Sleepwalking Into Lock-in? Avoiding Wrongs to Future People in the Governance of Solar Radiation Management Research

Catriona McKinnon

Professor of Political Theory

University of Reading

Whiteknights, Reading RG6 6AA

c.mckinnon@reading.ac.uk

+44(0)118 378 8502

Abstract

Arguments are advanced for two ways in which we can avoid the reckless endangerment of future people in the governance of solar radiation management (SRM) research, which could happen through lock-in to SRM deployment from research. SRM research is at an early stage, one at which the mechanisms of lock-in could start to operate. However, lock-in fit to endanger future people could be slowed or stopped through targeted governance. Governance of SRM research that does not include provisions to detect, slow, or stop lock-in fails the test of an intergenerationally adequate precautionary principle, and research governed without these provisions cannot itself be justified as a precaution against the impacts of climate change.

Key words: Solar radiation management, intergenerational justice, precaution, lock-in, geoengineering, climate justice.

Introduction

Since Paul Crutzen made solar radiation management (SRM) a respectable topic for enquiry by the scientific community (Crutzen 2006), there has been a huge increase in the interest paid to SRM by natural scientists, and by scholars concerned to assess the social and economic costs and benefits of the technologies.[1] Ethical questions that arise for research have been given far less attention than those that arise for deployment.[2] This may be because it is taken as obvious that there is nothing special about research into SRM that grounds distinctive ethical concerns that should inform governance of the research, especially in the early stages of modelling, lab experiments, and small-scale field

[1] Despite the fact that only 8% of the lay public were able to correctly define the term 'geoengineering' in a 2010 survey (Corner *et al.* 2012). More recent work suggests greater levels of public awareness (Wibeck *et al.* 2017).
[2] There are exceptions (Hamilton 2013, Morrow *et al.* 2013).

testing. On this view, singling out SRM research as requiring governance in line with distinctive ethical concerns without applying that same standard to other scientific research programmes is lopsided and unfair (Heyward 2015). The zeitgeist with respect to SRM research seems to be that the most that is justified right now is self-regulation, perhaps using guiding principles and codes of conduct, such as the Oxford Principles (The Royal Society 2009, Rayner *et al.* 2013, Hubert *et al.* 2016). Furthermore, many reports on governance to date advocate an 'allowed zone' for research (Cicerone 2006, Morrow 2017, p. 12-13).

Questions about the governance of research are urgent: a prominent field experiment is due to take place within the next year or two and the pace may pick up thereafter.3 There are signs that geoengineering is catching the attention of legislators and climate diplomats: the US Senate appropriations committee recently passed a spending bill that included funding to support geoengineering research (Nuccitelli 2016); there have been recent US house subcommittee hearings on geoengineering; and the head of the Intergovernmental Panel on Climate Change (IPCC), Hoesung Lee, has said that Panel should be examining geoengineering - including its governance - very seriously (Goldberg 2016). These developments sit alongside reports of concerns expressed by major figures in the SRM research community that research is 'under provided', and that governance is needed to make good on this.4

Here, I argue that in order to take seriously our moral obligations to future people, the governance of research into SRM must, at the very least, contain provisions to enable the shackling, or even the shutdown, of research programmes under conditions in which it is clear that such research cannot stand as a plausible precaution against climate change. These are conditions in which the research itself creates dangers requiring precautionary action. I argue that if any governance proposal for SRM research does not include provisions

3The experiment will be run through Harvard's Solar Geoengineering Research Program. See https://geoengineering.environment.harvard.edu/
4 See Geoengineering Research Governance Project (GRGP) Oxford Workshop on a Code of Conduct for Responsible Geoengineering Research – Anna-Maria Hubert | FCEA 2017)

and mechanisms that enable this action, it should be rejected by those considering it because it recklessly endangers the interests of future people. These requirements are far more demanding than any existing code for self-regulation by the SRM research community, and they imply that there should be no a priori allowed zone for research, including early stage indoor research involving modelling and lab experiments.[5]

