus clouds that form via homogeneous freezing is not certain, but it could be the dominant mechanism in many places. By introducing *ice nucleating particles* (INPs) in regions where homogeneous freezing occurs, we could cause a larger fraction of cirrus clouds to form through heterogeneous freezing. The ice particles could grow larger, and there would be fewer of them; the same amount of water spread over less surface area results in less optical thickness, and more infrared radiation could escape. Additionally, the heavier particles would fall out of the atmosphere more quickly, reducing cloud lifetime. The result would be fewer, thinner cirrus clouds, and a net cooling effect on the planet.

A potential advantage of CCT (over SAI or MCB) is that CCT is the only one of the three which increases the amount of outgoing longwave radiation rather than reducing incoming shortwave radiation, thus affecting the atmosphere through a more similar mechanism to greenhouse gases. That said, CCT would still not have the same spatial or seasonal pattern of radiative forcing as well-mixed atmospheric greenhouse gases, and would thus still lead to changes in regional temperature and precipitation patterns. More critically CCT is the least well understood of the three methods considered here. CCT is based on the assumption that cirrus clouds form primarily through homogeneous freezing, but the extent to which this is true is an open question in atmospheric science. Thus far, CCT has only been explored in climate models, as there is no natural analog for it; in climate models, the percent of natural cirrus that form homogeneously varies hugely depending on the physics, even in the same model. Due to these uncertainties, there is a possibility that an attempt at CCT would actually warm the planet; overseeding - supersaturating a region where no more cirrus would naturally form - would result in the formation of additional cirrus rather than thinning them. Additionally, the mechanism of CCT itself is not well understood; more information is needed regarding the concentration, temperature, and supersaturation dependence of cirrus ice nucleating particles. Finally, the logistics of CCT deployment are also unclear, i.e. what material, particle size, and dispersion mechanism would be most effective, and how well a deployment could target regions in which the heterogeneously-formed cirrus that was created replaced homogeneously-formed cirrus that otherwise would have occurred, rather than forming new cirrus where there otherwise would have been none. As such, one cannot say confidently how well CCT would work, or what the climate effects would be if it did.

The overall potential of CCT to reduce global warming is not large. It is estimated that cirrus clouds currently warm the planet by 5 W/m^2; this provides a hard upper bound on the change CCT could bring about. CCT would be most effective at high latitudes, but cirrus are most common at low latitudes, and thus the maximum achievable effect of CCT would be a reduction of about 2 W/m^2. The effect would be spatially diverse, and little is known about how the climate or ecosystems would respond. Due to the uncertainties discussed above, it is unclear what the realistically achievable effect is, and it differs between climate models.

Due to the low understanding of CCT and the relatively small amount of forcing it could achieve, we cannot realistically expect to rely on CCT to offset global warming in the next few decades.

5. Frequently Asked Questions

In this section, we discuss many of the common questions or concerns relating to SRM, including both some societal dimensions (such as potential for moral hazard) and physical climate impacts (such as ozone loss or agricultural yield). Some of the information contained in this section has already been touched on above, but we include it here in greater depth for the sake of easy-to-find organization. Some of the issues described below apply to all three SRM methods (i.e. moral hazard), while others are more relevant to some methods than others and may not apply to some methods at all (e.g., ozone depletion for SAI).

5.1. Framing

The answer to any question about SRM has to be “it depends” - not just on the method of SRM considered, but also on how it is being deployed, and against what climate background. There is considerable misinformation and confusion regarding the expected physical impacts of SRM; much of this happens because the questions are framed incorrectly - in other words, the questions asked by science and policy are often different, and it is important to understand why.

