


Doing research on environmental issues or responses is usually an easy call for policy-makers and gets wide political support, even if there's disagreement what to do about the issue. But there is now one big exception: research on solar geoengineering (SG). SG would cool the Earth, temporarily and imperfectly offsetting some of the climate effects of elevated greenhouse gases, by reflecting a bit (around 0.5% to 1%) of incoming sunlight. The way to do this that now looks most promising would be to spray a fine mist of reflective aerosols in the upper atmosphere. I've written here before with [introductions to SG](#), what is known about it from the [limited research that has been done](#), and the [shape of the major arguments](#) about it. I won't repeat those here.



Potential solar geoengineering methods. Source: Carbonbrief

Now, as climate alarm grows, the debate over SG is changing in two ways. First, it's getting [much more media coverage](#), including serious treatments in [prominent news](#) and [commentary outlets](#), [several new books](#), and a central role in [several recent works of science-fiction](#) or science-adjacent fiction. Second, long-stuck arguments about research on SG may be shifting. A hopeful sign of this movement is the thoughtful [report](#) released in March by a Committee of the National Academies, which recommended starting a US Federal research program and gave detailed suggestions on research priorities, funding, and governance arrangements. A less hopeful recent sign was the latest encounter between activists and scientific groups over small proposed field experiments, also last Spring. These experiments can give more realistic insight than modeling or lab studies into how (and whether) SG would work, with what environmental impacts. Experiments proposed so far

are all so tiny in physical scale and strength of perturbation that they have vanishingly small environmental impacts, yet have still met strong opposition. This recent conflict is over a Harvard group's [proposal](#) to launch a balloon from a space launch station in northern Sweden, to release about 1 kilogram of material into the stratosphere. The [flight](#) was [postponed after objections](#) from several Swedish environmental groups and the Swedish part of the Saami Council, a trans-national organization of the indigenous Saami people.

A recent issue of Science Magazine addressed the controversy over SG research, with two




separate “Policy Forum” essays, plus a guest Editorial by me. The two [Policy Forum](#) essays offer specific proposals for high-value research in hope of broadening support for research. [My editorial](#) reflected on the meaning and motivations of the fight over research, and suggested how underlying concerns about research should be addressed. I argue that research opponents raise concerns that are valid, but do not justify the actions they propose, blocking research or subjecting it to such stringent and aspirational governance requirements as to be functionally equivalent to prohibition. Indeed, I argue that the longer the opponents continue to succeed at stalling research, the more they increase the risk of the harms and injustices that they fear.

Underlying the fight over research are divergent views on whether using SG could ever be acceptable. Because the prospect of its use is entirely hypothetical, there is little evidence to resolve these disagreements and the debate is rather strange. On the opposing side, the strongest statements are [categorical](#) and [unconditional](#): It must [never be used](#), under any conditions, on account of its imperfect climate correction, its environmental side effects, and its time-scale mismatch with greenhouse gases (atmospheric lifetimes of months versus centuries to millennia), plus various indirect risks related to how it might be used recklessly

or unjustly, or over-relied on to the detriment of emissions cuts or other essential climate responses. On the other, more favorable, side, there aren't any similarly categorical statements. The [strongest supporters note](#) that SG might augment other responses to reduce climate risks in ways they alone cannot, and might [particularly](#) benefit the most vulnerable, with [risks](#) and [limitations](#) that can be mitigated to acceptable levels. Emphasis on "might."

With such deep uncertainties and high stakes, you might expect strong agreement to support research. Points in favor of research include the virtual certainty that some future climate decisions will be influenced by the prospects for SG, good or ill; the possibility that it may, subject to various uncertainties and conditions, be beneficial; and the practical impossibility of ensuring that no one ever uses it. These suggest the need for both scientific and technical research, and analysis and consultations regarding how to govern SG prudently and legitimately, and how to anticipate and manage potential future demands to use it.

But proposals to expand SG research have thus far gained little support and attracted sharp opposition – lately including not just attacks on scientific groups proposing experiments, but also on the [recent NAS Committee report](#).



