

Is Aerosol Geoengineering Ethically Preferable to Other Climate Change Strategies?

Pre-Print Version

Toby Svoboda (Assistant Professor of Philosophy, Fairfield University)

Forthcoming in *Ethics & the Environment* (2012), Indiana University Press

Introduction.

In this paper, I address the question of whether aerosol geoengineering (AG) ought to be deployed as a response to climate change. First, I distinguish AG from emissions mitigation, adaptation, and other geoengineering strategies. Second, I discuss advantages and disadvantages of AG, including its potential to result in substantial harm to some persons. Third, I critique three arguments against AG deployment, suggesting reasons why these arguments should be rejected. Fourth, I consider an argument that we ought to adopt that response to climate change which would result in the least net harm to persons. I suggest that *under certain conditions*, such as a climate emergency scenario, the least harmful climate change strategy *could* involve deployment of AG. The implication is that, despite the risks of harm associated with it, in certain situations AG could be (or be part of) a climate change strategy that is ethically preferable to other available strategies.

Mitigation, Adaptation, Geoengineering.

Climate change could result in severe harm for some persons¹ through various impacts, including sea-level rise, ocean acidification, droughts, and an increase in the frequency of severe weather events (IPCC 2007). Partly because anthropogenic emissions of greenhouse gases are the driving forces behind these potentially harmful impacts, climate change raises a number of

ethical issue (Singer 2004, 14-50; Gardiner 2004). But what ought to be done to reduce or, if feasible, eliminate the harm that climate change threatens to cause? Often discussed responses include mitigating greenhouse gas emissions (Nordhaus 2001) and adapting to the impacts of climate change (Lobell et al. 2008). Another option is geoengineering, or the intentional, large-scale manipulation of the Earth's environment (Keith 2001). In recent years, a growing number of scientists has called for more research on geoengineering as a potential response to climate change (Crutzen 2006; Shepherd 2009; MacCracken 2006; Keith, Parson, and Morgan 2010).

Geoengineering techniques can be divided into two categories. First, carbon dioxide (CO₂) removal would reduce the quantity of greenhouse gases in the atmosphere, either by trapping emissions at their sources or by drawing CO₂ directly out of the atmosphere, such as by using air scrubbers or growing large phytoplankton blooms in the oceans in order to absorb atmospheric CO₂ (Shepherd 2009, 9-23). Once captured, the CO₂ could be sequestered underground or in the deep ocean. If it works as planned, this would cancel the warming effect the sequestered CO₂ would otherwise have had in the atmosphere. However, the costs of some forms of CO₂ removal, such as capturing it at its source and sequestering it, are currently projected to be quite high (Shepherd 2009, table 5.1). Further, the effectiveness of some forms of CO₂ removal, such as fertilizing the oceans with iron in order to foster the growth of CO₂-absorbing phytoplankton blooms, is itself questionable (Strong et al. 2009).

The second category of geoengineering is solar radiation management, which would alter the quantity of solar radiation that is absorbed by the Earth's surface. In the context of climate change, SRM could be utilized to induce global cooling, thus counter-balancing the warming effect of atmospheric greenhouse gases. If it works, this form of geoengineering could avoid some of the potentially harmful impacts of climate change, such as sea-level rise caused by

melting ice sheets (MacCracken 2009). Possible solar radiation management techniques include installing orbital space mirrors to reduce the quantity of solar radiation entering Earth's atmosphere (Keith 2001) and marine cloud brightening to increase cloud albedo and reflect some solar radiation back into space (Jones, Haywood, and Boucher 2011).

For the remainder of this paper, I will focus on what is perhaps the most often-discussed geoengineering technique in the scientific literature, namely AG, which would inject sulfate precursor (e.g., SO₂) particles into the stratosphere, thus increasing the Earth's albedo and reducing the quantity of solar radiation reaching the Earth's surface. Aerosol geoengineering would simulate the effects of a volcanic eruption (Brewer 2007), reflecting some incoming sunlight back into space and thus providing a kind of sun shade that leads to global cooling. Since the degree of cooling can be altered depending on the quantity of aerosols injected into the stratosphere (Irvine, Ridgwell, and Lunt 2010), AG could be used as a global thermostat, turning down the temperature to counter-balance the global warming caused by increased concentrations of atmospheric greenhouse gases. In this way, AG could avoid some (but not all) of the potentially harmful impacts of climate change.

Advantages and Disadvantages of Aerosol Geoengineering.

Aerosol geoengineering is expected to have certain advantages over other responses to climate change. Authors who have urged further research on AG have noted that it appears to be fast-acting (Keith, Parson, and Morgan 2010), inexpensive (Barrett 2008), and capable of averting or diminishing some of the harmful impacts of climate change (MacCracken 2009). First, sulfate precursor aerosols seem to have a rapid impact on global temperatures, providing a significant cooling effect shortly after being injected into the stratosphere (Robock, Oman, and

Stenchikov 2008). AG would be much faster than emissions mitigation, for example, since reducing or even eliminating emissions would not have an immediate, substantial effect on global temperatures. Given the long atmospheric lifetime of CO₂ of centuries to millennia (Archer and Brovkin 2008), past emissions would continue to contribute to climate change.² Second, AG is expected to be relatively inexpensive, costing much less than emissions mitigation (Barrett 2008; Teller et al. 2003). Further, AG could also be much less expensive than adaptation, depending on the severity of the impacts to which humans would need to adapt (IPCC 2007).³ Third, if it works, AG could be used to avert or reduce certain potentially harmful impacts of climate change, such as sea-level rise that threatens persons living in coastal regions (MacCracken 2009) or reductions in agricultural productivity (Morgan and Ricke 2011, 13).