In order for a question about SRM to be framed correctly (and therefore answered properly), three things need to be considered: the nature of the strategy, the background climate, and the comparison being made. Therefore, if someone were to ask, “how will SRM affect precipitation,” the answer would depend on each of the following:

1. What is the SRM strategy being considered? For example, in a deployment of SAI, the impacts will depend on the latitude, altitude, and season of injection, as well as the material being used. Adding aerosols at high northern latitudes in the spring to cool the Arctic will affect at least the rest of the hemisphere, but the results would be quite different than a “symmetric” strategy injecting near both poles, or a “global” strategy including injections at low latitudes. A deployment of MCB or CCT would have similar considerations.
2. What is the background climate state when SRM is deployed? In addition to the type of SRM and strategy being considered, the impacts will also depend on how much cooling is desired, which depends on both the year and the emission levels of CO_2 and other greenhouse gases. For

example, climate scientists conducting a simulation of SRM may use a high-emissions background scenario with high warming (such as RCP8.5 or SSP5-8.5), and then counteract it with lots of SRM, thus achieving a high "signal-to-noise ratio" (SNR). In this case, the high SNR would be useful in discerning the effects of SRM from the high natural variability of the earth's climate; however, if the changes were assumed to be descriptive of changes under a more moderate emissions scenario (and therefore more moderate cooling), then the conclusions could be misleading from a policy perspective. Likewise, a different group of scientists more concerned with policy or economics might conduct an SRM simulation starting in 2040 with a background they consider more consistent with current emission and mitigation trends, such as RCP4.5 or SSP2-4.5, and thus use less SRM. Such a simulation would offer less information about changes to the climate because the SNR is lower, but it might offer more information about policy, or about how climate impacts in a specific country might change under this more "realistic" SRM simulation.

3. What comparison is being made: are we comparing SRM to not using SRM - two hypothetical future worlds under the same assumptions about future emissions of CO_2 and other greenhouse gases (for example, the year 2100 with SRM to 2100 without SRM) - or are we comparing the effects of SRM to an earlier world with the same global mean temperature and neither further increases in greenhouse gases nor use of SRM (for example, 2100 with SRM to 2040 without SRM)? Scientific papers often focus on the latter to understand how SRM affects the climate differently from greenhouse gases, but from a policy perspective, this comparison would only be relevant if achieving the same temperature target through emissions cuts alone was being considered as a practical alternative choice (depending on the temperature target used in any particular climate model simulation of SRM, this is likely impossible). The former comparison is appropriate if the choice of whether or not to use SRM is assumed completely decoupled from subsequent choices on emission reduction, but reality is potentially more complex and more uncertain (see discussion of "moral hazard" below).

Hopefully, this brief discussion demonstrates the complexity of SRM, and that it is therefore impossible to make blanket binary statements on any question. The three points above only cover the physical dimension; if the human dimension is also considered, the discussion becomes much more complicated.

5.2 Temperature and related changes

As noted in the discussion on climate response for SAI, projecting the response to any of the methods discussed here requires a climate model, and is uncertain due to both uncertainties in predicting the "forcing", and in predicting how the climate system responds to that forcing. However, any of the methods would cool the climate, and as a result we know the sign of many impacts, e.g. on cryosphere (sea ice, permafrost, contributions to SLR from Greenland and Antarctic melt), or thermosteric sea level rise, or heat waves, and there is a reasonable physical basis for reducing many precipitation changes as well as noted in next section. However, in addition to depending on the method, and how the method is implemented, the details of how SRM would affect some impacts are unclear. For example, the net effect on the Greenland ice sheet might depend on shifts in precipitation, or shifts in cloud cover, or shifts in the seasonal cycle of temperature for example.

The radiative forcing from well-mixed greenhouse gases is relatively uniform across latitudes and across seasons, the radiative forcing from SRM will be less so. The forcing from SAI is inherently relatively "smooth" over space and time because the lifetime of the aerosols in the stratosphere means that they have time to spread considerably. However, that is not true for either MCB or CCT, both of which would have much larger variation in the forcing from region to region (e.g., MCB is likely effective over only about 10% of Earth's surface, requiring comparatively larger forcing over those smaller regions to achieve a global effect). The atmosphere transports heat effectively - reducing sunlight over a patch of ocean will start cooling that patch, but that will change the flow of air, bringing in warmer air from elsewhere, and spreading the cooling out over a larger region. The climate response to MCB applied over 10% of the Earth's surface will not be as uniform as the climate response to SAI, but it will be more uniform than the forcing itself; the effect of this spatially heterogeneous forcing is currently unclear.

