Forthcoming in *Environmental Values* ©The White Horse Press http://www.whpress.co.uk

Aerosol Geoengineering Deployment and Fairness

Toby Svoboda

Abstract

If deployed, aerosol geoengineering (AG) could involve unfairness to both present and future parties. I discuss three broad risks of unfairness that an AG deployment policy might carry: (1) causing disproportionate harm to those least responsible for climate change, (2) burdening future parties with the costs and risks of AG, and (3) excluding some interested parties from contributing to AG decision-making. Yet despite these risks, it may be too hasty to reject AG deployment as a potential climate change policy. I argue that since it is very unlikely that a completely fair climate change policy will be pursued, we have ethical reason to prefer some "incompletely fair" policy. Given various facts about our world, it might be the case that some AG policy is ethically preferable to many other feasible climate change policies, even if AG carries deeply problematic risks of unfairness.

Introduction

Deployment of aerosol geoengineering (AG)—or injecting reflective aerosols into the Earth's stratosphere in order to induce global cooling, thus compensating for global warming caused by increased anthropogenic greenhouse gases—has the potential to avert some of the harmful impacts of anthropogenic climate change. However, AG also faces difficult ethical questions, such as whether its deployment would violate various norms of justice (Svoboda et al., 2011). In this paper, I trace some concerns regarding AG and fairness. Specifically, there are certain risks that, if deployed, AG could involve substantial unfairness to both present and future parties. I identify and discuss the following such risks:

Some Risks of Unfairness in AG Deployment

1. Disproportionate harm to parties least responsible for climate change.
2. Burdening future generations with costs and risks of AG.
3. Exclusionary decision-making on AG deployment.

Risks of these types are already recognized as ethically problematic in the literature on geoengineering ethics (Preston, 2012), but further analysis is needed to determine the conditions (if any) under which AG might be ethically preferable to other options despite such risks. After briefly discussing each of them, I examine possible measure to reduce risks of such unfairness in AG deployment. Most of the literature on the ethics of AG and other "solar radiation management" techniques focuses on their ethically problematic aspects. While it is important to identify such issues, it is also worthwhile to consider whether these problems can be ameliorated, something that is rarely done in the current literature. I argue that we should search for ways of ameliorating ethically problematic features of AG, given both that AG might be deployed in the future and that some AG policies could involve less risk of unfairness compared to others. All else being equal, an AG policy that involves reduced risks of unfairness would be ethically preferable to an AG policy that involves greater risks of unfairness, so it is presumably worthwhile to identify less risky AG policies. Indeed, a driving rationale for AG research is that deployment thereof could prove attractive should efforts to mitigate emissions fail to avoid dangerous climate change (Crutzen, 2006). The idea here is that under certain conditions some policy involving AG may be better (or less bad) than other available options, all things considered. This rationale can (but need not) rely on appeals to emergency, which contend that AG may be able to avoid some "climate emergencies," such as threshold collapses in the climate

system. Most proponents of this rationale for researching AG view this as a possible scenario in the future, but some have already declared emergency (AMEG, 2012). Whatever we think of arguments of this variety, ethicists who only criticize AG for its problematic features risk overlooking the comparative point that AG may be ethically preferable to other options in certain (most likely future) scenarios, whether or not those scenarios constitute genuine emergencies. As I will argue, this comparative point deserves consideration. Unless we see a serious commitment to mitigation soon, we can expect a future in which our climate policy options are all ethically problematic—assuming we have not reached that point already. In such a case, it would not be enough only to point out the ethical deficiencies of all these options, for we would also need some ethical guidance in deciding which ethically problematic options to consider seriously.

On the whole, this paper aims to make three new contributions: first, to identify promising ways of reducing recognized risks of unfairness that AG could carry; second, to introduce a conception of incomplete fairness as a useful tool in thinking about the ethics of potential responses to climate change, particularly in circumstances in which completely fair options are not available; and third, to argue that, perhaps surprisingly, there is a strong *prima facie* case to be made that AG could be quite attractive from an ethical point of view under such circumstances, which are likely to hold in the future if aggressive, large-scale emissions mitigation does not occur soon.

I should mention a few caveats. First, I focus exclusively on issues of fairness regarding *deployment* of AG, setting aside the many important ethical issues regarding research and testing of AG (Morrow et al., 2009; Tuana et al., 2012). Second, fairness is presumably not the only ethically relevant consideration pertaining to AG, but rather just one important consideration that deserves close study. Accordingly, my exclusive focus on fairness in this paper should not be

taken to imply that these other considerations (e.g., regarding how AG would impact overall welfare) are unimportant. Third, I do not address the question of whether AG deployment would be ethically permissible, nor under what conditions it would be permissible (if any)—rather, I examine some ethically worrisome features of AG and suggest ways of addressing them. Whether AG ought to be deployed, perhaps despite substantial risks of unfairness, is a question distinct from those I am able to address here.

Still, one might ask whether it is worth examining how to reduce these risks of unfairness if we are not even sure that AG deployment would be permissible. In response, I would note that since AG is receiving serious attention from respected scientific bodies (IPCC, 2014; Shepherd et al., 2009) and that some scientists are calling for near-term deployment of AG (AMEG, 2012), we should look for ways to ameliorate ethically problematic features of AG now, since certain agents might deploy it in the future even if doing so is not ethically permissible. In such a case, it would be useful to have identified AG policies that carry less risk of unfairness than others, since less risky policies would be ethically preferable to more risky policies. This is why the idea of incomplete fairness is important, for it provides a way to think through the respective merits and deficiencies of different response to climate change in cases in which all options are ethically problematic. This is not a pointless exercise, since leading proponents of researching AG recognize that it raises ethical issues (Keith, 2013), and so it is reasonable to hope they would share a desire to reduce risks of unfairness in AG. Indeed, many arguments for researching AG imply that deploying it in the future may be ethically preferable to other available options (e.g., because it involves the least harm), and this suggests proponents of such arguments may be sympathetic to attempts at reducing risks of unfairness.