Arguments against research vary, but most have a common core. They start with the limitations and risks of using SG noted above. But because no one is proposing to use SG now or soon, these concerns don't imply opposing research without additional steps of reasoning. Some opponents don't bother with these extra steps: they confound research with deployment, or suggest that all objections to the latter apply

unchanged to the former. Others distinguish the two activities and extend their arguments to research in two ways. For those convinced that deployment [could never](#) be [warranted](#), at any time or under any conditions, these extra steps are easy because [research](#) is [superfluous](#). Why bother researching something you already know with certainty would never be acceptable to do?

Others argue that doing research will trigger [social](#) lock-in processes that promote unexamined continuance, expansion, and eventual deployment, even if unbiased assessment of risks and benefits says this is unwarranted. This is basically a slippery-slope argument, which is usually not applied to all SG research, but to field experiments that would make active environmental perturbations, no matter how small. The most common form of these arguments asserts that research alone would lead to over-confident reliance on SG and so [impair support](#) for the essential first-priority responses of cutting emissions and adaptation. The mechanism most frequently invoked for this is "[normalization](#):" Treating SG as a normal response - one to be considered, critically scrutinized, assessed for benefits and risks along with other responses - makes its use more likely even if unwarranted by the evidence.

A third group of arguments asserts that solar geoengineering research itself [does direct](#) harm. Research can do harm. This is obvious in such areas as use of human subjects in medical research, and could clearly be the case if future SG research proposals involved large-scale, non-trivial environmental modifications. But for current and near-term research proposals, whose direct risks and environmental impacts are negligible, these arguments of direct harm from research unpack into either a variant of the lock-claim, or the risible claim that research does direct harm by virtue of its being done by the wrong researchers, institutions, or funders. At best, this latter argument asserts that democratic values related to process and participation have priority over scientific values in the design and implementation of research; in other cases, it is just *ad hominem* attack dressed in progressive symbolism.

Here's where the debate gets confusing. With the argument already shifted from "could it ever be OK to use this stuff" to "is it OK to research it," it then (mostly) shifted again, to what kind of research governance would be required to make proceeding with research acceptable. Scientific research is not conducted in an anarchy. All research is [governed](#), mostly through normal processes of peer review, research-program management, and compliance with applicable laws and regulations. Some research areas that pose special risks - e.g., research involving human subjects, or research using radioactive materials or infectious organisms - is subject to additional control, based on the obvious and direct harms it could cause. For SG, however, the key question is what additional research governance is called for on account of the novelty of the potential interventions, the concern

they trigger, and the possibility that research could lead to harms, via some combination of lock-in (so it can't be stopped or controlled once it starts), and the prospect of people misusing the knowledge or capability that comes from it (thus undermining emissions cuts, or leading to future use of SG that is reckless, partisan, or unjust).

Even the strongest supporters of research [endorse some additional oversight](#) and care: the question is how much, of what kind. There is virtual unanimity in support of research controls that include: limits to the scale and intensity of environmental field experiments; enhanced requirements for disclosure of research parties, aims, methods, and results; international coordination and information sharing; citizen consultations on the broader implications of SG, conducted in parallel with and linked to the research programs; and firm preset break-points, at which research programs would be re-assessed to consider changes including tighter restrictions or cessation. The specifics are not settled, but there is not much disagreement on any of these matters.

So where is the fight? Mainly it's about proposals for research controls much stronger than these. Typical demands include a legally binding [international](#) research [governance](#) system, which must be adopted before research proceeds or research programs are established; and public engagement processes for every research program or activity, which would be empowered to require changes in proposed research activities or prohibit them. These often invoke idealized governance norms based on universal direct democracy. The strongest forms of these governance proposals – binding international research governance, adopted before research is done or programs established, including consultative processes empowered to prohibit proposed activities – would carry such extreme burdens of delay, uncertainty, and cost, that they are for practical purposes equivalent to a prohibition on the research activities subject to them.

This is where my [editorial](#) comes in. Assessing the arguments over SG research and its governance, it argues that the concerns that motivate opposition are valid and require responses — but that they don't warrant the proposed remedy, either shutting down research or governing it so tightly that it's equivalent to a prohibition. It bases this argument on three prominent and systematic characteristics of the arguments against research.