An additional consequence of the heat transport is that with SAI in particular, you cannot just cool one country or one region. MCB would likely allow more localized effects, though there would still be transboundary effects from any deployment large enough to be useful.

5.3. Precipitation changes

CO_2 and other greenhouse gases increase temperatures everywhere, with larger increases over land and larger increases in the Arctic in particular. Because warmer air can hold more water (the relationship is governed by the Clausius-Clapeyron relation and the August-Roche-Magnus approximation, which says that the atmosphere can hold about 7% more water for every 1°C rise in temperature), this also leads to changes in precipitation: global mean precipitation increases, and different regions see varying increases and decreases (with a typical "wet-get-wetter and dry-get-drier" pattern). Most of these precipitation changes are a direct result of increased temperatures, and thus to first order, SRM that cools the planet would reverse these. However, the different mechanism (reflecting sunlight vs trapping heat), the different spatial and seasonal patterns, and additional physical aspects (such as stratospheric heating for SAI) mean that no form of SRM would simply reverse climate change, and there would be residual changes in both temperature (uneven cooling) and precipitation (shift of regional patterns). SRM would "over-compensate" the changes in global mean precipitation relative to the changes in global mean temperature (i.e., relative to the timeline of global warming and resultant precipitation increase, SRM reduces precipitation more than it reduces temperature), but this is not particularly relevant because this decrease in global mean precipitation is a small component of the projected residual regional changes.

Regional precipitation changes due to SAI have been particularly emphasized for tropical monsoon regions (India and Africa), where the amount and seasonality of the precipitation is particularly important for agriculture and ecosystems in general. In India, several climate model simulations have shown that the increase in precipitation under climate change would be over-compensated by SAI that aims to fully compensate global mean temperature changes; of course, a smaller amount of SAI might fully compensate for the change in precipitation but leave some residual warming. (Note that this result is not robust across different climate model simulations.) It has also been observed that there were reductions in precipitation after large volcanic eruptions, but in this respect a volcanic eruption is not a good analogue for SAI; the rapid increase in forcing produced by large volcanic eruption would cool continents much faster than the ocean, and since monsoonal precipitations are largely driven by the temperature difference between the ocean and land, the changes in monsoonal circulation due to a volcanic eruption would be much different than for a deployment of SRM, where the forcing would be likely to increase much more gradually with time.

One case where significant precipitation changes could occur is if SRM were deployed in only a single hemisphere. Cooling one hemisphere more than the other would change atmospheric heat transport, and the location of the Intertropical Convergence Zone (ITCZ), a band of heavy precipitation around the tropics; a similar effect has been noted after some volcanic eruptions. This possibility indicates that there is at least one "bad" way to deploy SRM - deploy in a single hemisphere. Of course, this effect could be corrected by deploying an appropriate amount in the opposite hemisphere as well.

More generally, changes in precipitation patterns can be explored with climate models, but as discussed above, several key observations need to be taken into account when interpreting the results: first, as noted above, what are the changes relative to? Is it merely an inability of SRM to restore the climate to the previous state, or are the changes actually produced by some features of SRM, such as the stratospheric warming produced by SAI? Second, are the changes even large enough to be detectable in the presence of natural variability under moderate amounts of SRM cooling? And third, any changes will depend on the particular SRM strategy that is being simulated: how much cooling is it producing? Is the strategy deliberately harmful (i.e. a large cooling just in one hemisphere, that would shift the Intertropical Convergence Zone and thus sensibly modify precipitation patterns in the tropics)? A recent study has shown that even just injecting sulfate in different seasons of the year might produce small modification in the changes to the Indian monsoon: while this is not a certainty, it opens up the possibility for understanding if particular strategies might avoid some of the regional changes to precipitation altogether.

Different methods of SRM would have different impacts on precipitation patterns that are as yet unclear. Heating of the lower stratosphere from aerosols such as sulfate will have some influence on atmospheric circulation and hence on precipitation patterns. The much more spatially heterogeneous radiative forcing from MCB or CCT, relative to SAI, would introduce different shifts in regional precipitation patterns; the fact that MCB would only be applied over the ocean also influences the resulting precipitation changes.