A Theory of Fairness

While it is perhaps intuitively obvious that the risks I discuss below involve unfairness, we need some theory of fairness in order both to explain why these risked outcomes would be unfair and to guide the crafting of a response to climate change that reduces risks of unfairness. For the purposes of this paper, I rely on John Broome's influential theory of fairness (Broome, 1990), which both fits with and explains our intuitions about clear cases of fairness or unfairness, doing so in an elegant and plausible fashion. Utilizing his theory is helpful because, while it fits with common-sense views about the risks I discuss, it also provides a plausible account of why the risked outcomes would be unfair if realized. In this section, I first sketch Broome's theory. I then distinguish between what I call complete and incomplete fairness, which respectively correspond to full and partial satisfaction of legitimate claims to some good according to strength of those claims.

On Broome's account, fairness is the proportional satisfaction of claims that various parties have upon some good. A claim is a particular kind of reason why some person should receive some good, namely a *prima facie* duty that is owed to that person herself. Such a duty entails that this person is owed some share of the good in question (Broome, 1990: 92). If multiple parties have claims upon some good, then fairness requires that good to be divided in proportion to the respective strength of these parties' claims. If all parties have an equally strong claim on some good, then fairness requires that good to be divided among them equally. If some party has a stronger claim on some good than another party, then fairness requires that the former receive a greater portion of that good than the latter, where the exact portion is determined by how much stronger the claim of the former is.

I will not attempt to provide necessary and sufficient conditions for what counts as a (legitimate) claim, but it is plausible to take desert to constitute one type (Broome, 1990: 93). Let us put aside potential non-desert claims for the moment. If you deserve some divisible good that I do not deserve, then you have a claim on that good while I lack a claim on it. If no one else deserves that good—and recalling that we are ignoring potential non-desert claims for the moment—then fairness requires that good to be allotted to you in full, since you are the only party possessing a claim on it. Alternatively, if some others deserve that good as well, then fairness requires that good to be allotted to deserving parties according to how much they deserve it. Given equal desert among all parties, an equal allotment of the good would be fair. Given unequal desert among parties, the good should be allotted in proportion to that desert. For example, if you deserve the good twice as much as I deserve it, and if no other claims are relevant, then fairness requires that your share of the good be twice as great as my share.

Needs might constitute another type of claim. For example, if some person needs an organ transplant in order to continue living, we might think that this person has a claim on available organs of the required kind. If so, then this person is owed an organ. Yet if there are fewer organs in the donation bank than there are persons in need of transplants, fairness requires that the bank of organs be allotted while taking into account the extent to which each party needs a transplant. For example, a fair distribution of limited organs might involve giving transplants to those whose need is most urgent (e.g., because otherwise they will soon die) while deferring transplants for those whose need is least urgent, assuming that there are not other types of claim that are relevant here.[1]

[1] Welfare might constitute a third type of claim (Broome, 1990). If so, then some party would have a claim upon any good that would contribute to that party's welfare, and this would entail a *prima facie* duty to provide this party with a share of the good in proportion to the strength of her welfare-based claim.

Different types of claim may be simultaneously relevant in a single case. Suppose that both desert and need constitute claims, and suppose that you have a desert-based claim on some good (but lack any other claim on it) while I have a need-based claim on that good (but lack any other claim on it). What does fairness require in this case? Since Broome holds that fairness requires satisfying claims in proportion to how strong they are, fairness involves granting all claims some degree of satisfaction, with stronger claims receiving greater satisfaction than weaker claims (Broome, 1990: 95). Applying this to the case just mentioned, fairness would require that my need-based claim and your desert-based claim be satisfied in proportion to their respective strength. If my need-based claim is much weaker than your desert-based claim, then it would be fair for me to receive a much smaller share of the good in question while you receive a much greater share of it. If these two claims are equal, then it would be fair for each of us to receive half of the good in question. And in cases in which our claims are of equal strength but the good at issue is indivisible (e.g., the last remaining organ in the donation bank), Broome suggests it would be fair to institute a lottery, for this would give each of us an equal chance of receiving the good (Broome, 1990: 99). While this is only a rough sketch of a theory of fairness, it will suffice for the purposes of this paper.

Let us say that a state of affairs is completely fair when all claims to some good or set of goods are satisfied in exact proportion to the strength of those claims.[2] If complete fairness holds with respect to some good, then each claimant receives a share of the good that is proportional to the strength of his or her claims relative to the claims of others. There are many cases in the actual world where complete fairness fails to hold. One cause of this is that individuals, institutions, and governments fail to comply with their duties to parties who have claims upon

[2] Arguably, complete fairness also requires that no person's rights (understood as side constraints) be violated, which is perhaps distinct from whether all claims receive proportional satisfaction (Hooker 2005). I shall not address this further, however, since my focus is on incomplete fairness in terms of the proportional satisfaction of claims.

various goods. Recall that some party's having a claim upon some good entails a *prima facie* duty to satisfy that claim in proportion to its strength. An individual, institution, or government that ignores or prevents such satisfaction therefore violates a duty to the claimant, who is owed some degree of satisfaction regarding her claim.[3] For example, if some community has a claim upon some natural resource (e.g., because it needs or deserves it), and this claim is ignored by some institution that harvests that resource and profits from it, then the institution has violated a *prima facie* duty to that community. This is so even if that institution also has some claim upon the natural resource, for it still would have failed to give proportional satisfaction to the claim of the community.