Yet proponents of further research also recognize that AG is “imperfect” (Keith, Parson, and Morgan 2010, 426). I examine four disadvantages of AG, namely that it would not address ocean acidification, that it could alter regional precipitation patterns, that it is subject to abrupt discontinuation, and that it is prone to unilateral deployment.⁴ Each of these disadvantages makes AG potentially harmful to some persons. Accordingly, the prospect of AG deployment raises ethical concerns (Gardiner 2010; Jamieson 1996; Svoboda et al. 2011). I consider each of these four disadvantages in turn.

First, increased ocean acidification is caused by elevated concentrations of atmospheric CO₂. As more CO₂ is emitted into the atmosphere, more is absorbed by the oceans, thus reducing ocean pH (Doney 2006; Doney et al. 2009). This change in ocean chemistry could interfere with the ability of some marine organisms (e.g., corals) to form shells of calcium carbonate, which in turn could have a deleterious impact of coral reef ecosystems (Hoegh-Guldberg et al. 2007). Unfortunately, AG would not address this problem, because it only compensates for global

warming by initiating some degree of global cooling. AG would not slow or reverse the increasing concentrations of atmospheric CO₂ that are driving ocean acidification.

Second, climate model simulations suggest that deploying AG could alter precipitation patterns, significantly reducing average annual precipitation in some regions (Matthews and Caldeira 2007). For example, AG could disrupt the Asian and African summer monsoons (Robock, Oman, and Stenchikov 2008) and potentially lead to droughts (Trenberth and Dai 2007). Precipitation change could also reduce agricultural productivity and access to freshwater in some regions (Robock, Oman, and Stenchikov 2008). The degree to which AG impacts precipitation might depend on the quantity of aerosols that is injected into the stratosphere, with higher quantities contributing to greater disruptions in precipitation (Irvine, Ridgwell, and Lunt 2010).

Third, once AG is deployed, there is a risk that it could be abruptly discontinued, leading to rapid climate change as aerosols dissipate (Goes, Keller, and Tuana 2011; Ross and Matthews 2009). Since sulfate aerosols have an atmospheric lifetime of only a few years (Rasch, Crutzen, and Coleman 2008), new stratospheric injections must be continually administered in order to maintain a constant concentration of aerosols and hence a stable cooling effect. If this maintenance should fail (e.g., due to war, terrorism, or technological malfunctions), the aerosols already in the stratosphere would disperse within several years, allowing global temperatures to increase rapidly (Goes, Keller, and Tuana 2011).

Fourth, given that AG is relatively inexpensive and does not require international agreement in order to be implemented, a single state could deploy AG unilaterally (Victor 2008; Victor et al. 2009). Unlike effective mitigation, which requires that various states cooperate in cutting their emissions, AG could be deployed by one state over the objections of other states.

However, since different regions could be affected in different ways by AG (Irvine, Ridgwell, and Lunt 2010; Robock, Oman, and Stenchikov 2008), unilateral deployment by one state could be contrary to the interests of other states. This raises the potential for geopolitical conflict (Victor et al. 2009). It is unclear how likely it would be that a single state would decide to deploy AG unilaterally, however. While Horton (2011) argues that states would have incentives to cooperate in deployment, Barrett (2008) notes that states especially vulnerable to climate change could have an interest in rejecting cooperation in favor of unilateralism. In this paper, I only note that unilateral deployment of AG is a potential scenario.

The Ethics of Aerosols Geoengineering.

The four disadvantages of AG discussed above have ethical significance. One reason for this is that an AG strategy could result in substantial harm to persons, which raises questions about whether that harm is ethically permissible. For example, ocean acidification could harm persons who depend on coral reef ecosystems for coastal protection or economic income from tourism (Hoegh-Guldberg et al. 2007), precipitation change could harm persons who experience drought or famine in affected regions (Robock, Oman, and Stenchikov 2008), discontinuous AG could harm future generations through rapid climate change (Goes, Keller, and Tuana 2011; Svoboda et al. 2011), and unilateral AG could harm persons in states that do not consent to AG deployment (Victor et al. 2009). Accordingly, AG requires ethical evaluation.

Svoboda et al. (2011) examine some of the ethical issues of AG by relying on a justice framework, arguing that AG faces serious challenges to satisfying the requirements of most if not all theories of justice (Rawls 1999; Dworkin 1981, 1981; Sen 1982; Arneson 1989; Olsaretti 2007). Specifically, they contend (1) that AG-induced precipitation change could harm some

persons much more than others, leading to a distributively unjust outcome;⁵ (2) that deploying AG could burden future generations with the risk of discontinuous AG, thus potentially violating intergenerational justice; and (3) that unilateral deployment of AG would exclude various persons affected by AG from contributing to the decision process concerning whether it should be deployed, thus potentially violating procedural justice.

While this analysis raises ethical concerns regarding the potential deployment of AG, it does not necessarily imply that AG ought not to be deployed. The mere fact that some climate change strategy threatens to harm some persons does not entail that this strategy ought not to be adopted. In order to determine whether a potentially harmful strategy ought to be adopted, further arguments should be considered. I will now examine four such arguments. The first three conclude that AG ought not to be deployed, whereas the fourth concludes that we ought to adopt the climate change strategy that would result in the least net harm to persons.

Risks of Harm Due to Aerosol Geoengineering (Argument 1).

If, contrary to my claim, it was the case that AG ought not to be deployed simply because it threatens to harm persons, something like the following argument (call it A1) would need to be sound:

P1: If we deploy AG, it could result in harm to some persons.

P2: We ought not to do anything that could result in harm to some persons.

C1: Therefore we ought not to deploy AG.

This argument is unsound because P2 is false, and P2 is false because it has very implausible implications. Since virtually all actions carry some risk of harm (i.e., they *could* result in harm to persons), P2 proscribes virtually all actions. If one doubts that virtually all actions carry some risk of harm to persons (either oneself or others), it is instructive to try searching for examples of actions that are risk-free. Someone who walks across the street risks being hit by car; by driving a car one risks hitting someone else; leaving one's home carries a risk of developing skin cancer due to exposure to ultraviolet radiation, even if one takes precautions; a doctor performing surgery risks injuring or killing her patient; and so on. Most of us often find these risks negligible and worth taking, but this suggests that we really do not accept P2. If one was to adopt P2 as a principle, it frequently would counsel inaction regarding even very low-risk activities.⁶ I assume that this is an unacceptable implication. Accordingly, we should reject P2 and hence A1 as well.