Finally, it is important to emphasize that while much is written about regional changes in mean precipitation, these changes are not always the most relevant variable to consider. Agriculture and ecosystem impacts will be more affected by overall water availability, roughly the balance between precipitation and evaporation. Since increased temperatures result in increased evaporation, the cooling from SRM affects this term as well; increased CO_2 levels in the atmosphere also reduce plant water needs, all else being equal, and thus precipitation changes alone cannot be considered in isolation from other factors. Similarly, longer-term drought risks might be more correlated with soil moisture, which roughly tracks precipitation minus evaporation rather than precipitation itself, though this is also affected by the frequency of extreme precipitation events (which lead to greater runoff). And risks from flooding, for example, are driven more by precipitation extremes.

To conclude, further research is needed, but overall there is a general agreement that the regional changes in precipitation would typically be smaller than those produced by an increase in surface temperatures; this follows both from climate modeling, and from basic physics. There is no modeling or physical basis to assert precipitation changes as a show-stopper for SRM applications, but shifts in precipitation clearly need to be evaluated as part of a more comprehensive effort in understanding impacts more thoroughly.

5.4. Arctic impacts and high-latitude SRM

As discussed previously, the different methods of SRM can, to differing extents, be molded to focus on one part of the world. In particular, SRM focused on high latitudes (henceforth "Arctic SRM," although most of what follows could also apply to the Antarctic) will preferentially cool the poles relative to the tropics. This could be beneficial because it would preserve certain aspects of the Arctic climate - sea ice, land ice, and permafrost - with lesser effects on lower latitudes.

Of the three methods of SRM, the effects of Arctic SAI are best understood, although little research has been done on SAI focused exclusively on the Arctic. The Arctic gets most of its energy through heat transport from lower latitudes, and Injecting directly over the poles has little effect because surface albedo is already high and aerosol lifetime is low. Rather, it may be more effective to inject at high-mid latitudes (i.e. 60°N) and reduce heat transport into the Arctic. As described above, any meaningful deployment of SAI will affect every person on the planet; SAI at low latitudes will still cool the Arctic, and high-latitude SAI will still affect the rest of the world. However, it is known that Arctic SAI recovers sea ice, preserves permafrost, and reduces ice sheet melt more efficiently than SAI at lower latitudes.

High latitudes have a much higher seasonal variation in the amount of sunlight they receive. While the effect is most pronounced at the poles, 60°N still receives 10 times more sunlight in the summer than in the winter. As such, the seasonality of injection is much more important for Arctic SAI than for low-latitude SAI; there is little point to producing aerosols in the winter months when there is little to no sunlight to reflect. It takes the average injected SO_2 molecule perhaps 3-6 weeks to oxidize and become an aerosol, meaning that injections concentrated in the spring months will result in aerosols that are present throughout the period of peak insolation in the summer. A study of 12 Tg of SO_2/yr injected at 60°N restored twice as much September sea ice when concentrating the injection in March, April and May compared to the same quantity distributed throughout the whole year.

Antarctic SAI would probably help to preserve the Antarctic ice shelf, but this area has been explored substantially less than Arctic SAI.

Because the tropopause is lower at high latitudes, the barriers to implementation for Arctic SAI are lower than for low latitudes. Namely, there currently exist aircraft capable of carrying payloads to the altitudes required for Arctic SAI.

The extent to which MCB could be focused on the Arctic is unclear. While there is relatively little open ocean north of 60°N, MCB focused on the middle latitudes (i.e., perhaps 20-40°N) would reduce heat transport into the Arctic. Since CCT is theorized to be most effective at mid-high and high latitudes, it is plausible that CCT would preferentially cool the Arctic.