Let us say that a state of affairs is incompletely fair when some claims to some good or set of goods are satisfied, but only to an imperfect degree. This captures the thought that some scenarios involving unfairness are clearly more fair (or less unfair) than others. Although a simple idea, the notion of incomplete fairness lets us make sense of this thought without ignoring the ethical deficiencies of scenarios that do not fully satisfy legitimate claims in proportion to their strength. To judge something as incompletely fair is to recognize that it exhibits both virtues and vices: on the one hand it satisfies (at least partially) some legitimate claims in proportion to their strength, but on the other hand it fails to satisfy other such claims (or fails to satisfy them in full). Incomplete fairness can encompasses a wide range. On the one hand, suppose that all claims to some good have received proportional satisfaction, save one claim that receives slightly less satisfaction than it should. On the other hand, suppose that half of all claims to some good have received proportional satisfaction while the other half of claims have received

[3] To be more precise, preventing satisfaction of some claim violates a *prima facie* duty, but it might be permissible in certain cases. Although considerations of fairness provide moral reasons for action, there may be countervailing moral reasons (e.g., regarding overall welfare) that over-ride them in certain cases. Given this possibility, preventing satisfaction of some claim, while always unfair, sometimes is morally permissible all things considered.

only partial satisfaction. While both scenarios are incompletely fair, the former is clearly more fair (or less unfair) than the latter. While I will not address whether it is ethically permissible to pursue policies expected to yield incomplete fairness, I do assume that, given two incompletely fair scenarios, there is an ethical reason to prefer the one that is more fair (or less unfair), all else being equal. As I shall illustrate below, the idea of incomplete fairness is helpful in thinking through potential AG proposals, for it offers a kind of middle ground that allows for nuanced positions that neither condemn nor endorse particular geoengineering techniques wholesale. This is useful because completely fair climate policies may not be available under some plausible future conditions, and in such cases we need tools that help us ethically evaluate the options available to us.

Risk 1: Disproportionate Harm to Parties Least Responsible for Climate Change

It is widely recognized that deploying AG on a global scale has the potential to create both winners and losers, with some parties (e.g., regions, states, or individuals) being made worse off and others better off with respect to some baseline. AG could lead to changes in precipitation, increased or decreased agricultural productivity, ozone depletion, economic damages, and so on (Kravitz et al., 2013; Tilmes et al., 2008; Goes et al., 2011). Since these impacts would not be uniform, and since different parties might be either harmed or benefited by the same impacts (or harmed or benefited to varying degrees), AG has the potential to result in net harm for some parties and net benefit for others. Importantly, the set of those benefited or harmed by AG—as well as the degree to which they are harmed or benefited—likely would be non-identical to those harmed or benefited under some different policy (e.g., one of adaptation, mitigation, or business-as-usual). One way to conceptualize this is that AG could affect both the

quantity and distribution of what I call "climate-sensitive goods" (see below). Given this potential, we should ask whether the outcome of deploying AG would be fair in terms of how such goods are altered in distribution and quantity. While these outcomes could violate norms of distributive justice in various ways (Svoboda et al., 2011), I will focus specifically on how AG might impact parties who bear little or no responsibility (neither causal nor moral) for anthropogenic climate change.

To take a specific example, computer model simulations suggest that AG could alter regional precipitation patterns, with reductions in average annual precipitation occurring in South America, Africa, and southeast Asia (Irvine et al., 2010). Such reductions in precipitation could impact agricultural productivity and freshwater availability (Haywood et al., 2013), thus harming individuals who reside in affected regions.[4] Incidentally, parties in these regions are among those least responsible (if they are at all) for the problem that AG is meant to address, namely dangerous climate change caused by anthropogenic emissions of greenhouse gases. Intuitively, it seems deeply unfair that those least responsible for this problem should be burdened with significant harms and costs that result from addressing it. The theory of fairness discussed above is well-placed to explain why. Let us use the term "climate-sensitive goods" to indicate goods that are likely to be affected (e.g., by increasing or decreasing their quantity or quality) by climate change, whether that climate change is facilitated by anthropogenic emissions or geoengineering technologies. Parties residing in regions at risk of precipitation change have claims upon various climate-sensitive goods, such as those they need (e.g., food and drinking

[4] Of course, we must be careful about concluding too much from these models, given their inherent uncertainties and the complexity of the systems they seek to represent. Perhaps AG-induced precipitation change would be less significant than some models currently suggest; perhaps AG would result in *increased* agricultural productivity by reducing heat stress on crops (Pongratz et al., 2012); or perhaps regions experiencing precipitation reduction would also receive ample benefits from AG that make them net "winners." Nonetheless, there is a risk here that precipitation change could result in substantial net harm in affected regions.

water). Since deployment of AG threatens to compromise some claimants' access to such climate-sensitive goods, AG might already involve a *prima facie* violation of some duty of fairness, for its deployment may result in a satisfaction of claims that is disproportionate to their strength. If so, then AG would be unfair in this respect.