Harm Resulting from Aerosol Geoengineering (Argument 2).

One might revise A1 in order to avoid this objection. Consider the following argument (call it A2):

P3: If we deployed AG, it would result in harm to some persons.

P4: We ought not to do anything that would result in harm to some persons.

C2: Therefore we ought not to deploy AG.

Unlike A1, this argument does not claim that AG merely *threatens* harm to persons (i.e., that it *could* harm them) but rather that, if deployed, it would in fact result in harm to persons. Hence, A2 is not prone to the objection that it implausibly proscribes any action or policy that merely

carries some *risk* of harm. Moreover, P3 seems true—it is virtually certain that *some* person or persons would be harmed by an AG strategy, whether due to ocean acidification, precipitation change, abrupt discontinuation, unilateralism, or some other impact. Hence, whether A2 is sound hinges on the truth or falsity of P4.

I contend that P4 is false because accepting it has very implausible implications. There are situations in which, as a matter of common sense, one is permitted to cause harm to a person, whether oneself or another. Examples might include self-defense against a malicious assailant or administering chemotherapy to a cancer patient. Each of these examples involves causing harm to some person, but it would be implausible to hold that a victim of attack is not permitted to defend herself against her assailant or that a doctor is not permitted to administer chemotherapy (which certainly causes the patient to be harmed, even if it leads to eventual benefits) to her cancer patient. If one holds that causing harm is justified in cases like these, then by consistency one should not accept P4, since the principle contained in that premise proscribes doing *anything* that would result in harm to persons.

Further, it is possible that there are scenarios in which all available courses of action would result in harm to persons. In such cases, complying with P4 would require one to reject all available courses of action, even if one of those courses of action was much less harmful than inaction. Yet this seems irresponsible. In a situation in which one has the option of either performing some slightly harmful action or allowing substantial harm to accrue by doing nothing, it seems that one ought to perform the slightly harmful action. Consider Peter Singer's famous example of a person who encounters a drowning child in a shallow pond (Singer 1972, 231). If one does nothing, the child will die, which would be substantially harmful to the child as well as his family and friends. Yet if one dives into the water to save the child, the rescuer's

shoes will be soiled, which would be slightly harmful to the rescuer. No one seriously accepts the view that one should not attempt to rescue the drowning child because it would ruin one's shoes. Yet this is the course of action that would be prescribed by P4 in this situation. If one ought not to do anything that would result in harm to some persons, and attempting to rescue the child would harm a person (namely oneself), then one ought not to attempt to rescue the child. But this conclusion is untenable, which suggests that P4 should not be accepted.

Arguably, climate change provides a scenario in which all available courses of action would result in harm for some persons. The available strategies include emissions mitigation, adaptation, geoengineering, or some hybrid of these (Wigley 2006). Emissions mitigation cannot prevent the harmful impacts resulting from the climate change to which we are already committed due to past emissions (Wetherald 2001), and realistic mitigation targets might not be sufficient to avoid additional harmful impacts from future emissions (Solomon et al. 2009).⁷ Adaptation measures could be insufficient to avoid all harm due to impacts from climate change, such as harm to persons living in small island developing states, whose territories could literally disappear due to sea-level rise (Byravan and Rajan 2010). As far as I am aware, all geoengineering techniques would probably result in harm for some persons, not only AG (Robock 2008). Finally, it is doubtful that a hybrid strategy can be found that would not harm some persons. For example, combining mitigation with AG (Wigley 2006), although it might reduce the harm that would result from either strategy on its own, probably would still harm some persons, e.g. through precipitation change. Thus, it is arguably the case that all available responses to climate change would result in harm to persons, although some could be substantially less harmful than others.

If it is the case that all available climate change strategies would result in harm to persons, and if one accepts P4, then by consistency one should hold that none of these strategies should be adopted and that we ought instead to do nothing about climate change. This would be to accept something very close to a “business as usual” approach, which would allow climate change to occur unabated. But of course a business as usual approach is expected to result in substantial harm to persons through sea-level rise, ocean acidification, droughts, severe weather events, etc. (IPCC 2007). Presumably, such unabated climate change would be more harmful than mitigation, adaptation, and (arguably—see below) some geoengineering strategies. Yet accepting P4 seems to entail acquiescing to this more harmful approach. If this seems implausible, then we should reject P4 and hence A2 as well.

Aerosol Geoengineering as More Harmful than Other Strategies (Argument 3).

Taking account of this difficulty with A2, one might put forward the following argument instead (call it A3):

P5: If we deployed AG as a climate change strategy, it would result in more harm to persons than some other available climate change strategy.

P6: We ought not to adopt any climate change strategy that would result in more harm to persons than some other available strategy.

C3: Therefore we ought not to deploy AG as a climate change strategy.

This argument avoids the problems associated with A2. First, unlike P4, P6 does not proscribe all harmful actions but instead holds that a climate change strategy should not be adopted if some

other available strategy is *less* harmful. This is more plausible than P4, because P6 implies that it could be permissible to adopt a climate change strategy that causes harm to persons, provided that such a harmful strategy is *less* harmful than all other available strategies. Accordingly, if all available climate change strategies would result in harm to persons, those who accept P6 need not reject all those strategies in favor of inaction.

Yet it is not clear whether P5 is true. As discussed above, mitigation strategies also could result in substantial harm to persons, given that mitigation cannot reduce committed climate change and that realistic emissions cuts might be insufficient to avoid many of the harmful impacts of climate change. Likewise, adaptation strategies could result in harm to persons as well, given that the more severe impacts of climate change might outstrip realistic adaptation measures. Whether the harm resulting from mitigation or adaptation strategies would be greater or less than the harm resulting from AG (or some other geoengineering strategy) is uncertain. Hence, it is unclear whether P5 is true.