5.5. Ozone

A known "side-effect" of using sulfate for SAI would be a reduction in high-latitude stratospheric ozone. Stratospheric ozone provides protection from ultraviolet radiation from the sun, as it absorbs the most harmful part of the spectrum (UV-C, which can penetrate deeply into organic tissue and is energetic enough to damage DNA) and partially absorbs a less harmful part (UV-A and UV-B) before they reach the surface. In the recent past, the increase in chlorine (Cl) and bromine (Br) gases in the atmosphere (resulting from the use of CFCs) has reduced this stratospheric ozone layer everywhere, but especially close to the poles. The Montreal Protocols banning those compounds have managed to stave off a further reduction, and ozone levels have already started a slow recovery, but the long atmospheric lifetime of these gases means that stratospheric ozone is not expected to return to pre-industrial levels until the second half of the century (this would take even longer, except for the upper atmospheric cooling produced by the CO_2 increase; this is projected to lead to a “super-recovery” of ozone later in the century).

The presence of sulfate aerosols in the stratosphere would modify some of the recovery processes of ozone in various ways, and the overall effect would depend also on how much Cl and Br gases are left in the atmosphere, and hence on when SAI was started. The overall influence is complicated by changes in different processes of ozone production and destruction, as well as changes in ozone circulation, and will depend on the spatial and seasonal distribution of the injected aerosols. The stratospheric warming in the tropics produced by the absorption of infrared radiation would accelerate some of the processes of ozone destruction, but at the same time, the presence of the aerosol would inhibit some other destruction processes that depend on the quantity of gaseous nitrogen (NO_x) and HO_x. Overall, changes to tropical ozone concentrations would likely be small. Closer to the poles, where ozone destruction by Cl and Br is preponderant, ozone concentrations would be reduced, and the magnitude of this reduction would depend on the amount of sulfate injected and transported therein.

Sulfate is not the only possible choice of aerosol material for SAI, and there are other materials that would be expected to have a smaller effect on ozone. MCB and CCT would not directly affect ozone.

5.6. Acid rain deposition

Industrial processes are responsible for a large portion of all sulfate emitted in the 20th century; this sulfate oxidizes into sulfate aerosols that mostly deposit close to where they were emitted. Since most emissions happen close to population centers, the problem of acid rain has been a significant environmental concern. Together with acidifying soils, the sulfate particles produced by industrial processes are very small and very harmful - mostly falling in the categories of Particulate Matter (PM2.5 and PM10). Some natural processes (effusive volcanoes, carbonyl sulfide from various land processes, dimethyl sulfide from the oceans) also result in tropospheric sulfate.

The use of sulfur for SAI would increase global sulfur pollution. However, due to the much longer lifetime of sulfate particles in the stratosphere (one year, instead of few weeks), only a small fraction of the sulfate that is currently emitted from the ground would be needed to obtain a much larger radiative effect: industrial pollution releases approximately 100 million tons of sulfur dioxide every year, but the amount of SO_2 needed for SAI to cool the climate by roughly 1°C is about 10 million tons, only one-tenth of this number. Additionally, due to the longer residence time in the stratosphere, the particles falling down from the stratosphere would be much larger than those produced by industrial processes and would thus not be in those sizes where they are more harmful for respiratory systems. Furthermore, the stratospheric circulation would spread it around much more evenly: for this reason, the localized increase of sulfate deposition from SAI would likely be very small compared to the overall deposition (and most of it would fall over the oceans). However, due to the more efficient spreading, there may be areas where usually little deposition from human sources arrives that might see a much larger increase, increasing soil acidification in areas previously considered pristine.

The use of materials other than sulfur for SAI would not cause acid rain, but the injected material would still return to the surface, with potential effects on ecosystems, infrastructure, and human health. MCB and CCT would not cause increases in acid rain (though the environmental effects of the materials used would still need to be evaluated; even increased airborne salt may have some effects on coastal ecosystems).

5.7. Photosynthesis

The primary goal of SRM is to reduce incoming solar radiation. To offset all anthropogenic warming to date, approximately 1% of sunlight would need to be reflected. While this is a huge amount of energy from a human perspective and would be a significant undertaking, it is a relatively small fraction from the perspective of photosynthesis.

