

The Ethics of Geoengineering: Moral Considerability and the Convergence Hypothesis

TOBY SVOBODA

ABSTRACT *Although it could avoid some harmful effects of climate change, sulphate aerosol geoengineering (SAG), or injecting sulphate aerosols into the stratosphere in order to reflect incoming solar radiation, threatens substantial harm to humans and non-humans. I argue that SAG is prima facie ethically problematic from anthropocentric, animal liberationist, and biocentric perspectives. This might be taken to suggest that ethical evaluations of SAG can rely on Bryan Norton's convergence hypothesis, which predicts that anthropocentrists and non-anthropocentrists will agree to implement the same or similar environmental policies. However, there are potential scenarios in which anthropocentrists and non-anthropocentrists would seem to diverge on whether a particular SAG policy ought to be implemented. This suggests that the convergence hypothesis should not be relied on in ethical evaluation of SAG. Instead, ethicists should consider the merits and deficiencies of both non-anthropocentric perspectives and the ethical evaluations of SAG such perspectives afford.*

Introduction

Geoengineering, or the intentional manipulation of the Earth's environment on a large scale,¹ could avoid some of the potentially harmful effects of climate change, such as increases in temperature and sea level.² However, geoengineering also threatens substantial harm of its own to both human beings³ and non-human entities. Given this risk of harm, geoengineering strategies require ethical evaluation. In evaluating geoengineering strategies, one might rely on Bryan Norton's convergence hypothesis, according to which anthropocentrists and non-anthropocentrists, despite their divergent theoretical commitments, will agree to implement the same or similar environmental policies.⁴ However, while the convergence hypothesis has many virtues and is worth considering in the context of geoengineering, I argue that it is plausible to expect anthropocentric and non-anthropocentric evaluations of some geoengineering policies to diverge in certain realistic scenarios. Accordingly, the question of which (if any) non-humans deserve moral consideration is relevant for ethical evaluations of geoengineering strategies, given that some such strategy might be morally acceptable by anthropocentric lights but not by non-anthropocentric lights.

This article examines geoengineering in the form of injecting sulphate aerosols into the Earth's stratosphere. In the first section, I sketch the science of sulphate aerosol geoengineering (SAG) and briefly discuss how SAG could avert some of the effects of anthropogenic climate change. In the second section, I consider some of the specific impacts and risks of SAG, including its potential to harm human beings and non-human

organisms. In the third section, I discuss Norton's convergence hypothesis, as well as why it might seem to be an attractive tool in the context of SAG. In the fourth section, I briefly sketch ethical evaluations of SAG from anthropocentric, animal liberationist, and biocentric perspectives, concluding that SAG is *prima facie* ethically problematic from each perspective, although for different reasons. In the fifth section, I develop a SAG policy that minimizes risks to humans and compensates human victims of SAG. I argue that, in realistic scenarios, it is reasonable to expect divergence between anthropocentrists and non-anthropocentrists regarding this policy. Accordingly, I suggest that it is unclear whether the convergence hypothesis correctly predicts anthropocentric and non-anthropocentric verdicts regarding particular SAG policies. I conclude that, in evaluating SAG, ethicists should not ignore the question of which entities are morally considerable, since the answer to that question could make an important difference for whether a particular SAG policy ought to be implemented.

1. Sulphate Aerosol Geoengineering

A geoengineering policy regarding climate change would utilize technological means to alter the planet's environment and thus avoid at least some of the effects of climate change. Geoengineering is distinct from both adapting to the impacts of climate change and attempting to avert climate change by mitigating greenhouse gas emissions, although geoengineering could be combined with mitigation,⁵ adaptation, or both in order to form a hybrid policy. Geoengineering is receiving increased attention from individual scientists and national governments, and the Intergovernmental Panel on Climate Change is set to study geoengineering in its next assessment report.⁶ However, ethicists have only just begun to consider the numerous ethical implications of geoengineering.⁷

Possible geoengineering strategies include solar radiation management techniques, such as cloud-whitening and placing space mirrors in orbit, both of which would reflect some incoming solar radiation into space and thus prevent an increase in global average surface temperature.⁸ Other geoengineering strategies include techniques to reduce concentrations of atmospheric carbon dioxide (CO₂), such as fertilizing the Earth's oceans with iron in order to encourage the growth of phytoplankton blooms, which would absorb and thus reduce atmospheric CO₂.⁹ However, perhaps the most often discussed version is sulphate aerosol geoengineering (SAG), a variety of solar radiation management that would inject sulphate precursor (e.g. SO₂) aerosols into the stratosphere, increasing the Earth's albedo and thus reflecting some incoming solar radiation into space.¹⁰ This would imitate the effects of a large volcanic eruption,¹¹ reduce the amount of solar radiation reaching Earth's surface, and thus induce global cooling. In effect, SAG could provide a planetary thermostat, allowing humans to control the Earth's temperature by altering the concentrations of sulphate aerosols in the stratosphere. This would allow humans to avert some of the effects of climate change, such as rising sea levels caused by melting polar ice sheets.¹²