Further, P5 might be true at a certain point in time yet false at another. It is a contingent matter of fact whether or not AG would be more harmful than other climate change strategies. For example, it is possible that an aggressive mitigation strategy implemented within the next five years would result in substantially less harm than an AG strategy deployed within the next five years. If so, then P5 would be true within the next five years. But imagine that an aggressive mitigation strategy is not adopted in the near future and that rates of greenhouse gas emissions continue to accelerate. In this not implausible scenario (Solomon et al. 2009), it might *become* the case that in the future a mitigation strategy would be more harmful than an AG strategy. Given that AG can quickly reduce global temperatures (Keith, Parson, and Morgan 2010), and given that mitigation cannot reduce committed climate change due to past emissions, P5 could be

false in the future even if it is true now. Accordingly, in evaluating A3, one would need to know when and under what conditions AG is being deployed in order to decide whether it is more or less harmful than other available climate change strategies.

Finally, it is controversial whether the principle contained in P6 should be accepted, because it implies that we ought to adopt a climate change strategy *only if* it would result in less harm than any other available strategy. A rival principle is that we ought to adopt the climate change strategy that is *on balance* most beneficial, i.e. the strategy that results in the greatest net benefit. Yet adopting the strategy that results in the greatest net benefit could entail rejecting the strategy that results in the least harm to persons, since that latter strategy might not deliver the greatest net benefit. This is because the resultant harm or benefit of some strategy is distinct from the resultant *net* harm or benefit of some strategy. Imagine a scenario in which climate change strategy A results in substantial harm but even greater benefit, such that it delivers a net benefit, whereas strategy B results in significantly less harm but produces no net benefit.⁸ In such a case, both the resultant harm and the net benefit of A is greater than the resultant harm and the net benefit of B. If one accepts P6, then by consistency one should hold that strategy B ought to be adopted rather than strategy A. Yet some will find this implausible.

If one is sympathetic to classic utilitarianism (Mill 2008), then it might seem that one ought to adopt strategy A, since on balance it is more beneficial than strategy B (e.g., because it maximizes happiness). Of course, one need not be a classic utilitarian or even a consequentialist in order to think that one ought to adopt strategies that are on balance more beneficial than their competitors. For example, Kantians could hold that A ought to be adopted because it is required by respect for persons. Conversely, negative utilitarians, who accept the ethical principle that one ought to minimize harms rather than maximize benefits (Popper 2002), might prefer strategy B

to strategy A. However, negative utilitarianism has certain implausible implications (Smart 1958), such as that a world with neither harms nor benefits is preferable to a world with minor harms and great benefits. It is not my purpose here to determine whether the principle contained in P6 is true but only to note that it is controversial, thus making A3 controversial as well.

Adopting the Strategy Resulting in Least Net Harm (Argument 4).

There is a way to avoid this controversy. Consider the following argument (call it A4):

P7: If every available climate change strategy would result in net harm to persons, then we ought to adopt that strategy that would result in the least net harm to persons.

P8: Every available climate change strategy would result in net harm to persons.

C4: Therefore we ought to adopt that response to climate change that would result in the least net harm to persons.

Given P7, this argument covers *only* cases in which all available climate change strategies are on balance harmful. In such a situation, so the argument goes, we ought to adopt that available strategy that would result in the least *net* harm to persons. Hence, the proponent of A4 can avoid the controversy encountered by the proponent of A3. Whereas the latter potentially would need to advocate a climate change strategy that is *on balance* less beneficial or even on balance more harmful than some other available strategy, the former holds only that *if* all available strategies are on balance harmful, then we should adopt that strategy that would result in the least net harm.⁹

As an ethical principle, P7 appears to be more plausible than P2, P4, or P6. Intuitively, it seems reasonable that, given an unfortunate choice among several harmful courses of action, one should adopt the course that is on balance least harmful. At the very least, it would seem that one has a very strong, *pro tanto* ethical reason to adopt the least harmful course of action in such situations. Whether one has an all-things-considered ethical reason to do so is perhaps less clear. Some consequentialists might have no quarrel with accepting this. However, others (particularly some non-consequentialists) might object to P7 that the net harm of possible actions is not the only ethically relevant factor to consider. For example, one might hold that distributive justice also matters in determining what climate change strategy ought to be adopted (Svoboda et al. 2011). Imagine two climate change strategies, A and B, both of which result in net harm to persons. Suppose that strategy A results in slightly more net harm than strategy B, but that the harm resulting from A is distributed justly among persons whereas the harm resulting from B is distributed unjustly among persons. Depending on one's commitments in ethical theory, one might hold that strategy A ought to be adopted instead of strategy B. For someone who holds this, even in a scenario in which all available climate change strategies would result in net harm, the fact that some particular strategy would result in less net harm than any other is not an all-things-considered ethical reason to adopt that strategy. This is because that least harmful strategy might result in a distributively unjust outcome.¹⁰ On this view, P7 would not be true as stated, because the fact that every available climate change strategy would result in net harm to persons would not be a *sufficient* condition for it to be the case that we ought to adopt that strategy that would result in the least net harm to persons.

Given the possibility of such a view, I will not here defend the claim that, in situations in which all available climate change strategies would result in net harm, one has an *all-things-*

considered ethical reason to adopt the strategy that would result in the least net harm. Instead, I suggest that, in such situations, one has at least a very strong, *pro tanto* ethical reason to adopt the least harmful strategy. Moreover, even those who deny that P7 is true could treat the principle stated in P7 as what Christine Korsgaard calls a “provisionally universal principle.” Unlike an “absolutely universal principle,” which holds without exception, a provisionally universal principle “applies to every case of a certain sort, unless there is some good reason why not” (Korsgaard 2009, 73). By treating P7 as a provisionally universal principle, one can recognize that it gives one not only a powerful, *pro tanto* ethical reason to adopt the strategy that would result in the least net harm, but also that, *all else of ethical relevance being equal among available strategies*, one ought to adopt that strategy. Further, given that the principle is *provisionally* universal, one can also recognize that there might be competing ethical reasons that could override this principle.