To first order, one might expect that a 1% decrease in sunlight reaching the surface would result in a 1% decrease in photosynthesis, but there are other relevant mechanisms; SAI not only reflects sunlight back into space, but also scatters light, which means that although direct sunlight decreases, diffuse light increases. The increase in diffuse light as a fraction of total incoming sunlight could be a net benefit for plants, as the non-direct portion of the incoming sunlight reaches otherwise shaded leaves, and thus allows for higher total photosynthesis (while currently sunlit leaves are often already saturated and a small reduction in direct light may not even reduce their photosynthesis). Thus, on average, one shouldn't expect a drop in photosynthesis due to SAI, but the changing ratio of direct and diffuse light may benefit different plants and thus affect ecosystems in unknown ways. Further, in any given location there might be many overlapping effects determining the overall sign of the changes: temperature, soil humidity, and cloud cover changes that modify the amount of diffuse radiation are just some examples of factors that might determine whether the productivity of plants would increase or decrease under SAI.

For MCB, because the forcing would only be applied over a relatively small fraction of the Earth's surface (~10% or less), then if the goal was to achieve global cooling, a much more significant reduction in total sunlight at the surface would be needed in those regions, with currently unknown effects on ocean biological productivity.

CCT does not directly affect incoming sunlight.

5.8. Agriculture

As discussed above (see "photosynthesis"), there are multiple factors determining whether SRM would increase or decrease crop productivity. The overall effect on crop yields would depend on CO_2 levels, heat stress, and water availability (not precipitation directly, but the difference between precipitation and evaporative demand); the latter also depends on irrigation potential. Crop yields are potentially influenced by the ability to alter what seed strains are planted as well. In a future with high warming from greenhouse gases, it is likely that the majority of crops would benefit from an SRM intervention reducing surface temperatures, even with a small reduction in incoming sunlight. Compared with the agricultural yields in current (or pre-industrial) conditions, the changes induced by how SRM might affect the system differently from increased greenhouse gases are likely small compared to the projected increase in output due to CO_2 fertilization. However, much more research is needed to understand the effect of SRM on various types of crops, which all have different sensitivities to water availability, sunlight, temperature and more.

5.9. Termination shock

If a deployment of SRM were abruptly terminated, then the climate would rapidly warm back towards roughly the same state that it would have been in had SRM never been deployed, with most of the warming occurring over a few years. The consequences of that rapid warming are generally assumed to be worse for ecosystems than had SRM never been deployed. Termination is clearly a risk, but how much of a risk depends on how much cooling is being done. In scenarios that gradually ramp up the forcing, it might be decades after deployment starts before the masked-warming is sufficient to make it too risky to suddenly stop. To first order, the risk does not depend on how the cooling is being done, but only on how much; there would be a very slightly slower initial warming from terminating SAI relative to terminating MCB simply due to the aerosol lifetime, but this is probably only relevant for reducing the risk associated with a 1-year interruption in activity.

The risk of termination provides a strong incentive if one country develops the technology for other countries to do so too.

Note that if, after many decades of deployment, some physical climate impact became evident that justified stopping deployment, then it is of course possible to gradually ramp down the amount of cooling rather than suddenly terminating.

5.10. Moral hazard (mitigation deterrence)

There is clearly a risk that use of SRM (and even the awareness that SRM is possible) could lead to less emphasis on mitigation, and this is perhaps the number one reason why some are opposed to SRM or even research into SRM. The long-term climate consequences if there is some significant mitigation deterrence could conceivably result in worse outcomes. Predicting future societal responses to SRM is of course a guess, and would still be a guess at the time of a hypothetical decision to deploy or not. There has been some small-scale research that is inconclusive even with regards to the sign of the effect (that is, some people react with “wow, if that’s what we’re going to be faced with if we don’t cut emissions, we better cut emissions faster”). Discussions of moral hazard are also fraught with the ethical questions associated with the idea of “we” as if humanity were a single actor – because the people who suffer from climate change impacts (and thus might benefit from SRM) aren’t the same people who might use SRM as an excuse to keep emitting, as well as questions of withholding knowledge because we don’t trust future people to use it.