But the problem is actually worse than that. Arguably, a party that is morally responsible for climate change thereby has a weaker claim on climate-sensitive goods than a party that is not so responsible, all else being equal. Suppose that two parties, A and B, begin with equally strong claims upon some good, G. Fairness would require that G be divided equally between A and B. Now suppose that A causes the value of G to be reduced and that A is morally responsible for doing so (e.g., through negligent management of G). It is reasonable to hold that A now has a weaker claim than B upon G. After all, it is A's fault that G is now less valuable, and we are supposing that B is faultless in the matter. Now fairness would require that B receive a greater share than A of what remains of G. It would be unfair for B to receive only half of G, but it would be even more unfair for B to receive less than half of G.

Anthropogenic greenhouse gas emissions are substantially damaging many climate-sensitive goods upon which various parties have claims. This reduces the strength of those claims on the part of parties morally responsible for that damage, such as high emitters residing primarily in developed countries, and so those not morally responsible for that damage now have a comparatively stronger claim on what remains of the damaged, climate-sensitive goods. If such stronger claims are not satisfied (e.g., because AG deployment further damages relevant, climate-sensitive goods), the result is a kind of compound unfairness: the initial damage is unfair because it fails to provide proportional satisfaction of the claims of some parties, and the subsequent damage is also unfair because it fails to provide proportional satisfaction of the now

stronger claims of those parties. This offers an explanation for why disproportionate harm to those least responsible for climate change seems particularly unfair. Now it would be premature to hold that AG deployment *would* involve such unfairness. Further, there may be cases in which it is morally permissible to bring about unfair outcomes (e.g., because some other important duty over-rides *prima facie* duties of fairness), so even if AG would result in unfair outcomes it does not automatically follow that it ought not to be deployed. Yet we should be mindful of the potential for certain AG policies to entail such unfairness. Imagine some policy that is *likely* to create a situation in which the winners of AG deployment (e.g., those who are spared both the harmful impacts of emissions-driven climate change and the harmful impacts of AG) are those most responsible for climate change, while the losers are those least (or not at all) responsible for climate change. Intuitively, this would be a policy that poses a substantial risk of unfair outcomes, and that is morally problematic even if deployment of that policy turns out to be justified, all things considered.

Risk 2: Burdening Future Generations with Costs and Risks of AG

There is also a suite of challenges specifically regarding AG deployment and fairness to future generations. First, we should ask whether deployment of AG on the part of some present generation would unfairly burden future generations. If, for example, the present generation deploys AG without also reducing its emissions, then future generations would have to maintain AG in order to avoid the harmful climate change that might occur if they allowed AG to lapse (Goes et al., 2011). Thus, a decision to deploy AG might be unfair to future generations by effectively forcing them to bear the costs (broadly construed to include any harmful side-effects of AG) of a policy meant to address a problem they did not create, namely anthropogenic climate

change. This might occur as a result of parochialism (Gardiner, 2011), such as if the present generation deploys AG to avoid paying the costs of mitigation without heeding or caring about the costs this imposes on future generations.

Further, maintaining AG post-deployment faces the so-called "termination problem." Aerosols used to reflect incoming solar radiation have a relatively short atmospheric lifespan (Rasch et al., 2008), meaning that new aerosols would need to be injected regularly in order to maintain a constant cooling effect. If this replenishment should cease for any reason (e.g., terrorism, war, or malfunction), the result would be extremely rapid global warming due to the anthropogenic greenhouse gases still in the atmosphere. There is a risk of AG being terminated prematurely, leading to substantial economic damages and other harms in the future (Baum et al., 2013; Goes et al., 2011). Thus, we might ask whether it is fair to put future generation at risk in this way.

Broome's theory of fairness is able to explain in a plausible fashion why such burdens to future generations could involve unfairness. Future parties presumably have claims upon various goods, including climate-sensitive ones. Some of those goods, such as natural resources, are shared across generations and are thus intergenerational goods. Fairness thus requires that a future party receive a share of an intergenerational good that is proportional to the strength of that party's claim upon that good. If a present party acts to compromise the future party's receipt of its fair share (e.g., by depleting some natural resource), that present party has acted unfairly by making it difficult or impossible for the future party's claim to be satisfied in proportion to its strength. Now AG deployment could compromise future parties' receipt of goods upon which they have claims, and so AG deployment risks unfairness. For example, since AG could decrease agricultural productivity and freshwater availability in some regions, there is a risk of it

compromising future parties' access to those goods, perhaps to such a degree that it would prevent the proportional satisfaction of their claims. Likewise, unintentional termination of AG in the future could limit the access of future parties to these and other climate-sensitive goods, given that the rapid global warming allowed by sudden termination of AG is expected to have far-reaching and substantial impacts (Goes et al., 2011).[5]

Risk 3: Exclusionary Decision-Making on AG Deployment

Finally, there is a serious concern regarding whether AG would be deployed in a procedurally fair fashion. If deployed on a global scale, AG would entail substantial impacts (potentially both harmful and beneficial) and risks for various parties around the world. We may identify these as interested parties, since they would have a stake in whether or not AG is deployed. Plausibly, fairness requires that each party to be substantially affected or put at non-negligible risk have an opportunity to contribute to decision-making on whether or not AG is deployed. Ideally, AG would be deployed only with the consent of such affected or at-risk parties, including members of Indigenous communities whose voices historically have not been heard on environmental matters in which they have substantial stakes (Powys Whyte, 2012). It might be the case that complete procedural fairness requires that such parties consent to AG deployment. However, since the preferences of various parties regarding climate policy are likely to diverge, it is perhaps unlikely that all would consent to the choice of a single response to