For example, if it is *not* the case that all else of ethical relevance is equal among available strategies (e.g., because the strategy resulting in the least net harm also would result in a distributively unjust outcome whereas other strategies would not), then there might be a good reason as to why the principle does not apply in that case. In such a situation, concerns about distributive justice might give one an ethical reason to adopt a strategy that does *not* result in the least net harm. I do not take a position in this paper on whether that is the case. However, even if this is true, the principle contained in P7, understood as provisionally universal, nonetheless offers a plausible guide: if every available climate change strategy would result in net harm to persons, then we ought to adopt that strategy which would result in the least net harm to persons, *unless* there is a sufficiently good ethical reason why we ought not to do so.¹¹

Perhaps the more controversial claim in A4 is P8, which holds that all available climate change strategies would result in net harm to persons. This premise is uncertain and, like P5 above, might be true at some points in time yet false at others. One might argue that P8 is false at present, perhaps because mitigation or adaptation strategies could still secure outcomes that are on balance beneficial to persons. Yet it is not difficult to envision realistic future scenarios in which P8 is true. For example, imagine that very little progress continues to be made in cutting emissions. At some point in such a future world, a mitigation strategy might be unable to avoid net harm to persons, given the committed climate change of past emissions. Further, the climate change in that future might be so severe that adaptation measures are likewise insufficient to avoid net harm to persons. Arguably, in such a future scenario, P8 would be true, even if it is false at present. Given the lack of progress with mitigation efforts, we might well be on course for such a scenario.

Which Strategy Would Result in the Least Net Harm?

I suggest that the principle contained in P7 is a more helpful guide to considering the ethics of AG than any of the principles contained in A1-3. This is because the principle in P7, whether it is conceived as an absolutely or provisionally universal principle, is more plausible than those in P2, P4, and P6. Moreover, many of those who propose further research on AG argue that it could be an attractive strategy in some climate emergency scenarios, such as those involving certain threshold events or “tipping points” in the climate (Crutzen 2006; Keith, Parson, and Morgan 2010). Such scenarios could be situations in which all available climate change strategies would on balance be harmful. For example, melting permafrost causing the release of vast quantities of methane that in turn lead to increased global warming through a

positive feedback loop (Heimann and Reichstein 2008), the shutdown of the oceans' meridional overturning circulation (Keller, Bolker, and Bradford 2004), or the collapse of polar ice sheets (MacCracken 2009) could result in net harm to persons, such as by causing substantial harm to some while causing little or no benefit for others. Moreover, it might be the case that neither mitigation nor adaptation strategies would be able to avoid this net harm. Especially if a threshold collapse is imminent, it could be too late either for mitigation efforts to avert a tipping point or for adaptation adequately to prepare for it.

In such a scenario, it makes sense to examine which available climate change strategy would result in the least net harm to persons. There is some prima facie evidence that AG could be (or be part of) that strategy in such situations. Given that it has a fast cooling effect (Keith, Parson, and Morgan 2010), AG might be able to avert otherwise-imminent threshold collapses before they occur, thus avoiding the harm that otherwise would result from such events.¹² For example, by rapidly cooling the Earth, AG might be able to halt the release of large quantities of methane from melting Arctic tundra (Morgan and Ricke 2011, 14), avert an impending collapse of polar ice sheets (MacCracken 2009; Irvine et al. 2009), or avoid a shutdown of the meridional overturning circulation (Crutzen 2006, 241; MacCracken 2006). Even if AG would result in net harm to persons, this might be substantially less than the harm that would result from allowing these threshold collapses to occur. Moreover, in such an emergency scenario, the resulting harm of AG might also be less than that resulting from mitigation, adaptation, or other geoengineering strategies, since non-AG strategies are unlikely to avoid or diminish the impacts of an *imminent* threshold collapse. Hence, in such a scenario, AG could be less harmful than all other available strategies. If the principle in P7 is an absolutely universal principle, then AG ought to be deployed in that scenario. This is because all available strategies would result in net harm to

persons, but AG would result in the least net harm. However, even if the principle in P7 is only a provisionally universal principle, it still provides a powerful, *pro tanto* ethical reason to deploy AG in such a scenario, although there might be countervailing ethical reasons why AG ought not be deployed. In that case, the weight of these competing ethical reasons would need to be considered in order to decide what ought to be done.

Further Research.

Much work remains to be done on AG, both scientific and ethical. Further research is needed to reduce uncertainty about the potential impacts of AG. This is important partly because it could reveal that AG would be significantly more or less harmful than currently thought. More research is also needed to devise ways in which the harm resulting from AG could be reduced. Examples might include coupling AG with emissions mitigation (Wigley 2006) in order to reduce the risk of abrupt discontinuation of AG in the future, compensating those who are harmed by AG deployment (Bunzl 2011), and/or establishing an international treaty to reduce the risk of unilateral AG (Victor 2008).¹³

Further research also is needed to address questions concerning both the permissibility of AG deployment (as I have begun to do here) and the ethics of researching AG (Morrow, Kopp, and Oppenheimer 2009). For example, field tests of AG would involve stratospheric aerosol injections and hence carry their own risks of harm, thus raising questions of what conditions must be met in order for field tests to be permissible, whether victims of such tests would be compensated, how research decisions ought to be made, etc. (Tuana et al. 2011). Moreover, Gardiner (2010) raises a number of worries about what he calls the “Arm the Future Argument,” according to which substantial *research* on AG should proceed soon, both because there is a high

probability of a climate emergency scenario in the future and because AG deployment would be the least bad option in such a scenario. Gardiner claims that this argument does not by itself provide a definitive case for conducting AG research because it overlooks important issues, such as that AG research could pose a moral hazard, that there could be non-geoengineering options for dealing with an emergency scenario, and that political inertia could lead to unjust uses of AG if it is researched, among others (Gardiner 2010, 291-295).¹⁴ In this paper, I have not addressed ethical questions regarding research on AG, but of course such questions are extremely important and require careful consideration.