However, SAG has been critiqued due to its potentially harmful effects for human beings.¹³ It has the potential to harm humans through alterations in regional precipitation, increased ocean acidification (unless SAG is accompanied by aggressive efforts to mitigate CO₂), ozone depletion, unilateral implementation by a powerful state, or abrupt

discontinuation that fosters rapid climate change.¹⁴ Given that SAG would threaten such harm to humans, it clearly requires ethical evaluation. However, no published assessment of SAG considers the ethical implications of its potential to harm non-humans through precipitation change, ocean acidification, and other impacts. It is a serious question whether these potential harms to non-humans should be taken into account when engaging in ethical evaluation of SAG. The answer to this question depends in part on whether non-humans are included within the scope of moral considerability, or whether non-humans deserve moral consideration from moral agents.¹⁵ I consider two non-anthropocentric positions in this article: animal liberationism,¹⁶ which recognizes the moral considerability of all and only sentient entities, and biocentrism,¹⁷ which recognizes the moral considerability of all and only living entities.¹⁸ I assume that a policy is *prima facie* ethically problematic if it threatens substantial harm to morally considerable entities. Further, I assume that if a policy is *prima facie* ethically problematic from a moral agent's perspective, she has an ethical reason to oppose implementing that policy, although this reason could be opposed by countervailing ethical reasons in favour of implementing that policy. Finally, I assume that, given a situation in which all available policies are *on balance* harmful (i.e. result in net harm) to morally considerable entities, a climate change policy ought to be implemented if it threatens less net harm to morally considerable entities than any other available policy. For example, in a climate emergency scenario, it might be the case that all available climate change policies (e.g. mitigation, adaptation, geoengineering, or some combination of these) would be on balance harmful to morally considerable entities. If this is the case, then that available policy which threatens the least net harm to morally considerable entities ought to be implemented.¹⁹

Given these assumptions, I show below that a SAG policy is *prima facie* ethically problematic according to anthropocentric, animal liberationist, and biocentric perspectives. Therefore, proponents of all three positions have ethical reasons to oppose the implementation of SAG. However, the ethical reasons that anthropocentrists, animal liberationists, and biocentrists respectively have against implementing SAG are not identical to one another. This differentiation of ethical reasons could lead to divergent verdicts concerning whether some particular SAG policy ought to be implemented. The potential for such divergence makes it unclear whether the convergence hypothesis correctly predicts anthropocentric and non-anthropocentric positions on regarding SAG implementation. Before turning to either explicit ethical evaluation of SAG or to the question of convergence versus divergence, however, I will first discuss in more detail the potential impacts and risks of SAG.

2. Impacts and Risks of Sulphate Aerosol Geoengineering

A policy that implements SAG has the potential to allow increased ocean acidification and to cause alterations in regional precipitation patterns. SAG also could be unilaterally implemented or unintentionally discontinued. All of these potential impacts and risks could be harmful for human beings, non-human animals, and plant life.

First, due to increased emissions of CO₂ into the atmosphere, the world's oceans are becoming more acidic. Atmospheric CO₂ is absorbed by the oceans and reacts with other chemicals, thus reducing ocean pH and upsetting the chemical balance of ocean

ecosystems.²⁰ Ocean acidity is expected to increase as atmospheric CO₂ continues to accumulate due to anthropogenic emissions.²¹ Granting that SAG would avoid some of the harms associated with climate change, such as rising sea levels due to melting polar ice sheets,²² it does not address ocean acidification because it does not mitigate anthropogenic CO₂ emissions. Instead, this version of geoengineering only compensates for the global warming caused by CO₂ and other greenhouse gases, allowing such gases to continue to accumulate in the atmosphere. If current emissions trends continue, the Earth's oceans will become increasingly acidic even if SAG maintains global temperatures at a predetermined value. Increased ocean acidity would threaten substantial harm to human beings who depend on marine ecosystems, such as those who rely on coral reefs for ecosystem services (e.g. coastal protection) and for income through fisheries and tourism.²³ It also would threaten substantial harm to non-human marine organisms, such as by reducing populations of calcifying organisms (e.g. plankton, crustaceans, and molluscs)²⁴ and thereby potentially disturbing food webs, although more research is needed to understand these complex relations.²⁵ Further, ocean acidification threatens the very existence of coral reefs, since high levels of absorbed CO₂ can kill corals (e.g. through coral bleaching),²⁶ which could substantially harm the flora and fauna (e.g. fish, seabirds, and pteropods) dependent on coral reefs for habitat.²⁷

Second, climate model simulations suggest that SAG could cause substantial reductions in average annual precipitation in some regions, particularly in Africa, South America, and south-eastern Asia.²⁸ While more research is needed on this issue, SAG-induced precipitation change could lead to droughts²⁹ and potentially reduce freshwater availability and agricultural productivity in certain regions,³⁰ such as by disrupting the African and Asian summer monsoons.³¹ Current models and empirical evidence from past volcanic eruptions suggest that SAG-induced precipitation change could substantially harm humans in this regard, such as by creating famines in some regions.³² Likewise, precipitation change could substantially harm terrestrial animals and plants, potentially leading to ecological transformations affecting organisms that depend on specific ecological niches in order to survive and flourish. However, it is deeply uncertain whether and to what degree SAG would cause such harm.³³ Presumably, this would depend in part on how well those affected organisms could adapt to changing environments. Nonetheless, it is plausible to expect that these impacts (e.g. droughts), if they occur, could harm non-human animals and plants in potentially far-reaching ways.

Third, since SAG is estimated to be relatively inexpensive compared to mitigating greenhouse gases,³⁴ and since any state with the necessary technology could implement SAG without international cooperation, there is a risk that some state could deploy SAG unilaterally.³⁵ Unilateral SAG has the potential to be substantially harmful to human beings, especially for non-residents of a state that chooses to implement it.³⁶ This is because SAG can be deployed to a greater or lesser degree by increasing or decreasing the quantity of aerosols injected into the stratosphere, thus determining the degree of global cooling that is induced.³⁷ A rogue state thus could choose to geoengineer a global climate that is in its own perceived self-interest but perhaps contrary to the interests of other states. In that case, even if unilateral SAG should alter the global environment in ways that benefit humans in that state's own territory, these alterations might be substantially harmful for humans in other states or regions. Further, various non-human organisms could be substantially harmed by unilateral SAG as well, such as through the introduction of temperatures, precipitation patterns, or ocean acidity levels that are

contrary to the interests of certain animals or interfere with the biological functioning of flora. Whether and to what extent these harms would be realized are both deeply uncertain,³⁸ but unilateral SAG does pose risks of these harms occurring.