Ultimately, while there is the possibility that an SRM deployment would reduce short-term incentives and motivation to reach net-zero emissions, the choice of “emissions cuts, CDR, or SRM” is not a binary one (any more than one needs to choose between wearing seat belts and using the brakes). Most notably, there is no evidence that an immediate worldwide halt to SRM research would substantially accelerate progress in mitigation or CDR, or that such a halt would be the difference between a world with unacceptable climate consequences and one without. In other words, as discussed above (see section 1.3, “the need for SRM research”), we may already be past the point where mitigation and CDR alone are enough to prevent unacceptable climate consequences, and halting SRM research won’t change that.

5.11. Public perception

While the idea of SRM has been discussed at the international level (the United Nations’ 2021 AR6 report on climate change briefly addresses SRM), no nation or group of nations currently has an official stance on SRM or SRM research. Likewise, different environmental groups have differing stances on SRM. While some work has been done to examine public perceptions of SRM, it is difficult to measure public attitudes accurately due to low levels of awareness. Additionally, given the nuanced nature of the topic, responses are influenced by the wording of the questions; for example, most respondents tend to prefer CDR when presented as a "solution" to climate change, but as discussed above, the choice between CDR and SRM is not a binary one. Different methods of SRM are not mutually exclusive, and as such, framing a question as "would you prefer CDR, SAI, or MCB" is not particularly useful, particularly given that the answer to that question should be influenced by efficacy and side-effects that are not currently known.

Perception of SRM (and SRM research) is influenced by perception of the risks posed by climate change compared to the risks posed by SRM. It is commonly thought that political conservatives and greenhouse gas interests support SRM as a means to continue polluting without consequence, but there is little evidence for this; rather, SRM research is most likely to be supported by those who consider climate change to be an important issue, and by those with higher levels of trust in science and scientists. Additionally, there is evidence that those living in developing countries in the Global South (who are more vulnerable to the consequences of climate change) are slightly more supportive of SRM research than those living in developed countries. Broadly speaking, respondents tend to oppose imminent SRM deployment, tentatively oppose field research, and more broadly support laboratory and modeling research. A common theme in survey responses is the idea of manipulating nature; many people express discomfort in the idea of humanity deliberately modifying the world on a global scale in such a way, an attitude which could stem from either practical concerns or more spiritualistic reasons. In the spring and summer of 2021, Harvard University prepared for a small field experiment in northern Sweden; however, the Indigenous Sámi people objected to the experiment, with one Indigenous leader arguing that humans should learn to live with their environment rather than manipulate it, and the experiment was ultimately postponed. As the technology evolves, public discourse surrounding it will likewise evolve, and continuing to understand public perceptions of SRM will be critical in the years to come.

5.12. Implementation challenges

For SAI, the most immediate challenge to implementation is lofting the material, almost certainly involving aircraft. At the equator, the tropopause is 17 km above the surface, and the desired injection altitude for low- and mid-latitude injections would probably be at least 20 km. The deployment vehicle would need to carry a payload to that altitude and maintain that altitude while that payload is dispersed; currently, there exists no aircraft capable of doing this. However, studies suggest that overcoming this obstacle would not be fundamentally difficult for a state or organization with sufficient resources (see "who could deploy?" below). A strategy focused on the Arctic might be easier; the tropopause is lower at high latitudes, and an injection at 60°N (for example) could inject at 15 km. This could feasibly be done with existing aircraft.

For CCT, the largest barrier to deployment is arguably the state of the research. Setting aside the relatively low upper bound of realistically achievable forcing, more needs to be known about the formation of cirrus clouds to ensure that an attempt at CCT does not end up warming the planet instead of cooling it. The deployment altitude for CCT would be much lower than for SAI, as cirrus clouds form in the upper troposphere, and the amount of material that would need to be deployed is likely much lower than for SAI. However, little research has been done as to the mechanism by which CCT might be deployed (i.e., piloted aircraft, autonomous drones, etc.). As such, while implementing CCT might be practically easier than SAI, a CCT deployment mechanism would have to be not only implemented, but designed first.

Unlike SAI or CCT, MCB can be deployed by ships; a 2008 Royal Society article included designs for a fleet of MCB-deploying ocean vessels, including specifications for rotors and turbines to spray seawater into the clouds (see "Sea-going hardware for the cloud albedo method of reversing global warming"). However, for a global-scale deployment, a large number of ships would be required, along with the facilities and people to maintain them. Deployment would also require development of spray nozzles that can produce the desired size distribution.