Closing Comments.

I have not argued in this paper that AG ought to be deployed. Instead, I have argued for a conditional claim: *if* we were in a situation in which all available climate change strategies are on balance harmful, and *if* AG was (or was part of) the strategy that would be on balance least harmful, *then* we have at least a powerful, *pro tanto* ethical reason to deploy AG.¹⁵ Moreover, I have suggested that certain climate emergency scenarios, particularly those involving threshold collapses, might be situations in which both these conditions are satisfied, although this would need to be determined in each particular case.

It is probably not the case that both these conditions are satisfied at present, and it very well may be the case that we ought immediately to adopt some climate change strategy that does not involve AG. For example, it is probably the case at present that we ought to pursue aggressive emissions mitigation. However, given current emissions trajectories, it is not implausible to expect that there could be a future scenario in which all available climate change strategies would result in net harm but that AG would be (or be part of) the strategy resulting in

the least net harm. The upshot of all this is that, despite the risks of substantial harm to persons associated with AG, it is not necessarily the case that AG ought not to be deployed. On the contrary, there might be cases in which AG is ethically preferable to available alternatives.¹⁶

¹ Throughout this paper, I focus on harm to persons rather than harm to non-persons. This is because (1) I lack space to consider harm to non-human animals, plants, or ecosystems, and (2) there has been little scientific research on the impacts AG could have on non-humans, thus making it difficult to comment on such potential harm. However, it should not be assumed that I deny that harms to non-persons are morally significant.

² Moreover, given that the atmospheric lifetime of sulfate precursor aerosols is measured in years rather than centuries or millennia (Robock 2000), the cooling effect of AG could be discontinued relatively quickly if it was found to have harmful impacts (Keith, Parson, and Morgan 2010).

³ To take just one example, the Intergovernmental Panel on Climate Change (IPCC) warns, “Towards the end of the 21st century, projected sea-level rise will affect low-lying coastal areas with large populations [in Africa]. The cost of adaptation could amount to at least 5-10% of Gross Domestic Product (GDP)” (IPCC 2007, 13).

⁴ Others disadvantages of AG include the potential for sulfate aerosols to cause ozone depletion (Tilmes, Müller, and Salawitch 2008) and the possibility that researching AG poses a moral hazard (Bunzl 2009).

⁵ I would add that ocean acidification also poses a serious obstacle to AG meeting the requirements of distributive justice, as it could harm some persons much more than others.

⁶ P2 would seem to be entailed by certain strong versions of the precautionary principle. For discussion, see Gardiner (Gardiner 2006).

⁷ Strictly speaking, some of these strategies would not *cause* persons to be harmed but would only allow them to be harmed. However, I lack space in this paper to consider whether there is any morally relevant distinction between doing and allowing (Singer 1972; Scheffler 2004).

⁸ I am not claiming that there is *in fact* some climate change strategy that both causes more harm than some other strategy and is on balance more beneficial than that other strategy. That is an open question.

⁹ Note also that a principle directing one to adopt that strategy resulting in the greatest net benefit would be of no use in such a situation, since no such strategy would be available.

¹⁰ One way to address or at least ameliorate this potential problem is to compensate those victims of AG who receive a distributively unjust allocation of harms. One might argue, for example, that beneficiaries of AG have a responsibility to remunerate those who are harmed by AG. This could go some distance in at least reducing violations of distributive justice.

¹¹ Stated this way, the argument would need to be reformulated in order to maintain the same conclusion (call this A5):

P9: If every available climate change strategy would result in net harm to persons, then we ought to adopt that strategy that would result in the least net harm to persons, unless there is a sufficiently good ethical reason why we ought not to do so.

P10: It is not the case that there is a sufficiently good ethical reason why we ought not to adopt that strategy that would result in the least net harm to persons

P11: Therefore, if every available climate change strategy would result in net harm to persons, then we ought to adopt that strategy that would result in the least net harm to persons.

P12: Every available climate change strategy would result in net harm to persons.

C5: Therefore we ought to adopt that strategy that would result in the least net harm to persons.

I am not here endorsing this argument. P10 is controversial and I do not know whether it is (or, in certain scenarios, would be) true. Moreover, A4 might be sound, in which case the extra premises provided by A5 would be superfluous.

¹² It is currently an open question whether imminent threshold collapses could be predicted prior to their occurrence. Such events might not be recognized with sufficient time left to deploy geoengineering in order to avert them (Tuana et al. 2011). This could make it difficult to *determine* at a particular time whether P8 is true, thus raising serious application problems for the proponent of deploying AG in emergency scenarios.

¹³ Such further research is highly unlikely to remove all uncertainty, of course. Nonetheless, research can reduce uncertainty about these various issues. This is surely a desirable goal, since it would allow us better to assess both the impacts of different strategies and their propensity to result in harm or benefit.

¹⁴ Initially, it might seem that Gardiner's critique also applies to my thesis in this paper, namely that it could be the case that AG ought to be deployed in climate emergency scenarios. However, this is not so. Gardiner's major objection to the "Arm the Future Argument" is that it does not by itself provide a definitive basis to proceed with AG *research*. This may well be true. However, this objection does not challenge my thesis that it could be the case that AG ought to be *deployed* in climate emergency scenarios.