Fourth, implementation of SAG carries a risk of unintentional discontinuation. Since sulphate aerosols have an atmospheric lifespan of only a few years,³⁹ they would need to be replenished continuously by new injections in order to maintain a constant level of global cooling. If this upkeep of aerosol levels should cease (e.g. due to war or terrorism), the aerosols in the stratosphere would disperse within a few years, thus eliminating the cooling counterweight to atmospheric greenhouse gases. This is expected to lead to very rapid global warming as global temperatures 'catch up' to the radiative forcing of CO₂ and other greenhouse gases.⁴⁰ Discontinuous SAG could cause substantial harm to human beings, with the subsequent rate of temperature increase potentially having deleterious impacts on agricultural production and other environmental systems that, as Kintisch writes, 'scientists could only imagine'.⁴¹ It is plausible to expect that the prospect of discontinuous SAG also threatens substantial harm to non-human organisms, given that rapid global warming presumably would have significant impacts on various ecosystems and the organisms who rely on them.⁴² However, the probability of discontinuous SAG is not clear. Barrett thinks it is unlikely because states that implement SAG have powerful incentives to continue it, particularly to avoid the severe economic and environmental problems that discontinuation would cause.⁴³ Alternatively, MacCracken worries that the benefits of continuing SAG might be lost on 'the typical citizen', thus raising concerns that SAG could be discontinued due to ignorance or apathy.⁴⁴ At present, it is not clear how probable discontinuous SAG would be, but it is clear that such discontinuation could result in substantial harm to human beings and non-human organisms.

3. A Place for the Convergence Hypothesis?

I will argue below that SAG is *prima facie* ethically problematic from anthropocentric, animal liberationist, and biocentric perspectives. Proponents of each of these perspectives have a unique set of ethical reasons to oppose implementation of a SAG policy. Yet one might ask why we should bother evaluating SAG from all three of these perspectives. Given that SAG threatens substantial harm to human beings and various non-human organisms, it might seem that anthropocentrists and non-anthropocentrists would converge in rejecting a SAG policy. Further, it might seem that these concordant ethical evaluations of SAG help to confirm Bryan Norton's convergence hypothesis, which holds that both anthropocentrists and non-anthropocentrists will at least roughly agree on what environmental policies ought to be enacted.⁴⁵ According to Norton, this is because 'policies serving the interests of the human species as a whole, and in the long run, will serve also the "interests" of nature, and vice versa'.⁴⁶ Thus, although anthropocentrists and non-anthropocentrists might differ in a given case on the theoretical justifications they offer in defence of their policy recommendations, they will converge in endorsing or rejecting the same policy. The convergence hypothesis is certainly worth considering in the context of geoengineering, because in many other cases it seems to provide a powerful, practical guide to environmental policy-making. Thus, it might appear superfluous to evaluate SAG from three different ethical perspectives. As the

convergence hypothesis predicts, perhaps anthropocentrists and non-anthropocentrists of various stripes will agree in rejecting SAG as a climate change policy. Such convergence for or against some policy would be a happy result, since it would mean that one could bracket the complex debates regarding anthropocentrism versus non-anthropocentrism and still render appropriate ethical evaluations of that policy.

This would be too quick, however. Below, I do not argue that anthropocentrists, animal liberationists, and biocentrists would all agree that SAG ought not to be implemented. I will argue only that SAG is *prima facie* ethically problematic from all three perspectives and that proponents of them have ethical reasons to oppose SAG's implementation. Importantly, some of the ethical reasons non-anthropocentrists have for opposing implementation of SAG are not shared by anthropocentrists. Now, it is plausible to expect that policy choices by those who have differing reasons will sometimes diverge from one another. As I argue below, there are realistic scenarios in which it is reasonable to expect that some particular SAG policy ought to be implemented according to anthropocentric lights, whereas it ought not to be implemented according to non-anthropocentric lights. Depending both on how a particular SAG policy is constructed and on various other circumstances, anthropocentrists and non-anthropocentrists could disagree on whether that policy ought to be implemented in some particular scenario, or so I argue.

A crucial point of potential divergence between anthropocentrists and non-anthropocentrists regarding SAG is that human beings can be compensated economically for many (although presumably not all) of the harms they suffer, whereas non-humans cannot be so compensated. Obviously, non-human organisms cannot benefit from economic remuneration for harms. While a human being who suffers from a fishery becoming less productive due to ocean acidification can be awarded damages that restore her losses, there is no similar compensatory mechanism available for marine animals who suffer from a lack of food or habitat due to ocean acidification. Later in this article, I sketch a SAG policy involving compensation to human victims. Arguably, under certain conditions, this policy would be endorsed by consistent anthropocentrists yet rejected by consistent non-anthropocentrists.