[5] Considerations of intergenerational fairness are complicated by the non-identity problem (Parfit, 1984). Choices on the part of present persons can alter which persons are born in the future, and this creates puzzles regarding our obligations to future persons, such as whether it is possible to harm them. While I cannot resolve the non-identity problem here, of course, there is a plausible response that helps secure the intuition that AG deployment carries risks of unfairness to future generations. We might reject the so-called "person-affecting principle," according to which an action is wrong only if it makes someone worse off (Parfit, 1984: 363, 378). Then, adopting an impersonal view for (at least some) duties of fairness, we can hold that we have a *prima facie* obligation to maintain fair distributions of goods regardless of the identities of the individuals who have claims upon those goods.

anthropogenic climate change, whether or not it would involve geoengineering. One thing a theory of procedural fairness should specify is how to preserve fairness (or as much of it as possible) in decision-making when universal consent is not feasible. It seems clear that it would be unfair to exclude an interested party from decision-making procedures regarding AG deployment. This does not entail that merely allowing such a party the opportunity to vote on potential policies would be sufficient to make AG deployment procedurally fair, of course. Genuine procedural fairness may require that interested parties have the opportunity to take part in deliberations at various stages of decision-making, including decisions about what policy options are to be on the table and what research trajectories are to be pursued (Powys Whyte, 2012). We may identify exclusion of at-risk parties from such deliberations as a form of procedural unfairness.

While perhaps not obvious at first, the theory of fairness I am relying on can explain why excluding interested parties from AG decision-making could be unfair. First, we can view the power to make and implement a decision as a divisible good, since some interested party can have more or less influence on some decision than another party. For example, a militarily dominant state might be able to impose economic sanctions on some other state unilaterally, taking into account neither the preferences of the target state nor its trading partners. Second, we can hold that various interested parties can have claims upon this good, namely the power to make and implement decisions. This suggests that fairness requires such decision-making power to be distributed according to the strength of those claims. If all claimants have an equal claim to make a decision on some issue, then fairness requires the decision making power to be divided equally among all claimants. If some claimant has a much stronger claim than another, then it would be fair for the former to have a much greater impact on the decision than the latter. And so

on. This is not the place to develop and defend an account of what constitutes a legitimate claim to make a decision, of course, but it is plausible to understand procedural fairness in terms of the theory of fairness used in this paper.

Unfortunately, AG seems particularly prone to procedural unfairness. Given the relatively low cost estimates for deployment (Barrett, 2008), and given the technical feasibility of deploying without broad participation (e.g., unlike cutting emissions, AG does not face a daunting collective action problem), there is potential for a single state to deploy AG unilaterally (Victor, 2008). Obviously, a genuinely unilateral decision by some party to deploy AG would preclude all other interested parties from contributing to the decision of whether or not AG is to be deployed. Since such interested parties presumably have claims to contribute to a decision on AG, so precluding them would fail to give proportional satisfaction to their claims and thus would be unfair. Since unilateral deployment is possible, there is a risk of procedural unfairness with respect to AG deployment.

While Horton (2011) suggests that the potential for unilateral AG deployment is greatly exaggerated, even multilateral decisions to deploy AG could still exclude many claimants and thus could be procedurally unfair. For example, Ricke et al. (2013) argue that coalitions who wish to deploy AG would have incentives to remain small, excluding broad participation in decision-making. Since different parties may have competing interests best served by different geoengineered climates, it might be difficult for large coalitions to agree on a particular AG policy (e.g., some parties might prefer less aggressive cooling than others). Thus, it may be in the interest of a small coalition of parties to deploy an AG policy that best serves their own preferences. Of course, this would exclude other parties from participating in the decision-

making process. If such interested parties have claims on the relevant decision-making power, excluding them would involve procedural unfairness.

How Might these Risks of Unfairness be Reduced?

I have identified three broad risks of unfairness associated with AG deployment. To be clear, I have not claimed that deploying AG *would* involve unfairness in any of these respects, especially since some AG policies could pose substantially different degrees of risk. I now will suggest some measures for reducing the risks of unfairness I have noted. This is a worthwhile pursuit for ethicists, given the recent surge in research on AG, interest from the United States Congress and from wealthy individuals (Bullis, 2009; Kintisch, 2010), the performance of a small-scale field test (Izrael et al., 2009), and even some calls for near-term deployment (AMEG, 2012). We need not assume that deployment of AG would be permissible, of course. Indeed, there may be a presumption against the ethical permissibility of geoengineering (Preston, 2011). Further, perhaps certain parties ought to have mitigated their emissions some time ago, and perhaps that would have been sufficient to avert harmful or unjust climate change, making geoengineering unnecessary. Yet it is precisely because we have failed on the mitigation front that some are arguing for research on AG and other geoengineering technologies, as well as serious consideration of their eventual deployment. Now if AG is deployed in the future, we would have good ethical reason to prefer some version that carries less risk of unfairness than others, all else being equal—at the least, an AG policy with reduced risk of unfairness would seem to be less ethically problematic than an AG policy with greater risk of unfairness, again assuming all else is equal. Since AG deployment in the future is a serious possibility, we should identify ways of reducing ethical problems associated with it, including risks of unfairness.