¹⁵ Some have asked whether, given uncertainty about climate change and the impacts of AG, it would be more appropriate to rely on an ethical principle that countenances the risks associated with different strategies. Yet while considerations of risk are important, I view them as distinct from identifying an appropriate ethical principle regarding whether or not AG ought to be deployed in certain scenarios. Due to uncertainty within some scenario, it might be unclear whether or not AG would result in less net harm than other strategies. In that case, there would be a risk involved with deploying AG, since it might turn out to result in more net harm than other strategies. However, this seems to be an epistemic issue, insofar as it could be very challenging in some scenarios to determine the relative harms and benefits of AG.

¹⁶ This has implications not only for potential deployment but also for conducting research on AG, since the prospect for AG to be the ethically preferable strategy in some scenarios arguably provides a *pro tanto* reason to study it further, although there also might be countervailing reasons against AG research (Gardiner 2010).

References.

- Archer, D., and V. Brovkin. 2008. "The Millennial Atmospheric Lifetime of Anthropogenic CO₂." *Climatic Change* 90 (3): 283-297.
- Arneson, Richard. 1989. "Equality and Equal Opportunity for Welfare." *Philosophical Studies* 56 (1): 77-93.
- Barrett, Scott. 2008. "The Incredible Economics of Geoengineering." *Environmental and Resource Economics* 39 (1): 45-54.
- Brewer, Peter G. 2007. "Evaluating a Technological Fix for Climate." *Proceedings of the National Academy of Sciences* 104 (24): 9915-9916.
- Bunzl, Martin. 2009. "Researching Geoengineering: Should Not or Could Not?" *Environmental Research Letters* 4 (4).
- Bunzl, Martin. 2011. "Geoengineering Harms and Compensation." *Stanford Journal of Law, Science & Policy* 4: 70-76.
- Byravan, Sujatha, and Sudhir Chella Rajan. 2010. "The Ethical Implications of Sea-Level Rise Due to Climate Change." *Ethics and International Affairs* 24 (3): 239-260.

- Crutzen, Paul J. 2006. "Albedo Enhancement by Stratospheric Sulfur Injections: A Contribution to Resolve a Policy Dilemma?" *Climatic Change* 77 (3-4): 211-219.
- Doney, S. C. 2006. "The Dangers of Ocean Acidification." *Scientific American* 294 (3): 38-45.
- Doney, Scott C., V. J. Fabry, R. A. Feely, and J. A. Kleypas. 2009. "Ocean Acidification: The Other Co2 Problem." *Annual Review of Marine Science* 1: 169-192.
- Dworkin, Ronald. 1981. "What Is Equality? Part 1: Equality of Welfare." *Philosophy and Public Affairs* 10 (3): 185-246.
- Dworkin, Ronald. 1981. "What Is Equality? Part 2: Equality of Resources." *Philosophy and Public Affairs* 10 (4): 283-345.
- Gardiner, Stephen M. 2004. "Ethics and Global Climate Change." *Ethics* 114 (3): 555-600.
- Gardiner, Stephen M. 2006. "A Core Precautionary Principle*." *Journal of Political Philosophy* 14 (1): 33-60.
- Gardiner, Stephen M. 2010. "Is "Arming the Future" with Geoengineering Really the Lesser Evil? Some Doubts About the Ethics of Intentionally Manipulating the Climate System." In *Climate Ethics*, edited by S. M. Gardiner, S. Caney, D. Jamieson and H. Shue. Oxford: Oxford University Press.
- Goes, Marlos, Klaus Keller, and Nancy Tuana. 2011. "The Economics (or Lack Thereof) of Aerosol Geoengineering." *Climatic Change* In press.
- Heimann, Martin, and Markus Reichstein. 2008. "Terrestrial Ecosystem Carbon Dynamics and Climate Feedbacks." *Nature* 451 (7176): 289-292.
- Hoegh-Guldberg, O., P. J. Mumby, A. J. Hooten, R. S. Steneck, P. Greenfield, E. Gomez, C. D. Harvell, P. F. Sale, A. J. Edwards, K. Caldeira, N. Knowlton, C. M. Eakin, R. Iglesias-

- Prieto, N. Muthiga, R. H. Bradbury, A. Dubi, and M. E. Hatzitolos. 2007. "Coral Reefs under Rapid Climate Change and Ocean Acidification." *Science* 318 (5857): 1737-1742.
- Horton, Joshua. 2011. "Geoengineering and the Myth of Unilateralism: Pressures and Prospects for International Cooperation." *Stanford Journal of Law, Science & Policy* 4: 56-69.
- IPCC. 2007. *Climate Change 2007: Impacts, Adaptation and Vulnerability*: Cambridge: Cambridge University Press.
- IPCC. 2007. *Climate Change 2007: The Physical Science Basis*. Cambridge: Cambridge University Press.
- Irvine, Peter J, Daniel J Lunt, Emma J Stone, and Andy Ridgwell. 2009. "The Fate of the Greenland Ice Sheet in a Geoengineered, High Co₂ World." *Environmental Research Letters* 4 (4): 045109.
- Irvine, Peter J., Andy Ridgwell, and Daniel J. Lunt. 2010. "Assessing the Regional Disparities in Geoengineering Impacts." *Geophysical Research Letters* 37 (18).
- Jamieson, Dale. 1996. "Ethics and Intentional Climate Change." *Climatic Change* 33 (3): 323-336.
- Jones, Andy, Jim Haywood, and Olivier Boucher. 2011. "A Comparison of the Climate Impacts of Geoengineering by Stratospheric So₂ Injection and by Brightening of Marine Stratocumulus Cloud." *Atmospheric Science Letters* 12 (2): 176-183.
- Keith, David. 2001. "Geoengineering the Climate: History and Prospect." *Annual Review of Energy and the Environment* 25: 245-284.
- Keith, David W., Edward Parson, and M. Granger Morgan. 2010. "Research on Global Sun Block Needed Now." *Nature* 463 (7280): 426-427.