None of this is to suggest that the convergence hypothesis is not an important practical guide to environmental policy-making. I suspect that there is often a convergence among the 'interests' of human beings and non-human entities. In such cases, Norton is surely right that one need not settle difficult theoretical disputes about the scope of moral considerability in order to decide which policies ought to be adopted. However, while the convergence hypothesis provides a helpful general guide to environmental policy-making, it might not hold in every particular case. Some of Norton's own formulations of the convergence hypothesis seem to allow for this. For example, he writes that, according to the convergence hypothesis, advocates of both anthropocentric and non-anthropocentric theories 'would approve many, perhaps all, of the same policies'.⁴⁷ Norton also has formulated the convergence hypothesis as 'the view that anthropocentrists and non-anthropocentrists will tend to propose similar policies'.⁴⁸ Given these two formulations, Norton leaves it open whether anthropocentrists and non-anthropocentrists would agree on *all* policies, implying that convergence on 'most' policies or a tendency to accept 'similar' policies is sufficient to establish the truth of the convergence hypothesis. On either of these formulations, accepting the convergence hypothesis is compatible with holding that there are some cases in which anthropocentric

and non-anthropocentric perspectives would counsel divergent policies. I suggest that, in certain plausible scenarios, the question of whether to implement SAG constitutes such a case. Thus, despite its many virtues as a general guide to environmental policy, the convergence hypothesis might not hold here. This makes it worth considering the particular reasons anthropocentrists and non-anthropocentrists have for endorsing or rejecting certain SAG policies.

4. Evaluating Sulphate Aerosol Geoengineering from Anthropocentric, Animal Liberationist, and Biocentric Perspectives

Among the few papers that examine the ethical implications of SAG,⁴⁹ all focus on its potential to harm human beings. From an anthropocentric perspective, SAG's potential to allow ocean acidification,⁵⁰ to alter regional precipitation patterns, to be deployed unilaterally, and to be abruptly discontinued all make it *prima facie* ethically problematic, given that each threatens substantial harm to human beings. Hence, given the assumption made above, all four of these risks provide ethical reasons for anthropocentrists to oppose SAG's implementation. However, the presence of these ethical reasons does not establish definitively that SAG ought not to be implemented. There might be counter-vailing ethical reasons in favour of implementing SAG that, from an anthropocentric perspective, outweigh the ethical reasons against implementing it. For example, perhaps all realistic options for responding to climate change are *prima facie* ethically problematic, and perhaps some SAG policy would result in the least net harm to human beings.⁵¹ Under such conditions, all things considered, it might be the case that some version of SAG ought to be implemented despite its threat of harm to human beings. Yet SAG has the potential to result in substantial harm to non-human animals and flora as well, as noted above. Accordingly, SAG is *prima facie* ethically problematic from various non-anthropocentric perspectives, with non-anthropocentrists having ethical reasons to oppose its implementation.⁵² However, proponents of differing views on the scope of moral considerability have distinct sets of ethical reasons for opposing SAG. In order to illustrate this, I briefly sketch two influential versions of animal liberationism and biocentrism.

According to Peter Singer's animal liberationism, all and only sentient entities are morally considerable, given that all and only such entities have interests that can be satisfied or frustrated by moral agents.⁵³ In particular, all and only sentient entities have 'subjective experiences', which are necessary in order for them to feel pleasure or pain.⁵⁴ Since plants lack subjective experiences, they do not have interests and thus do not deserve moral consideration. Conversely, many humans and non-human animals do have subjective experiences and hence interests, thus making them deserving of moral consideration. According to Singer's non-anthropocentric view, it follows that moral agents ought to take into account the interests of both sentient humans and sentient non-human animals, such as in ethical assessment of SAG policies.

This view has implications that distinguish animal liberationist ethical evaluations of SAG from anthropocentric ones. The fact that SAG threatens substantial harm to sentient non-human organisms (e.g. seabirds in the case of ocean acidification, terrestrial animals in the case of precipitation change) provides animal liberationists with an ethical reason to oppose implementing SAG. Importantly, anthropocentrists lack this particular

reason, given that they do not take sentient non-humans to be morally considerable and thus do not view substantial harm to such entities as *prima facie* ethically problematic. As with anthropocentrism, however, animal liberationists' reasons for opposing SAG do not guarantee that it ought not to be implemented. Again, it might be the case that there are countervailing ethical reasons that favour implementing SAG, and these reasons might outweigh the ethical reasons against implementing it. However, from an animal liberationist perspective, it is clear that a policy that avoids some or all of these risks of substantial harm to sentient animals, all else being equal, is ethically preferable to SAG.

According to Paul Taylor's biocentrism, every living entity is morally considerable because each is a 'teleological center of life', where an entity is a teleological centre of life if 'its internal functioning as well as its external activities are all goal-oriented, having the constant tendency to maintain the organism's existence through time and to enable it successfully to perform those biological operations whereby it reproduces its kind and continually adapts to changing environmental events and conditions'.⁵⁵ This includes all living organisms because they are all goal-directed in the relevant sense, i.e. they tend to perform their biological functions, reproduce their kind, and adapt to their environments. Hence, in addition to humans and sentient non-human animals, non-sentient animals and plants are all morally considerable on Taylor's account. To evaluate SAG from this non-anthropocentric perspective requires taking into account potential harms to all living entities who would be affected by SAG.

Biocentric ethical evaluations of SAG are distinct from both anthropocentric and animal liberationist evaluations. First, as with animal liberationists, the fact that SAG threatens substantial harm to non-human sentient organisms makes SAG *prima facie* ethically problematic by biocentric lights, thereby affording biocentrists an ethical reason against implementing SAG. This is a reason that anthropocentrists do not share, given that they do not view harm to non-humans as ethically problematic in its own right. Second, the fact that SAG threatens substantial harm to non-animal organisms (e.g. corals in the case of ocean acidification, terrestrial flora in the case of precipitation change) gives biocentrists another ethical reason to oppose implementing SAG. This is a reason that neither anthropocentrists nor animal liberationists share, insofar as proponents of these views do not take harm to non-sentient organisms to be ethically problematic in its own right.