In what follows, I shall assume that aggressive mitigation to achieve the Paris Agreement's 1.5C or 2C targets in the next 50 years is very unlikely. This explains why a call for accelerated research into SRM may gain traction outside of the SRM research community. Even if all parties achieve their Nationally Determined Contributions (NDCs) we may still see temperature rises of at least 2.6C by 2100 (Rogelj *et al.* 2016), and many climate experts now accept that there will be significant overshoot of temperature targets in the decades prior to 2100 (Rogelj *et al.* 2015, CAT 2016, Schleussner *et al.* 2016).[6] Knowing this helps to explain the potential appeal of SRM, which could (its advocates claim) cool the planet quickly and relatively cheaply, thus enhancing our chances of meeting the Paris Agreement targets by buying more time for mitigation (Keith 2013, MacMartin *et al.* 2014).[7] I shall also take it for granted in what follows that we have moral obligations to future people, including those not yet born (Gosseries and Meyer 2009, McKinnon 2012). I shall assume that whatever theory we have of these obligations, they prohibit the reckless endangerment of the fundamental interests future people have in secure access to water, food, shelter, as well as the protection of their physical integrity and social needs.[8]

[5] Thus, I address at least two of the 'key unresolved questions' of SRM governance identified by David Morrow (Morrow 2017), related to the objectives of climate engineering governance, and the role of precaution in that governance.
[6] Factoring in the ratcheting up of emissions reductions as per the Paris Agreement could affect this range, but it would not be unduly pessimistic (given the record of history) to be sceptical about the prospects for effective ratcheting up.
[7] For a skeptical take on these claims see (Barrett *et al.* 2014).
[8] For many, Henry Shue's conception of 'basic rights' captures this basket of social, political and environmental preconditions for human life with dignity (Shue 1996).

My arguments here do not support a moratorium on research. Instead, I see them as pushing us towards being much more honest about how much our present institutions would have to change in order for them to provide a basis for the ethically adequate governance of SRM research: simply building out from the institutions we have at present is unlikely to enable the right bar to be met (Gardiner 2014). The right attitude to this fact is not despair. Rather, we should see the present vacuum of governance for SRM research as an unusual opportunity to be bold in building new institutions, structures and processes fit to provide robust and ethically adequate protection to future people.[9]

The problem of governing SRM research

The governance of SRM research can be more or less 'shackling' or 'stimulating'. Governance that is very stimulating is fit to enable SRM research programmes to progress at speed to large-scale field tests, with a view to the development of deployable technology. Such governance could: facilitate a global coordination of research councils to make SRM research a priority; place no (or very minimal) requirements on researchers to engage in public consultation about proposed field experiments; authorise, support and direct SRM research in a narrow but deep 'club' agreement; or authorise unilateral research funded by private people, or corporations. In contrast, very shackling governance frameworks are those fit to significantly slow down research programmes, perhaps to stasis. Such governance could: require the consent of all affected by large-scale experiments; prohibit continuation of a research programme if all the risks it creates cannot be fully specified in advance; prohibit research unless all countries ratify a demanding global convention for the governance of research activities; or, place stringent limits on the financial support that can be given to SRM research in any national context.

---

[9] I do not think that wide-eyed optimism about the prospects for this opportunity being grasped is supported by the history of climate politics, nor its current state. Nevertheless, hope for this outcome can be consistent with, and justified at the same time, as pessimism about it (McKinnon 2014).

How ought we evaluate governance proposals along the shackling-stimulating spectrum in ways that capture what we owe to future people? One problem in thinking about this question is that there is deep uncertainty about the consequences for future people of any type of governance we adopt. Given this, the context in which choices about SRM governance are being made is sometimes presented as requiring a 'risk-risk trade-off' (Reynolds and Fleurke 2013): there are serious risks created by doing the research and deploying SRM, but also serious risks created by not doing the research and not deploying, given the realities of climate change. This frame suggests that the evaluation of governance proposals must focus on weighing the risks created by different proposals. The animating thought behind this frame is that once we decide on how to rank these risks, we will have a rudder to guide our choice of governance proposals.