- Keller, Klaus, Benjamin M. Bolker, and David F. Bradford. 2004. "Uncertain Climate Thresholds and Optimal Economic Growth." *Journal of Environmental Economics and Management* 48 (1): 723-741.
- Korsgaard, C.M. 2009. *Self-Constitution: Agency, Identity, and Integrity*: Oxford University Press.
- Lobell, David B., Marshall B. Burke, Claudia Tebaldi, Michael D. Mastrandrea, Walter P. Falcon, and Rosamond L. Naylor. 2008. "Prioritizing Climate Change Adaptation Needs for Food Security in 2030." *Science* 319 (5863): 607-610.
- MacCracken, Michael C. 2006. "Geoengineering: Worthy of Cautious Evaluation?" *Climatic Change* 77 (3-4): 235-243.
- MacCracken, Michael C. 2009. "On the Possible Use of Geoengineering to Moderate Specific Climate Change Impacts." *Environmental Research Letters* 4 (4).
- Matthews, H. Damon, and Ken Caldeira. 2007. "Transient Climate--Carbon Simulations of Planetary Geoengineering." *Proceedings of the National Academy of Sciences* 104 (24): 9949-9954.
- Mill, John Stuart. 2008. "Utilitarianism." In *Utilitarianism and on Liberty*: Blackwell Publishing Ltd.
- Morgan, M. G., and K Rieke. 2011. *Cooling the Earth through Solar Radiation Management: The Need for Research and an Approach to Its Governance* Geneva: International Risk Governance Council.
- Morrow, David R., R. E. Kopp, and M. Oppenheimer. 2009. "Toward Ethical Norms and Institutions for Climate Engineering Research." *Environmental Research Letters* 4.
- Nordhaus, W. D. 2001. "Global Warming Economics." *Science* 294.

- Olsaretti, Serena. 2007. "Distributive Justice and Compensatory Desert." In *Desert and Justice*.
Oxford: Oxford University Press.
- Popper, K.R. 2002. *The Open Society and Its Enemies: The Spell of Plato*: Routledge.
- 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 (L02809).
- Rawls, John. 1999. *A Theory of Justice*. revised ed: Harvard University Press. Original edition,
1971.
- Robock, A. 2000. "Volcanic Eruptions and Climate." *Reviews of Geophysics* 38 (2): 191-219.
- Robock, A., L. Oman, and G. L. Stenchikov. 2008. "Regional Climate Responses to
Geoengineering with Tropical SO₂ Injections." *Journal of Geophysical Research--
Atmospheres* 113.
- Robock, Alan. 2008. "20 Reasons Why Geoengineering May Be a Bad Idea." *Bulletin of the
Atomic Scientists* 64 (2): 14-18.
- Ross, Andrew, and H. Damon Matthews. 2009. "Climate Engineering and the Risk of Rapid
Climate Change." *Environmental Research Letters* 4 (4).
- Scheffler, Samuel. 2004. "Doing and Allowing." *Ethics* 114 (2): 215-239.
- Sen, Amartya. 1982. "Equality of What?" In *Choice, Welfare and Measurement*, edited by A.
Sen. Cambridge: Cambridge University Press.
- Shepherd, John. 2009. *Geoengineering the Climate: Science, Governance and Uncertainty*.
London: The Royal Society.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (3): 229-
243.

- Singer, Peter. 2004. *One World: The Ethics of Globalization*. New Haven: Yale University Press.
- Smart, R. N. 1958. "Negative Utilitarianism." *Mind* LXVII (268): 542-543.
- Solomon, Susan, Gian-Kasper Plattner, Reto Knutti, and Pierre Friedlingstein. 2009. "Irreversible Climate Change Due to Carbon Dioxide Emissions." *Proceedings of the National Academy of Sciences* 106 (6): 1704-1709.
- Strong, Aaron, Sallie Chisholm, Charles Miller, and John Cullen. 2009. "Ocean Fertilization: Time to Move On." *Nature* 461 (7262): 347-348.
- Svoboda, Toby, Klaus Keller, Marlos Goes, and Nancy Tuana. 2011. "Sulfate Aerosol Geoengineering: The Question of Justice." *Public Affairs Quarterly* 25 (3): 157-180.
- Teller, E., R. Hyde, M. Ishikawa, Nuckolls J., and Wood L. 2003. Active Stabilization of Climate: Inexpensive, Lowrisk, near-Term Options for Preventing Global Warming and Ice Ages Via Technologically Varied Solar Radiative Forcing. Lawrence Livermore National Library.
- Tilmes, Simone, R. Müller, and R. Salawitch. 2008. "The Sensitivity of Polar Ozone Depletion to Proposed Geoengineering Schemes." *Science* 320 (5880): 1201-1204.
- Trenberth, Kevin E., and Aiguo Dai. 2007. "Effects of Mount Pinatubo Volcanic Eruption on the Hydrological Cycle as an Analog of Geoengineering." *Geophys. Res. Lett.* 34 (15): L15702.
- Tuana, N., R. Sriver, T. Svoboda, R. Tonkononjenkov, P. Irvine, J. Haqq-Misra, and K. Keller. 2011. "Towards Integrated Ethical and Scientific Analysis of Geoengineering: A Research Agenda." *Ethics, Policy & Environment*.
- Victor, D. G., M. G. Morgan, J. Apt, J. Steinbruner, and K. Ricke. 2009. "The Geoengineering Option." *Foreign Affairs* 88 (2): 64-76.

- Victor, David G. 2008. "On the Regulation of Geoengineering." *Oxford Review of Economic Policy* 24 (2): 322-336.
- Wetherald, R. T., Stouffer, R. J., Dixon, K. W. 2001. "Committed Warming and Its Implications for Climate Change." *Geophysical Research Letters* 28 (8): 1535-1538.
- Wigley, T. M. L. 2006. "A Combined Mitigation/Geoengineering Approach to Climate Stabilization." *Science* 314 (5798): 452-454.