However, as with anthropocentrism and animal liberationism, these ethical reasons of biocentrists are not necessarily decisive against the implementation of SAG. Again, this is because there might be countervailing ethical reasons in favour of implementing it. For example, increased concentrations of atmospheric CO₂ can facilitate the flourishing of plants.⁵⁶ Since SAG would not mitigate CO₂ emissions, its potential to benefit plants by allowing increasing concentrations of this greenhouse gas arguably gives biocentrists an ethical reason in favour of implementing SAG. Further, the presence of atmospheric aerosols also might benefit plants by producing higher levels of diffuse sunlight, thus increasing photosynthesis.⁵⁷ Of course, these potential benefits to flora could be outweighed by potential harms, such as precipitation changes that alter their environments. To determine whether SAG ought to be implemented would involve considering the competing ethical reasons for and against deployment, which is beyond the scope of this article. Nonetheless, from a biocentric perspective, it is clear that a policy that avoids some or all of these risks of substantial harm to non-sentient organisms, all else being equal, is ethically preferable to SAG.

5. Policy A and the Potential for Divergence

Given these distinct sets of ethical reasons, there might be cases in which anthropocentrists and non-anthropocentrists would disagree on whether some SAG policy ought to be implemented. Consider an organized, international response to climate change (call it 'Policy A') that contains the following features: (1) it implements SAG while allowing anthropogenic greenhouse gas emissions to continue unmitigated, (2) it minimizes the risk of precipitation changes that would be harmful to humans, (3) it provides fair economic compensation for human beings who are harmed by SAG's effects, and (4) it contains effective safeguards against discontinuous and unilateral geoengineering. There are, of course, challenges to constructing and implementing a policy with these features. For example, although SAG could be implemented with varying degrees of aerosol injections, some of which could affect precipitation levels in potentially less harmful ways than others,⁵⁸ it might be difficult for the international community to reach consensus on a particular degree of SAG to pursue. Further, determining appropriate compensatory mechanisms is not straightforward, as there would be difficult questions about who should compensate whom, how much victims should be compensated, and what conditions must be met in order to merit compensation. Moreover, minimizing the risk of accidental discontinuation of SAG would require careful planning, perhaps by varying the technologies and resources needed for injecting sulphate aerosols into the stratosphere and by having effective contingency plans in place to respond to various scenarios that might occur. Finally, it presumably would not be easy to institute and enforce a ban on unilateral geoengineering,⁵⁹ although an international treaty with deterrent provisions (such as economic sanctions) is a possibility. Despite these challenges, it is possible for humans to adopt such a policy, and since I am not advocating Policy A but only using it to illustrate a potential divergence between anthropocentrists and non-anthropocentrists, meeting these challenges is both beyond the scope of this article and unnecessary for my argument. Finally, to be clear, I do not assume that all human interests can be codified adequately in economic terms, nor do I assume that all SAG-related harm to humans can be recompensed adequately through economic compensation. Instead, I suggest both that such compensation can go a substantial distance toward alleviating harm to human victims of SAG and that a SAG policy involving such compensation is ethically preferable to one that lacks compensation.

For anthropocentrists, Policy A removes or reduces many of the most ethically problematic features of SAG. By minimizing risks to humans and compensating those who are harmed, anthropocentrists have fewer (and perhaps less strong) ethical reasons to oppose implementing SAG than they otherwise would possess. This is not to say that there are no ethical problems with Policy A. However small, there would still be risks of discontinuous and unilateral SAG. Further, appropriate compensation for harm might not be possible in all cases, such as death or the loss of cultural practices. Moreover, given a 'weak' version of anthropocentrism that attributes value to non-human nature,⁶⁰ it might not be possible to compensate humans for the loss of certain irreplaceable natural entities or phenomena, such as an endangered species that is valued aesthetically by human beings.⁶¹ These considerations arguably give anthropocentrists ethical reasons to oppose implementing SAG. However, under certain conditions, it might be the case that all available climate change policies have ethical problems, and Policy A might

prove the least problematic one available. This probably is not the case at present. Presumably, it is not too late for a mitigation policy effectively to avoid many of the harmful effects of climate change, and arguably such a mitigation policy is ethically preferable to Policy A, in part because the former involves less risk of substantial harm to humans.

However, if humans continue to make little progress in mitigating greenhouse gas emissions, there could be a future scenario in which climate change is beginning to have substantially harmful impacts, such as drastic sea-level rise, a substantial increase in the frequency and intensity of dangerous weather events, and decreased food production in some regions.⁶² In such a scenario, given the inertia of the climate system and the long atmospheric lifetime of CO₂,⁶³ a mitigation policy would be unable to counter the current and imminent harmful effects of climate change. Under these conditions, due to the potential of SAG to reduce global temperatures quickly and avoid many otherwise-impending harms, Policy A might be the least ethically problematic policy available from an anthropocentric perspective. Indeed, some scientists who urge more research on SAG tout its ability to respond quickly to impending climate change should humans continue to avoid a path of aggressive emissions mitigation.⁶⁴ For these reasons, assuming an anthropocentric perspective, arguably there are potential scenarios in which Policy A ought to be implemented, even granting that there might be harms that cannot be economically compensated for, such as those tied to the loss of weakly anthropocentric values.

