Who May Geoengineer II: justice and delegation for solar radiation management under conditions of geopolitical rivalry

Patrick Taylor Smith
Principia Advisory
Georgetown University

Brief Summary/Outline:

Solar radiation management (henceforth SRM) is especially fraught from the standpoint of global justice. An extensive and diverse set of uncertainties—from the distribution of side effects to the political and social consequences of deployment—make the use of SRM a considerable risk. And ideally, we would want a global governance regime embodying both procedural and substantive legitimacy to enable collective self-determination about the nature, appropriateness, and extent of SRM use in the long term. Unfortunately, such a global governance regime is unlikely to develop, especially given rising levels of geopolitical rivalry. Despite this, I have argued elsewhere that there is a secondary pathway by which SRM use might be justified. Some actors may have a revolutionary normative permission to engage in SRM. However, that permission only applies to agents that a) purport to represent those subject to considerable climate injustice b) are generally free of climate injustice and c) can reasonably be understood to legitimately speak for oppressed people. An important caveat of this normative permission to engage in risky SRM is that it requires that the potential users of SRM to *succeed* in mitigating climate injustice in order for their use to be justified. "Good intentions" are insufficient. This is both a political and an engineering problem: we—and, by extension, potential SRM actors—do not know whether SRM will work and, relatedly yet distinctly, whether it will be allowed to work.

This permission structure creates a normative problem: the *permission-capability asymmetry*. The agents that are most likely to have a normative permission to use SRM—those subject to the most severe climate injustice, such as the Marshall Islands—are also the least likely to have the ability to engage in SRM. The very elements that make them vulnerable to climate harms—poverty, lack of engineering know-how, lack of access to knowledge/patents, state incapacity, and so on—are also likely to undermine their ability to engage in effective SRM. And conversely, those nations and actors that have most benefitted from the burning of fossil fuels—and, therefore, have the wealth, research institutions, engineering capacity and so on—are those that can develop and eventually deploy SRM. So, there is a fundamental tension: those that most plausibly could use SRM successfully to mitigate climate harms—if SRM could in fact do this—are those with the weakest normative permission to do so while those with the strongest normative permission are the least likely to be successful.

This is a serious problem in at least two ways. First, this tension increases the likelihood that uses of SRM will be unjust. Risks of SRM harms will be imposed on vulnerable people around the world by agents who lack the permission to do so justly. Furthermore, powerful agents may use SRM as a cheap but unjust substitute for the other duties they have to respond to the climate injustice, such as providing adaptation funding or mitigating their emissions, in ways that are inapplicable to weaker and more vulnerable climate actors. Second, there will inevitably be bias if powerful actors use SRM unilaterally on their own behalf. Even if we grant that these powerful actors have good intentions and want to mitigate global climate risk for everyone, it is inevitable that actors will be biased in favor of their perceived interests. Thus, we would expect that the details of SRM deployment by powerful actors to be suboptimal in relation to the interests of those with the relevant normative permissions. For example, research increasingly suggests that injection site location plays a significant role in the potential effects—both positive and negative—of SRM, but it is likely that agents will select the injection site that reduces reputational, legal, pollical, financial, or environmental risks to themselves rather than sites that will maximize benefits to others.

There seem to be four potential responses to this asymmetry. First, we could simply ban the deployment of SRM. Second, we could be indifferent about who has a normative permission to engage in SRM. Let's set these first two aside for the moment as both are, in a way, disrespectful of the entitlement that oppressed peoples may have to engage in SRM. They are both ways of saying that it is simply *too hard* to figure out how to make that entitlement work so it should be abandoned.

The third strategy is *capacity-building*. Here, we shrink the capacity gap by empowering those with the relevant normative permissions. By building up their SRM related engineering and financial wherewithal, they could decide for themselves how to whether to exercise their normative permission. This is, obviously, an attractive strategy, especially since scientific capacity is fungible. If potential SRMers decide not to use their permission, they can use their newfound capacity to serve their other interests. Also, having those with the normative permission do the actual work of deploying the SRM will mitigate concerns about bias. As a result, I have argued that SRM researchers have a duty to take steps to build research and engineering capacity around the world on both justice and epistemic grounds.

There are nonetheless significant limitations to the capacity-building strategy. First, there is practical concern that it is very difficult to develop home-grown scientific capacity in impoverished states. If development was straightforward, it would already be done. This is especially true if there are considerable ethnic, class, linguistic or other cleavages in the target society. And for nations that are sufficiently small, it is simply infeasible to create a domestic engineering industry that could do the work. What's more, there appears to be little appetite in the global north to engage in a crash scientific, economic, and engineering program to provide those nations with climate intervention capacities that are highly controversial. What's more, developing the relevant capacities in only one or a few of the nations with the normative permissions risks its own kind of bias: how will other nations—who are in the same oppressed position, disagree on the chosen deployment, but lack their own capacity—be represented?