First, we might ask if AG deployment could be technically designed such that it substantially reduces the risks discussed above. First, computer model simulations suggest that AG's impact on precipitation patterns could depend on both the location and quantity of aerosol injected into the stratosphere. Haywood et al. (2013) find that "large asymmetric stratospheric aerosol loadings concentrated in the Northern Hemisphere are a harbinger of Sahelian drought whereas those concentrated in the Southern Hemisphere induce a greening of the Sahel." This indicates that, compared to northern hemisphere deployment of AG, deployment in the southern hemisphere could involve a reduced risk to parties living in the Sahel. Further, Irvine et al. (2010) find that AG's impact on the hydrological cycle could be tied to the quantity of aerosols injected, with lower quantities leading to reduced impacts on precipitation. If that is the case, then less aggressive cooling meant only to slow (rather than reverse) global warming (Keith, 2013) may be more fair (or less unfair) than more aggressive cooling.

This suggests that, depending on their technical design, some AG policies could carry a reduced risk of failing to satisfy some parties' claims upon climate-sensitive goods. For example, the risk of southern hemisphere AG resulting in drought for low emitters in the Sahel may be less than the risk of northern hemisphere AG doing so. If that is the case, then we would have a *prima facie* ethical reason to prefer the former AG policy over the latter one, since at least in this one respect the former would be less likely to violate the legitimate claims of parties living in the Sahel. From the perspective of incomplete fairness, we would have an ethical reason to prefer southern hemisphere deployment, for it may go further than northern hemisphere deployment in partially satisfying legitimate claims to those goods that are sensitive to alterations in precipitation patterns. Of course, further research would be needed to determine what other risks of unfairness southern hemisphere deployment might carry, as well as how such risks compare to

those of other AG policies. The general point, however, is that technical design decisions might go some distance in ameliorating ethically problematic features of AG. Given a set of possible AG policies, we might investigate their respective risks of unfairness, attempting to identify that AG policy with the lowest such risk, as well as identifying further areas of research to this end.

Second, Wigley (2006) has proposed that AG might be coupled with emissions mitigation, with the former serving as a short-term response meant to "buy time" for instituting sufficient cuts in emissions, after which AG could be phased out. Such a policy potentially could reduce risks of unfairness to future generations, since it would not require AG to be maintained into the distant future. Compared to a stand-alone AG policy, this hybrid policy would decrease the number of future persons whose claims to climate-sensitive goods are not satisfied in proportion to their strength, and it would decrease the number of future persons at risk due to the termination problem. While this still could put near-term future parties at risk of not receiving proportional satisfaction of their claims, and while realization of that risk would be unfair to affected future parties, it presumably would be ethically preferable to a stand-alone AG policy that is unfair to a far greater number of future parties.[6] Relying on my account of incomplete fairness, this would be because a coupled AG-mitigation policy may on the whole satisfy legitimate claims to some set of climate-sensitive goods to a greater degree than a policy employing AG alone. By reducing the duration of AG, it ameliorates the risk of unfairness to future generations due to the termination problem (relative to long-term deployment), and thus it is less likely that future persons will have their legitimate claims to climate-sensitive goods thwarted by termination.

[6] There would still be a question of whether AG deployment would involve a moral hazard, reducing incentive to cut emissions by shielding some parties from some of the harmful impacts of climate change. If so, then even if AG was deployed with the intention to mitigate emissions as well, we may worry that motivation to mitigate would be weakened after deploying AG. But it is not clear that such moral hazard arguments are compelling (Morrow, 2014).

In addition to risk-reducing technical design decisions, we also should look for ways to reduce the risk of unilateral or otherwise exclusionary decisions regarding AG deployment (and perhaps AG research as well). One option is for the international community to adopt and enforce a treaty prohibiting exclusionary decisions of this kind, such as by instituting sanctions on parties that violate this prohibition. Of course, sufficiently powerful parties may be able to ignore such a prohibition if it suits their interests. Alternatively, Horton (2011) suggests that tactics drawn from "international management theory" might encourage multilateralism on AG. He argues that persuasion could be used to convince a skeptical state that some AG policy is actually in its national self-interest. If such persuasion is conducted in good faith and operates in multiple directions (e.g., with developing countries persuading developed ones and not just the reverse), the risk of procedural unfairness may be reduced. Treating the power to make decisions on AG as a divisible good, such measures could be used to increase the share of that good on behalf of parties who might otherwise be excluded from decision-making altogether or who enjoy less influence on decision-making than fairness requires. Partially satisfying some party's legitimate claim to contribute to decisions on AG would be incompletely fair and thus ethically preferable to that claim receiving no degree of satisfaction, all else being equal.

While certain technical design decisions and political tactics could reduce risks of unfairness, we may safely assume that these measures would be insufficient to render AG deployment perfectly fair. Accordingly, we should ask how AG-related unfairness should be addressed and (if possible) rectified. Perhaps such unfairness could be ameliorated by providing economic compensation to affected parties (Bunzl, 2011). If some party does not receive satisfaction that is fully proportional to the strength of a relevant claim to some climate-sensitive good, compensation might "make up" the difference, providing payment that is equal to that

portion of the good that was not (but should have been) received. However, designing a compensation system that is itself fair faces some difficult challenges, including determining who ought to pay compensation, how much ought to be paid, and who ought to receive compensation (Svoboda and Irvine, 2014). There are also limits to what kinds of harm economic compensation could redress, since certain losses (e.g., death or the disappearance of one's culture) seem irreparable in economic terms. While compensation could in principle ameliorate some unfairness resulting from AG deployment, it would be a controversial and difficult matter to determine what would constitute fair compensation. Nonetheless, all else being equal, an AG policy that includes compensation for victims of unfairness would be ethically preferable to an AG policy that provided no such compensation. Again appealing to incomplete fairness, it is plausible to expect that even imperfect compensation will go some distance in satisfying the legitimate claims of parties to climate-sensitive goods, and it may do so to an appreciable extent. The idea of incomplete fairness allows us to make sense of the common-sense judgment that we may have ethical reason to adopt a policy involving both AG and compensation even if that policy involves unfairness, for on the whole it may do better than the alternatives in providing the greatest overall (albeit partial) satisfaction of such claims.