It is not obvious, however, that non-anthropocentrists would concur in this ethical evaluation of Policy A. First, Policy A would allow increased ocean acidification, which, as noted above, would probably harm many marine calcifying organisms as well as marine animals and plants that rely on such organisms for habitat and food. Policy A also could alter regional precipitation levels, potentially causing drought and thus harming non-human animals and plants in affected regions. In accordance with Policy A, nothing would be done to avoid or address these potential harms to non-humans. Further, unlike humans who are harmed by these impacts, non-human animals and plants could not be compensated economically for these harms, even if humans wished to do so.⁶⁵ This is a serious ethical problem from both animal liberationist and biocentric perspectives. Accordingly, proponents of these perspectives have various ethical reasons to oppose implementing Policy A, namely that it threatens substantial harm to non-human animals and flora. Conversely, anthropocentrists lack these ethical reasons, viewing this harm to non-humans as ethically irrelevant in its own right.⁶⁶

This substantial differentiation among the ethical reasons possessed by anthropocentrists and non-anthropocentrists regarding Policy A makes it unclear whether the convergence hypothesis is correct in predicting that both would agree concerning whether particular SAG policies ought to be implemented. While this lack of clarity certainly does not refute the convergence hypothesis, much less discredit it as a helpful guide in many cases, it does suggest that ethicists should not rely on it as a shortcut to ethical evaluation of SAG. This is because it could be the case that some SAG policies (e.g. Policy A), under certain conditions, ought to be implemented given an anthropocentric perspective but ought not to be implemented given a non-anthropocentric perspective. To be clear, I am not claiming merely that it is logically possible for such divergence to occur in some scenario that one can imagine. Rather,

it is arguably the case that such divergence would occur in a plausible scenario in the near future, namely one in which dangerous climate change has begun. In such a situation, there is good reason to suspect that anthropocentrists and non-anthropocentrists would prefer different policies.

One might object that, in a scenario in which dangerous climate change is already occurring and threatens to become much worse, implementing SAG would result in the least net harm not only to human beings, but to non-human animals and flora as well. According to this objection, there would be a convergence among anthropocentric and non-anthropocentric evaluations of SAG after all, insofar as a non-anthropocentrist might hold that Policy A ought to be adopted because it would result in less net harm to non-humans than any other available option. Presumably, we can imagine cases in which such convergence would occur. However, the important question here is whether there would be such convergence in realistic future scenarios. A major reason to doubt this is that the harm SAG causes to human beings can be assuaged through economic compensation, but the harm SAG causes to non-humans cannot be so assuaged. In some scenario, SAG might be acceptable to anthropocentrists only if humans who are harmed receive appropriate compensation. For example, in some future scenario, Policy A might result in the least net harm to humans only because compensation reduces the harm that otherwise would result from implementing that same policy without compensation. However, given that such compensation cannot reduce harm to non-humans, it is plausible to expect that there will be cases in which the least (on balance) harmful strategy vis-à-vis humans is Policy A, whereas the least (on balance) harmful strategy vis-à-vis non-humans is not Policy A. If so, then it is likewise plausible to expect divergence between the policy preferences of anthropocentrists and non-anthropocentrists in such cases.

While the foregoing is not sufficient to falsify it, it is at best unclear whether the convergence hypothesis correctly predicts verdicts regarding SAG. This suggests that the convergence hypothesis should not be assumed when ethically evaluating SAG. Instead, ethicists should consider the respective merits and deficiencies of various anthropocentric and non-anthropocentric evaluations. If non-humans deserve moral consideration, then whether a SAG policy threatens substantial harm to non-humans could be an important factor in whether that policy ought to be implemented. This possibility provides sufficient ground for ethicists to take non-anthropocentric evaluations of SAG seriously, whether or not one ultimately agrees with such evaluations. Given the recent surge of interest in SAG and other geoengineering strategies, appropriate ethical evaluation may be an urgent matter.

Toby Svoboda, Department of Philosophy, Fairfield University, 1073 North Benson Road, Fairfield, CT 06824, USA. tobysvoboda@gmail.com

Acknowledgements

I would like to thank Peter J. Irvine, two anonymous reviewers, and Christopher Bennett for very helpful feedback on drafts of this article.