The other major barrier to deployment is the political one. All forms of SRM are an extremely contentious issue, and wildly varying opinions about SRM can be found across the general public, environmental groups, and climate scientists. Since any meaningful SRM deployment would affect everyone on the planet, ideally, everyone on the planet would get a say; however, given that we cannot even agree on whether SRM should be researched (or, in some cases, that global warming is a problem at all!), the odds of reaching a global consensus one way or the other are near zero, at least in the near term. As global warming continues to worsen, it is likely that SRM will become a greater topic of discussion on the international stage; regardless, the only way a nation or group of nations could deploy is if they believe the benefits of SRM outweigh the physical and political risks. On the other hand, given that SRM can be deployed from more or less anywhere in the world, once somebody decides to deploy SRM, it is very difficult to stop them barring direct military intervention. Of course, political and economic sanctions could be implemented, but a deployer may be big enough to weather the political storm, or vulnerable enough to climate change that they have nothing else to lose.

Cost is a lesser consideration. Most estimates for SAI (there is little cost estimate data for MCB, and little to none for CCT) fall in the range of single billions to tens of billions of USD per year. This would likely be relatively easy for most developed nations and international unions to afford (there is debate over whether smaller countries or even wealthy individuals or corporations could independently conduct SRM - see "who could deploy?" below). Note that the costs of SRM would include not only the direct costs of manufacturing, materials, and dispersal, but also budgets for monitoring and analysis. Additional costs could include compensation for nations or indigenous groups disproportionately impacted by SRM or who did not have a voice in the decision to deploy, extra money budgeted for healthcare costs related to SAI particle deposition or ozone depletion, or even legal costs should the effort to deploy be challenged in court.

5.13. Who could deploy?

The question of who could deploy SRM is largely answered by who could overcome the political, financial, and technological barriers discussed in the previous section (see above).

The low direct costs raise the question of whether even moderately small countries or wealthy individuals could deploy. Any deployment that actually affected climate (as opposed to a symbolic effort) would involve a sustained effort lasting for years. For SAI, for example, this would presumably involve many aircraft flights a day. As a last resort, this would of course be trivial to shut down unless there was at least the implicit backing of a major power. Even a wealthy individual needs access to airbases and a supply chain, and maintaining these would also require implicit backing of the nation that operations were taking place in.

For high-latitude strategies, aircraft exist today that could reach sufficient altitude; anyone could buy used business jets. For lower-latitude strategies the tropopause is much higher and as described above no existing aircraft could be used and a new aircraft would be required. Designing and building the aircraft itself is not particularly difficult. However, there are relatively few global suppliers of aircraft engines that could operate at the required altitudes. Developing a new engine from scratch would take years and could be beyond the capabilities of all but a relatively few number of countries.

The technology behind MCB or CCT would present few challenges, and could be deployed at least at small scale by nearly any country. For MCB any effort large enough to have a global influence would involve operations in international waters.

5.14. Detecting a deployment

If someone chose to deploy SAI, it would be trivial to detect. Long before the stratospheric aerosol levels were large enough to actually affect the climate, they would be large enough to be easily detected above the background variability of the stratosphere. Furthermore, the scale of the operation required may be small enough to be economically viable, but it would still be a large enough operation that the regular flights would be readily observed.

A small-scale MCB deployment within territorial waters of an individual country might more easily be kept secret, but it would also not have global ramifications. Any MCB effort large enough to have global ramifications would likely be trivially detectable just due to the sheer number of vessels required.

5.15. What about other methods?

Other methods of reflecting sunlight have been proposed, such as changing the reflectivity of land surfaces, but these methods are unlikely to be capable of a reduction of the required magnitude; for example, achieving a 1% reduction in sunlight would require a reflective area of about 8 million square kilometers, roughly equal to the area of the continental USA. Similarly, a space-based mirror between the Earth and Sun would need an area of approximately 2 million square kilometers; for example, this would require constructing an average of 100 square kilometers a day for the next 50 years. We do not further discuss either of these ideas here.