First, they all rely on mechanisms of risk that are indirect, arising from bad decisions policy-makers or other actors could make based on knowledge or capability derived from the research. Concerns of this type have arisen in many areas of research and technology, notably going back to the first [Asilomar conference](#) on recombinant DNA experiments in 1973. In some cases, such indirect risks based on bad uses of knowledge or capability make

a compelling case for controls. For example, research pursuing new capabilities that would make some dangerous new form of weapon widely available would make a compelling case for tight governance. A similar recent case is the gain-of-function research on viruses with human pandemic potential that was subject to a [funding moratorium](#) between 2014 and 2017 - although this was controversial, because the same research may help develop vaccines or treatments. Compared to these cases, claims that SG research could lead to harms rely on weaker and more indirect mechanisms of harm: If this capability is developed, decision-makers might over-rely on it to the detriment of other necessary responses, use it in a panic, or try to pursue unjust or destabilizing regional advantage.

Because these harms are all mediated by political behavior, it is impossible to conclusively refute claims that they “could” happen. Yes, these bad things could happen. But how likely are they, and how much can their likelihood be reduced? For that matter, how much more likely are they if research proceeds than if it does not, since suppressing research now cannot guarantee against future use? The only way to learn more about the likelihood and severity of these risks, and potential means to limit them, is to do the research, in parallel with research, planning, and consultation on how to build the needed governance capacity. These activities might lead to development of better understanding and control of both direct and indirect risks, but they are also precisely the activities that opponents identify as causing the risk. With this double-bind, the decision to reject research cannot be the result of a reasoned argument based on evidence; it is, rather, an axiom, not subject to elaboration or examination.

The second common characteristic of arguments against research is that they rely on mechanisms of harm that operate at global scale. Yet they are advanced to support local opponents rejecting proposed research in their location, even when the research presents no local risks or impacts and complies with all relevant laws and regulations. The proposed Harvard experiment in Sweden gives an illustration. Of the objections raised to the experiment, only one made any claim to being local, and it was purely procedural and so minor as to appear pretextual. The researchers did not do local consultation before a preliminary balloon flight that would have tested equipment but introduce no material into the environment. They instead [planned \(PDF\)](#) local consultation before a second, experimental flight that would introduce ~ 1 -2 kg of material in the stratosphere. Other than this lack of sufficiently early consultation, all other objections concerned indirect, global-scale mechanisms of harm. Yet opponents claimed standing to make these objections, and demand the project be delayed or cancelled, based on their connection to the place where the experiment was to be conducted. Other than the procedural objection, they [did not every try](#) to claim local risks or harms. And such claims of local harm would be

particularly hard to sustain in this location, since the two proposed balloon flights would have been from [a space facility](#) that has conducted 650 balloon launches and 550 launches since its establishment in 1978, without similar objections. Opponents later reinforced their reliance on global-scale risks by stating that such experiments should not be conducted [anywhere](#). This amounts to a claim for a veto over proposed research [conducted anywhere](#), sustained by a thin claim of special local standing but without (to analogize to US law) any claim of particularized local injury. It is notable that a parallel claim for standing could be made for those who might benefit from the research and its results. This disparity is especially sharp given the prospect - suggested by climate-model studies of SG, although by no means proven - that those [most vulnerable](#) to [climate impacts](#) might stand to [benefit most](#) from using SG. The only difference is which stance - the research may proceed, or it may not - gets the benefit of being the default presumption, relative to which the other stance requires justification. When the evidence is so thin that neither stance can mount a persuasive justification, then the implicit default always wins. As a wise colleague of mine once said in a very different context — fights over empirical research on the effectiveness of social policy — whoever owns the null hypothesis wins.

Third, arguments against SG research take a stance of extreme precaution, rejecting research based on a long list of harms - direct or indirect, likely or unlikely - that “could” follow from the research. But in demanding extreme precaution for decisions about SG research, they fail to apply similar precaution, or any at all, to other related risks that are implicated in these decisions. Extreme precaution can be justified, even required, under certain conditions, but the conditions are quite restrictive. The activity in question must present a significant risk of severe harm, which cannot be mitigated except by not doing the activity; and the alternative path taken if the activity is rejected must be known with high confidence to be acceptable.