NOTES

- 1 David Keith, 'Geoengineering the climate: History and prospect', *Annual Review of Energy and the Environment* 25 (2001): 245–284.
- 2 Michael MacCracken, 'On the possible use of geoengineering to moderate specific climate change impacts', *Environmental Research Letters* 4,4 (2009).
- 3 Alan Robock, '20 reasons why geoengineering may be a bad idea', *Bulletin of the Atomic Scientists* 64,2 (2008a): 14–18.
- 4 Bryan Norton, *Toward Unity Among Environmentalists* (New York: Oxford University Press, 1991), p. 240.
- 5 T. M. L. Wigley, 'A combined mitigation/geoengineering approach to climate stabilization', *Science* 314,5798 (2006): 452–454.
- 6 Paul J. Crutzen, 'Albedo enhancement by stratospheric sulfur injections: A contribution to resolve a policy dilemma?' *Climatic Change* 77,3–4 (2006): 211–219; John Shepherd, *Geoengineering the Climate: Science, Governance and Uncertainty* (London: The Royal Society, 2009); Wigley op cit.; Charles Hanley, 'As climate talks drag on, more ponder techno-fixes', *ABC News* 2010, online at <http://abcnews.go.com/print?id=12312227>; Frank Rusco & John Stephenson, *Climate Change: A Coordinated Strategy Could Focus Federal Geoengineering Research and Inform Governance Efforts* (Washington, DC: United States Government Accountability Office, 2010).
- 7 Martin Bunzl, 'Researching geoengineering: Should not or could not?' *Environmental Research Letters* 4, 4 (2009); Stephen M. Gardiner, 'Is "arming the future" with geoengineering really the lesser evil? Some doubts about the ethics of intentionally manipulating the climate system' in S. M. Gardiner, S. Caney, D. Jamieson & H. Shue (eds) *Climate Ethics* (Oxford: Oxford University Press, 2010); Dale Jamieson, 'Ethics and intentional climate change', *Climatic Change* 33,3 (1996): 323–336; Toby Svoboda, Klaus Keller, Marlos Goes & Nancy Tuana, 'Sulfate aerosol geoengineering: The question of justice', *Public Affairs Quarterly* 25,3 (2011): 157–180.
- 8 Oliver Morton, 'Great white hope', *Nature* 458,7242 (2009): 1097–1100.
- 9 Morton op cit.
- 10 Crutzen op cit.
- 11 Peter G. Brewer, 'Evaluating a technological fix for climate', *Proceedings of the National Academy of Sciences* 104,24 (2007): 9915–9916.
- 12 Peter J Irvine, Daniel J Lunt, Emma J Stone & Andy Ridgwell, 'The fate of the Greenland ice sheet in a geoengineered, high CO2 world', *Environmental Research Letters* 4,4 (2009); Michael MacCracken, 'Geo-engineering: Worthy of cautious evaluation?' *Climatic Change* 77,3–4 (2006): 235–243.
- 13 Robock 2008a op cit.
- 14 Marlos Goes, Klaus Keller & Nancy Tuana, 'The economics (or lack thereof) of aerosol geoengineering', *Climatic Change* 3,4 (2011): 791–825; Peter J. Irvine, Andy Ridgwell & Daniel J. Lunt, 'Assessing the regional disparities in geoengineering impacts', *Geophysical Research Letters* 37,18 (2010); H. Damon Matthews & Ken Caldeira, 'Transient climate — carbon simulations of planetary geoengineering', *Proceedings of the National Academy of Sciences* 104, 24 (2007): 9949–9954; A. Robock, L. Oman & G. L. Stenchikov, 'Regional climate responses to geoengineering with tropical SO2 injections', *Journal of Geophysical Research — Atmospheres* 113, D16 (2008):15; Robock 2008a op cit.; Svoboda, Keller, Goes & Tuana op cit.; Simone Tilmes, R. Müller & R. Salawitch, 'The sensitivity of polar ozone depletion to proposed geoengineering schemes', *Science* 320,5880 (2008): 1201–1204; David G. Victor, 'On the regulation of geoengineering', *Oxford Review of Economic Policy* 24,2 (2008): 322–336.
- 15 Kenneth Goodpaster, 'On being morally considerable', *The Journal of Philosophy* 75,6 (1978): 308–325.
- 16 Peter Singer, *Animal Liberation* (New York: Harper Collins, 2002).
- 17 Goodpaster op cit.; Paul Taylor, *Respect for Nature: A Theory of Environmental Ethics* (Princeton, NJ: Princeton University Press, 1986).
- 18 Of course, there are various other non-anthropocentric positions one might consider as well, such as ecocentric views that are concerned with the integrity or flourishing of ecosystems rather than the welfare or flourishing of individual organisms. While SAG could impact marine and terrestrial ecosystems in potentially harmful ways, limitations of space prevent me from considering this matter in detail. Moreover, as I illustrate later in this article, focusing on animal liberationism and biocentrism is sufficient to show that there is a potential divergence between anthropocentric and non-anthropocentric positions regarding SAG policies. Hence, considering other non-anthropocentric positions is not necessary in order to support my claim that ethical evaluations of SAG should not rely on the convergence hypothesis, instead considering the question