Ethical Comparison of Climate Policies

Under ideal conditions, we would not adopt an incompletely fair climate policy: when possible, we have ethical reason to prefer a policy that renders full satisfaction of all legitimate claims to both climate-sensitive goods and to participation in decision-making regarding climate policies. But geoengineering is arguably receiving serious attention precisely because ideal conditions do not hold, or at least that they may not hold in the relatively near future. High

emissions on the part of some human beings have eroded various climate-sensitive goods and, if left unchecked, will commit us to substantial warming that promises further erosion of such goods in the future. Under such non-ideal conditions, AG may do better than emissions mitigation or adaptation alone when it comes to serving both overall welfare and incomplete fairness, because it could be used to curb and to slow the dangerous warming to which past emissions will have committed us (Keith, 2013). Depending on its severity, such warming might damage both the quantity and quality of climate-sensitive goods, affecting both the welfare of various parties and how those goods are distributed. There are limits to what mitigation and adaptation can do to prevent the erosion and unfair distribution of such goods. Due to committed warming from past emissions and inertia in the climate system, we may face a future scenario in which all available policies would result in substantial net harm (Svoboda, 2012). In such a scenario, this harm could outstrip the adaptive capacities of human communities, and mitigation of future emissions would not help prevent committed warming from past emissions. Conversely, AG could have a fast impact on the climate system, potentially curbing global warming with several years (Moreno-Cruz and Keith, 2013). In a bad future scenario of this sort, there is a *prima facie* plausible case that AG could reduce net harm. If so, then we would have an ethical reason to deploy AG in such a scenario.[7]

One common ethical concern with AG is that it carries substantial risks of unfairness, as have seen, but this makes it important to consider specific types of AG deployment that might carry reduced risks of unfairness. It will not do to point out ethical worries about AG in some of its guises and then reject AG as such on that basis, for policies involving AG can differ in many ways. We would be right to worry about a policy that would deploy AG indefinitely or treat it as

[7] More precisely, we would have a *pro tanto* ethical reason to do so. Whether we would have an all-things-considered reason to deploy would depend on many other matters, and I do not take a position on that question here.

a strict alternative to mitigation, for example, both of which would carry substantial risks of unfairness to present and future persons. It is more interesting, however, to think about policies that are sensitive to the risks of unfairness AG might carry. If these risks can be successfully ameliorated, then a major ethical concern regarding AG will lose some of its force. Presumably, no AG policy would ever be fully unproblematic in an ethical sense, but neither would mitigation or adaptation if past warming had committed us to dangerous, net harmful warming. In that case, assuming low risks of unfairness, AG's ability to cool the planet rapidly may tip the balance of ethical considerations in its favor. Importantly, this line of reasoning does not require the (controversial) appeal to the idea of a climate emergency. Rather, it may be that future warming promises to erode climate-sensitive goods gradually, and AG could be ethically attractive because it might slow or reduce this gradual erosion. If it could do so without carrying substantial risks of various types of unfairness, we may have good ethical reason to prefer it over other climate policies that do not involve AG. For example, in such a scenario, a policy that prioritizes the legitimate claims of those least responsible for climate change (e.g., by deploying AG in the southern hemisphere), injects a quantity of aerosols intended to achieve moderate reduction (rather than reversal) in warming, includes a commitment to mitigation of present and future emissions, provides compensation for parties unfairly impacted by AG, and allows for broad participation in decision-making may have much to recommend it ethically.

None of the foregoing should be taken as an endorsement of AG deployment. Indeed, there are serious ethical concerns regarding both field-tests and full-scale deployment of AG (Tuana et al., 2012). Nonetheless, given that some scientists are already calling for near-term deployment (AMEG, 2012), it is important to consider whether and how ethically problematic features of AG deployment policies could be ameliorated. Should we find ourselves in a bad

future scenario, we might have good ethical reason to prefer an AG policy that minimizes risks of unfairness as much as possible.

References

AMEG. 2012. "AMEG Declaration of Emergency." London: Arctic Methane Emergency Group, http://www.ameg.me/index.php/about-ameg/13-ameg-declaration-of-emergency (accessed 11 December 2014).

Barrett, S. 2008. "The Incredible Economics of Geoengineering." *Environmental and Resource Economics* **39**(1): 45-54.

Baum, S. D., T. M. Maher Jr., and J. Haqq-Misra. 2013. "Double Catastrophe: Intermittent Stratospheric Geoengineering Induced by Societal Collapse." *Environment Systems & Decisions* **33**(1): 168-180.

Broome, J. 1990. "Fairness." *Proceedings of the Aristotelian Society* **91**: 87-101.

Bullis, K. 2009. "U.S. Congress Considers Geoengineering." *MIT Technology Review*, http://www.technologyreview.com/view/416187/us-congress-considers-geoengineering/ (accessed 8 May 2014).

Bunzl, M. 2011. "Geoengineering Harms and Compensation." *Stanford Journal of Law, Science & Policy* **4**: 70-76.