A fourth potential solution—with both its unique benefits and risks—is that of *delegation*. This solution requires those with the relevant normative permission to outsource the exercising of their entitlement to an outside agent or group of agents that are better placed to meet practical and normative requirements, such as monitoring and safety. This is commonplace in the realm of collective defense, where an organization such as NATO is used to delegate normative permissions of self-defense to states with shared values and superior capabilities. Estonia cannot protect itself from Russia alone, but it can delegate its self-defense rights to a collective security community with the needed capabilities. This strategy has obvious advantages in terms of resource allocation. Those parts of the world with an excess of resources can be placed at the disposal of those parts of the world that need them. Furthermore, delegation makes coalition building amongst the oppressed actors much easier. We do not need to develop multiple SRM programs but rather a coalition of interested nations that can interact with those agents possessing the money and know-how to engage in SRM, assuming that the research program is successful.

Nonetheless, key questions remain to be answered.

- **What is delegated:** We should not treat research/deployment/monitoring as a single undifferentiated responsibility that is either delegated or not. It is entirely possible to separate these responsibilities—and to create divisions within each category—such that only some elements are delegated, and others are withheld. In the defense space, a nation can delegate protection against ground invasion or nuclear deterrence while nonetheless taking direct responsibility for air defense or intelligence or cybersecurity.
- **Who is the delegate:** The nature of the delegate matters. The risks and benefits of security guarantees change if they are the result of a bilateral treaty with a superpower (such as the strategically ambiguous US guarantees of Taiwanese sovereignty), a collective guarantee from a multilateral institution (such as NATO), or a specific capability provided by a private entity (such as Starlink or BMD interceptors).

There are, despite these promising analogies, significant concerns with the delegation strategy. These include:

- **Principal-agent problems:** Legitimate delegation typically involves a system that ensures that the delegated agent will be appropriately attuned to the interests of the delegator. For example, a legal system that enforces fiduciary obligations for those who manage my trust helps ensure that their decisions align with my interests. Yet, such a system seems to be lacking in the context of SRM-delegation.
- **Public authority/Private Action**: There are government functions that cannot be legitimately outsourced to private agents if they are essential to citizenry engaging in collective self-determination. So, for example, a state may plausibly delegate functions where there is a clear possibility of full values substitution by a regulated private actor. Yet, there are essential, political functions—such as voting or criminal justice—where the private delegation would undermine the status of citizens as free and equal.

- **Lack of Control**: The delegated agent must be under "substantive control" of the delegator. However, the inequality of the relationship between the less powerful delegator and the more powerful delegate makes such control difficult to maintain. Can, for example, the Marshall Islands really exercise control over their delegated agent if that agent is Tesla/SpaceX or the European Union?

So, what are the necessary preconditions for legitimate/just delegation of SRM research and eventual deployment, given these concerns:

- **Alignment:** The structure and nature of the delegation mechanism need to include elements—such as authority over funds, who makes the final determination that relevant objectives have been met or concerning the elements of potential—that ensure that the delegate aligns with the delegator. These could include a tranche release structure, regular independent audits, and transparent decision-making.
- **Nature of delegate institution:** It is unlikely that delegation to private actors—with a profit motive and fiduciary responsibilities outside the delegation relationship—will be legitimate *unless* the delegating agents have an ownership stake in the delegate entity. So, while directly and completely delegating SRM responsibility to SpaceX/Tesla or something similar will generate serious concerns, it may be the case that a constructed entity where both the delegator and the delegates participate might succeed. This, however, would only work if one felt that mitigating climate harms via SRM was the sort of government responsibility that could be privatized/outsourced. If we think of SRM as more akin to various mitigation/adaptation responsibilities, then it may be, in principle, outsourceable. For example, a government can contract with a private company to build a seawall. If we conceptualize SRM more as kind defense or assurance against climate injustice, then private outsourcing looks less normatively attractive. For example, paying McKinsey to represent one's interests at COP may be an abdication of political responsibility *even if* it had considerable benefits.
- **Substantive control:** Delegation of SRM responsibilities will likely only be just if the agent with the normative permission has key levels of substantive control over how the delegate behaves. Whether less powerful nations will be able to exercise this sort of control will likely depend on several factors. Does the more powerful delegate use its own legal system—or participation in international, multilateral institutions—to create pre-commitment mechanisms that make ignoring the delegator expense. Does the delegator represent or reflect the will of a multilateral coalition of agents with the relevant normative permissions or is it a single nation? Is the internal structure of the delegate institution designed in such a way as to ensure control by the delegator? Does the delegator have the internal capacities to exercise that control in the first place? These are empirical questions that cannot be answered in the abstract.