None of these conditions applies to SG research. The severity and likelihood of the identified risks following from SG research, and whether they can be effectively mitigated if research proceeds, are simply unknown. And if SG research does not proceed, one possible alternative path is that some future actor [still deploys](#) SG in response to severe future climate impacts. In this case, blocking research will not have prevented future use, but only ensured such use takes place under conditions of greater ignorance about SG methods and their effects, and is thus more likely to be ineffective or dangerous. Or if blocking research does preclude future use, then the alternative is whatever climate trajectory would apply if no one ever thought of SG, moderated by whatever trajectory of other responses - emissions cuts, atmospheric greenhouse-gas removals, and adaptation - is achieved.

There is no basis for confidence that this alternative path will be acceptable, let alone

widely preferred to the distinct constellation of uncertain possibilities if SG research does proceed. Climate-change risks are severe and growing, even if uncertain in their precise character, timing, and incidence. And as the window to achieve the Paris targets gets ever narrower, the intensity and rapidity with which other, non-SG responses must be deployed to achieve these increasingly hold the prospect of severe disruptions. The roughly 50% cuts in global emissions in ten years needed to meet the 1.5°C target with limited overshoot (i.e., global heating doesn't go much above 1.5°C or stay there long, before large-scale carbon removals pull it back down) represents a more rapid energy-system transformation than has ever been achieved even within a single nation. Yet even the far-too-weak mitigation measures enacted thus far have triggered serious political opposition in many jurisdictions. As the set of acceptable energy technologies gets drawn tighter, as many climate activists propose - e.g., excluding nuclear power, carbon capture, large-scale biofuels or large new hydroelectric facilities - the more extensive the changes in behavior required in addition to technological changes, and the more difficult and disruptive the transition is likely to be. CDR and adaptation, also essential, both face their own limits. The enormous and rapid expansion of CDR needed to bolster 1.5°C scenarios appears uncertain both as to feasible rates of deployment and scaling, and potential impacts and opposition. And serious anticipatory adaptation measures, even for risks as well known as coastal inundation from sea-level rise, remain few, weak, and highly contentious. One need look no further than the California Coastal Commission's multi-year efforts to get vulnerable coastal towns to do rational advance planning of how they will handle the inevitable retreat, to disabuse you of the notion, "oh, we'll easily adapt to the changes."

These warnings do not in any way disparage the need for extreme efforts on emissions cuts, atmospheric removals, and adaptation. All of these must be ramped up rapidly, major progress has been achieved in some areas, and the prospects for further expansion are substantial. Rather, I am only arguing the possibility that these may fail to deliver the scale and speed of transformation needed to ensure that SG would never be warranted. These other responses may fall short, shifting more risk to climate-change impacts. Or they may be pursued with an extremity or clumsiness that bring different risks of harm or disruption. The required changes to world energy systems are often characterized as transformative, even revolutionary — and this is correct. But recognizing their revolutionary character surely carries the responsibility to recognize that revolutions often fail, and often carry terrible harms, risks, and injustices even when they succeed. These prospects do not warrant giving up on other responses, or blithely relying on deployment of SG to save us. But nor can assumptions of near-magical optimism about the ease, scalability, and benign impacts of these other responses justify refusing even to learn about SG when it might offer great reduction in climate harm, especially for the most vulnerable.



A highly precautionary stance is warranted in making decisions about climate response. But applying extreme precaution selectively to one potential response, even blocking research needed to understand it, will simply shift risk: maybe less risk from SG deployment (although even that is uncertain), but greater risk from climate impacts and from forcing other responses beyond points of social feasibility. The overall effect of suppressing SG research may thus be to increase total climate-change risks — and possibly also to make the distribution of risks even more unjust, by shifting risk toward those most at risk from climate impacts, most lacking in resources needed for effective adaptation, and most likely to have their interests overlooked in global strategies for extreme emissions cuts and carbon removal.