Crutzen, P. J. 2006. Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma? *Climatic Change* **77**(3): 211-220.

Gardiner, S. M. 2011. "Some Early Ethics of Geoengineering the Climate: A Commentary on the Values of the Royal Society Report." *Environmental Values* **20**(2): 163-188.

Goes, M., N. Tuana and K. Keller. 2011. "The Economics (or Lack Thereof) of Aerosol Geoengineering." *Climatic Change* **109**(3-4): 719-744.

Haywood, J. M., A. Jones, N. Bellouin and D. Stephenson. 2013. "Asymmetric Forcing from Stratospheric Aerosols Impacts Sahelian Rainfall." *Nature Climate Change* 3(7): 660-665.

Hooker, B. 2005. "Fairness." *Ethical Theory and Moral Practice* **8**(4): 329-352.

Horton, J. 2011. "Geoengineering and the Myth of Unilateralism: Pressures and Prospects for International Cooperation." *Stanford Journal of Law, Science & Policy* **4**: 56-69.

IPCC. 2014. *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report to the Intergovernmental Panel on Climate Change* [final draft], http://www.ipcc.ch/report/ar5/wg3/ (accessed 8 May 2014).

Irvine, P. J., A. Ridgwell and D. J. Lunt. 2010. "Assessing the Regional Disparities in Geoengineering Impacts." *Geophysical Research Letters* **37**.

Izrael, Y. A., et al. 2009. "Field Experiment on Studying Solar Radiation passing through Aerosol Layers." *Russian Meteorology and Hydrology* **34**(5): 265-273.

Keith, D. 2013. *A Case for Climate Engineering*. Cambridge, Massachusettes: MIT Press.

Kintisch, E. 2010. "Bill Gates Funding Geoengineering Research." *Science Insider*, http://news.sciencemag.org/scienceinsider/2010/01/bill-gates-fund.html (accessed 8 May 2014).

Kravitz, B., et al. 2013. "Climate Model Response from the Geoengineering Model Intercomparison Project (GeoMIP)." *Journal of Geophysical Research: Atmospheres* **118**: 8320–8332.

Moreno-Cruz, J. B., & Keith, D. W. 2013. Climate Policy Under Uncertainty: A Case for Solar Geoengineering. *Climatic Change* **121**(3): 431-444.

Morrow, D. R. 2014. "Ethical Aspects of the Mitigation Obstruction Argument Against Climate Engineering Research." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **372**(2031).

Morrow, D. R., R. E. Kopp and M. Oppenheimer. 2009. "Toward Ethical Norms and Institutions for Climate Engineering Research." *Environmental Research Letters* **4**(4).

Parfit, D. 1984. *Reasons and Persons*. New York: Oxford University Press.

Pongratz, J., D. B. Lobell, L. Cao and K. Caldeira. 2012. "Crop Yields in a Geoengineered Climate." *Nature Climate Change* **2**(2): 101-105.

Powys Whyte, K. 2012. "Indigenous Peoples, Solar Radiation Management, and Consent." In Preston (ed.), *Engineering the Climate: The Ethics of Solar Radiation Management*. Lanham, Maryland: Lexington Books. 65-76.

Preston, C. J. 2011. "Re-Thinking the Unthinkable: Environmental Ethics and the Presumptive Argument against Geoengineering." *Environmental Values* **20**:4: 457-479.

Preston C. J. (Ed.). 2012. *Engineering the Climate: The Ethics of Solar Radiation Management*. Lanham, Maryland: Lexington Books.

Rasch, P. J., P. J. Crutzen and D. B. Coleman. 2008. "Exploring the Geoengineering of Climate Using Stratospheric Sulfate Aerosols: The Role of Particle Size." *Geophysical Research Letters* **35**.

Ricke, K. L., J. B. Moreno-Cruz, and K. Caldeira. 2013. "Strategic Incentives for Climate Geoengineering Coalitions to Exclude Broad Participation." *Environmental Research Letters* **8**(1).

Shepherd, J. 2009. *Geoengineering the Climate: Science, Governance and Uncertainty*. London: The Royal Society.

Svoboda, T. and P. J. Irvine. 2014. "Ethical and Technical Challenges in Compensating for Harm Due to Solar Radiation Management Geoengineering." *Ethics, Policy & Environment* (in press).

Svoboda, T. 2012. "Is Aerosol Geoengineering Ethically Preferable to Other Climate Change Strategies?" *Ethics & the Environment* **17**:2: 111-135.

Svoboda, T., K. Keller, M. Goes and N. Tuana (2011). "Sulfate Aerosol Geoengineering: The Question of Justice." *Public Affairs Quarterly* **25**(3): 157-180.

Tilmes, S., R. Müller and R. Salawitch. 2008. "The Sensitivity of Polar Ozone Depletion to Proposed Geoengineering Schemes." *Science* **320**(5880): 1201-1204.

Tuana, N., R. Sriver, T. Svoboda, R. Olson, P. Irvine, J. Haqq-Misra and K. Keller. 2012. "Towards Integrated Ethical and Scientific Analysis of Geoengineering: A Research Agenda." *Ethics, Policy & Environment* **15**(2): 136-157.

Victor, D. G. 2008. "On the Regulation of Geoengineering." *Oxford Review of Economic Policy* **24**(2): 322-336.

Wigley, T. M. L. 2006. "A Combined Mitigation/Geoengineering Approach to Climate Stabilization." *Science* **314**(5798): 452-454.