- of anthropocentrism versus non-anthropocentrism. For a defence of ecocentrism, see J. Baird Callicott, 'Holistic environmental ethics and the problem of ecofascism' in his *Beyond the Land Ethic: More Essays in Environmental Philosophy* (Albany, NY: SUNY Press, 1999): 59–76.
- 19 It should be noted here that I am not assuming negative utilitarianism, or the view that moral agents ought to adopt those policies that produce the least harm. Rather, my assumption is that, in cases in which all available policies are on balance harmful, moral agents ought to adopt that policy which is on balance the least harmful. For discussion and defence of this position, see Toby Svoboda, 'Is aerosol geoengineering ethically preferable to other climate change strategies?' *Ethics & the Environment* (forthcoming).
 - 20 Scott C. Doney, V. J. Fabry, R. A. Feely & J. A. Kleypas, 'Ocean acidification: The other CO₂ problem', *Annual Review of Marine Science* 1 (2009): 169–192; J. Raven *et al.*, *Ocean Acidification Due to Increasing Atmospheric Carbon Dioxide* (London: Royal Society, 2005).
 - 21 James C. Orr *et al.*, 'Anthropogenic ocean acidification over the twenty-first century and its impact on calcifying organisms', *Nature* 437,7059 (2005): 681–686.
 - 22 MacCracken *op cit.*
 - 23 O. Hoegh-Guldberg *et al.*, 'Coral reefs under rapid climate change and ocean acidification', *Science* 318,5857 (2007): 1737–1742.
 - 24 Singer *op cit.* 2003, p. 57.
 - 25 Victoria J. Fabry, Brad A. Seibel, Richard A. Feely & James C. Orr, 'Impacts of ocean acidification on marine fauna and ecosystem processes', *ICES Journal of Marine Science: Journal du Conseil* 65,3 (2008): 414–432; see John M. Guinotte & Victoria J. Fabry, 'Ocean acidification and its potential effects on marine ecosystems', *Annals of the New York Academy of Sciences* 1134,1 (2008): 320–342.
 - 26 Hoegh-Guldberg *et al.* *op cit.*
 - 27 Hoegh-Guldberg *et al.* *op cit.*; Orr *et al.* *op cit.*
 - 28 Matthews & Caldeira *op cit.*; Robock *et al.* 2008 *op cit.*
 - 29 K. E. Trenberth & A. Dai, 'Effects of Mount Pinatubo volcanic eruption on the hydrological cycle as an analog of geoengineering', *Geophysical Research Letters* 34 (2007).
 - 30 Although cf. J. Pongratz, D. B. Lobell, L. Cao & K. Caldeira. 'Crop yields in a geoengineered climate', *Nature Climate Change* 2 (2012): 101–105.
 - 31 Robock *et al.* 2008 *op cit.*
 - 32 Brewer *op cit.*, 9915.
 - 33 See Lynn M. Russell *et al.*, 'Ecosystem impacts of geoengineering: A review for developing a science plan', *Ambio* (2012). Online at: <http://www.springerlink.com/content/u562q0218461k416/>
 - 34 Barrett *op cit.*; Alan Robock, 'Whither geoengineering?' *Science* 320 (2008b): 1166–1167; Alan Robock, Allison Marquardt, Ben Kravitz & Georgiy Stenchikov, 'The benefits, risks, and costs of stratospheric geoengineering', *Geophysical Research Letters* 36,19 (2009).
 - 35 Victor *op cit.*
 - 36 See Victor *op cit.*
 - 37 Irvine *et al.* *op cit.*, 2010.
 - 38 See Russell *et al.* *op cit.*
 - 39 P. J. Rasch, P. J. Crutzen & D. B. Coleman, 'Exploring the geoengineering of climate using stratospheric sulfate aerosols: The role of particle size', *Geophysical Research Letters* 35 (2008).
 - 40 Goes, Keller & Tuana *op cit.*
 - 41 Eli Kintisch, *Hack the Planet: Science's Best Hope — or Worst Nightmare — for Averting Climate Catastrophe* (Hoboken, NJ: John Wiley and Sons, 2010), p. 74.
 - 42 Russell *et al.* *op cit.*
 - 43 Barrett *op cit.*, p. 50.
 - 44 MacCracken *op cit.*, p. 238.
 - 45 Norton *op cit.*, 1991.
 - 46 Norton *op cit.*, 1991, p. 240.
 - 47 Bryan Norton, *Sustainability: A Philosophy of Adaptive Ecosystem Management* (Chicago, IL: University of Chicago Press, 2005), p. 508.
 - 48 Bryan Norton, 'Convergence, noninstrumental value and the semantics of 'love': Comment on McShane', *Environmental Values* 17,1 (2008): 5.
 - 49 Gardiner *op cit.*; Goes, Keller & Tuana *op cit.*; Robock 2008a *op cit.*; Svoboda, Keller, Goes & Tuana *op cit.*
 - 50 While SAG itself would not cause increased ocean acidification, a SAG-only policy (i.e., one accompanied by neither reductions in CO₂ emissions nor removal of atmospheric CO₂) would allow ocean acidity to

- increase. I lack space in this article to consider whether there is a morally relevant distinction between doing and allowing. See Samuel Scheffler, 'Doing and allowing', *Ethics* 114,2 (2004): 215–239.
- 51 See Svoboda op cit., forthcoming.
- 52 It should be noted that my analysis does not overlook the possibility of weakly anthropocentric views, which could afford some kind of value to non-human nature and thus allow that human beings can be harmed simply through the loss of valued natural entities or phenomena. For example, ocean acidification threatens to harm human beings not only economically but also by destroying natural landscapes and biotic communities that some humans value for their own sakes. I consider some further implications of weak anthropocentrism later in this article. For discussion of a weakly anthropocentric view, see Eugene Hargrove, 'Weak anthropocentric intrinsic value', *The Monist* 75,2 (1992): 183–207.
- 53 Singer op cit., 2002.
- 54 Peter Singer, 'Not for humans only: The place of nonhumans in environmental issues' in A. Light & H. Rolston, III (eds) *Environmental Ethics: An Anthology* (Malden, MA: Blackwell, 2003), p. 57.
- 55 Taylor op cit., pp. 121–122.
- 56 Amanda Pereira De Souza *et al.*, 'Elevated CO₂ increases photosynthesis, biomass and productivity, and modifies gene expression in sugarcane', *Plant, Cell & Environment* 31,8 (2008): 1116–1127.
- 57 Lina M. Mercado *et al.*, 'Impact of changes in diffuse radiation on the global land carbon sink', *Nature* 458,7241 (2009): 1014–1017.
- 58 See Irvine *et al.* 2010 op cit.
- 59 Although cf. Joshua Horton, 'Geoengineering and the myth of unilateralism: pressures and prospects for international cooperation', *Stanford Journal of Law, Science & Policy* 4 (2011): 59–69.
- 60 See Hargrove op cit.
- 61 This is a controversial matter, of course. Perhaps weak anthropocentrists can be compensated for the loss of irreplaceable, valued natural entities, such as by providing them with the economic means to experience and value other natural entities. I do not take a position on whether this is possible, as doing so would require much further discussion than I have space for here.
- 62 Intergovernmental Panel on Climate Change, *Climate Change 2007: The Physical Science Basis* (Cambridge: Cambridge University Press, 2007).
- 63 See D. Archer & V. Brovkin, 'The millennial atmospheric lifetime of anthropogenic CO₂', *Climatic Change* 90,3 (2008): 283–297.
- 64 See Crutzen op cit.
- 65 Non-humans could receive some kind of remuneration, e.g. in the form of relocation, but this would not be the economic compensation under discussion here.
- 66 This does not mean that anthropocentrists ignore harms to non-humans if those non-humans are valued by human beings. For example, an anthropocentrist can disapprove of someone who injures another person's pet, but in that case it is the harm to the other person (insofar as she loses something of value to herself) that is morally relevant, not the harm to the pet.