There's no way to quantify this authoritatively, but I share what seems to be the widespread sense that the range of possible climate futures, even in the near term, is growing wider; and that this range of possibilities increasingly includes degrees of possible societal disruption that are difficult to imagine, which because of their extremity are barely getting on the radar of serious debate. This characterization applies to the range of plausible futures for climate change overall, and also for how SG might relate to those. To refuse to admit the possibility that some state or other powerful actors might use it; or the possibility that using it might even, on balance, point toward human and ecological futures less bad than the alternatives; or even worse, to imagine that one's own political and normative commitments that assure you that using it would be bad somehow ends the discussion, all strike me as dangerous delusion. This doesn't mean that use of SG is either likely or desirable. It certainly does not mean anyone is planning to use SG. It merely means that these possibilities cannot be responsibly ignored.

In the short term, responsibly considering these possibilities means doing the research to learn whether and how they might work, how effectively they can delay or reduce impending harms of climate change, what impacts and risks they may create, and how these can be managed or reduced. It also means thinking hard in advance about effective mechanisms to control the use of these technologies - somehow maintaining a balance between the aspiration to advance human welfare, justice, and environmental protection; and pragmatic judgments of what is possible to achieve in something resembling the geopolitical world as it is.

In terms of concrete implications for near-term action, these considerations all bring me quite close to where the National Academies Committee came down.

- Research on solar geoengineering must proceed - including small field experiments insofar as these are judged scientifically valuable to address questions not well

addressed in to model, lab-bench, or passive-observational studies.

- The bulk of the research should be in public programs, strategically directed to provide knowledge needed to inform climate risk assessments and decisions, in jurisdictions with strong institutions and norms that support research accountability and public benefits. The NASEM Committee recommendation for a US national program under the GCRP is sound, but the United States should not be the only nation doing this research. Parallel programs should be established in other nations in both the Global North and South.
- The design and management of these research programs should incorporate governance that goes beyond normal peer review and compliance with applicable regulations
- These additional research governance requirements should address: transparency and disclosure regarding project actors, aims, methods, and results; disclosure and/or control of private interests and IP; limits on the scale and intensity of field experiments; and periodic requirements to re-assess and adjust programs, including the possibility of cessation. These last two governance elements specifically aim to reduce the risk of thoughtless continuance or expansion, or lock-in leading to unconsidered full-scale use.
- International consultations should begin immediately, but informally. Initially, these will mainly link research programs, focusing on coordination of research plans and exchange of information about methods, activities, and results. These may also provide one early vehicle to begin consultations on the more severe international governance problems that would be posed by future deployment proposals. Formal international institutions or requirements should not, however, be pre-conditions for national research programs. These are too slow and unwieldy to establish or to subsequently change, and what they would need to do depends on what is learned from the research.
- Informal international consultations on the larger-scale, longer-term governance challenges of SG should also start promptly. These governance deliberations must interact and exchange information with research programs, but the processes must remain distinct. The governance challenges of SG use are sufficiently novel that there is a need for explorations that are broad in both participation and substantive ideas explored. This suggests the value of multiple parallel processes, diverse in form, content, and participation (experts, officials, citizens). These governance explorations should be linked to and informed by parallel developments in scientific and technical research, but not be primarily scientific or technical in focus. They should provide vehicles to consider the totality of climate response, the role of SG within it, the requirements for effective, prudent, and legitimate international control of SG, and the

broad societal meaning and potential consequences of the unprecedented step to active, intentional interventions in the global environment. Detailed design of these processes will be challenging, and is one of the points left under-specified by the NASEM Committee report, on which further ideas and analysis and experimentation are needed.

The opponents of SG research are correct in asserting the unprecedented nature of the proposed technological interventions, their broad implications, and the need for public information and deliberation. The potential risks and harms they raise represent valid concerns. But because these are all potential outcomes mediated by bad political decisions, they are impossible to authoritatively refute — yet that is the standard of demonstration the opponents seek, by virtue of identifying these as a long list of bad things that “could” happen. If risks of misuse of SG existed in a vacuum, with no linkage to other related risks, the stance of suppressing SG research, whether through explicit prohibition or governance as a proxy for prohibition, might be valid. But they don’t. Given the gravity of related climate-change risks, remaining uncertainties, and potential limits to other responses, there is no responsible alternative to pursuing research and governance consultations for SG. This must be done with caution, with awareness of the associated risks, and with common-sense controls to limit these risks, even recognizing their unavoidable imperfection. But it must be done.