Using this frame, the risks to be traded off can be thought of in the following way. On the one hand, if we govern in shackling ways that prevent or slow SRM research such that a deployable technology is not developed in the next century or so, warming of 2.6C or higher could be on the cards and we are deeply uncertain about the long-term consequences of that. At worst, the effects warming at 2.6C or above by 2100 could be catastrophic (Hauer *et al.* 2016, Mora *et al.* 2017): one nightmare scenario is speedy runaway warming, perhaps as a result of climate sensitivity having a much higher value than best estimates to date (Friedrich *et al.* 2016, Tan *et al.* 2016). In the face of this danger, stimulating governance that would kickstart a vigorous SRM research programme as soon as possible is sometimes presented as the best way to 'arm the future'. By stimulating SRM research through governance now, the argument goes, we can bequeath to future people a technology that: could enable them to buy additional time for an aggressive mitigation effort; limits climate impacts on the world's poorest people (Horton and Keith 2016, Keith and Irvine 2016); and thereby minimises future adaptation needs. By arming the future in this way, the argument goes, we can benefit both future people who might otherwise have had these adaptation needs, and also future people who would have to bear the cost of providing for these needs (Gardiner 2010).

The risks on the other side of the trade-off relate to the possible long-term consequences of deployment. If we govern to enable and stimulate SRM research so that it aims at delivering a deployable technology as soon as possible, we could lock in a pathway that commits us to deployment, and we are deeply uncertain about the long-term consequences of that (Robock 2008). We do not know whether the long-term impacts of deployment of SRM would be beneficial or catastrophic. Furthermore, we do not know whether governance that stimulates research would create 'moral hazard', nor how severe and long-lived this effect would be (Hale 2012, Baatz 2016). That is, we do not know whether SRM research would dilute commitment to emissions reductions (such as it is) and lead to a reduction in the rate of decrease, or instead (as some in the SRM research community hope) would galvanise political will on emissions reductions and buy much needed time for humanity to aggressively mitigate. Indeed, one of the calls for more research on SRM is often justified in terms of the need to reduce these uncertainties, so as to be better placed to weigh the risks of SRM against the risks of climate change in a policy-relevant way (Reynolds and Fleurke 2013).

How not to evaluate governance of SRM research

Faced with deep uncertainty in all directions about the consequences of governance that would stimulate or shackle research, can an ethical concern for the interests of future people provide guidance with respect to the evaluation of proposals for the governance of SRM research in the present? One way forward is to try to eradicate all uncertainty – or as much as we can – in order to employ techniques of cost-benefit analysis to evaluate risks created by different governance pathways so that they can be ranked. Whatever apparent progress can be made with cost-benefit analysis, it is premised – at the very least – on assumptions about the values of variables that generate the social discount rate that are unjustified given a proper ethical concern for the interests of future people (Gardiner 2011, Caney 2014, Jamieson 2014). Putting that aside, cost-

benefit analysis can enable full evaluation of competing stimulating and shackling governance proposals only if the techniques can be applied across the board to the range of dangers created by each pathway. If this proves possible then there is not deep uncertainty associated with the outcomes of all types of governance proposals, and we can choose paths that minimise expected damage to future people. Those who think that we know (or can come to know in a time frame enabling evaluation of governance frameworks needed now) the exact probabilities of all the outcomes created by doing, and refraining from, SRM research, can stop reading now (Goes *et al.* 2011, Bellamy *et al.* 2012).

For those who take uncertainty and our obligations to future people seriously (Tedsen and Homann 2013), a different approach is to assess governance proposals according to how well they meet standards of precaution in the face of future climate impacts (Elliott 2010, Morrow 2014) . There are at least two scenarios in which to make this case. First, SRM research could stand as a precaution against these possible impacts as a result of parties' failures to fulfil their pledges under the Paris Agreement. Second, the research could stand as a precaution against these impacts despite parties' successes in fulfilling their pledges. In the second case, the future climate change to which we are already committed, combined with the inadequacy of the present NDCs with respect to maximising the likelihood of exceeding 2C (let alone 1.5C) by 2100, could create a future world of climate catastrophe. In the first case, the parties' failure to achieve their inadequate NDCs could make things even worse.

In the face of these possible futures, the precautionary argument for governance to stimulate research is sometimes extended as follows. Without a well-researched, tested, safe, deployable SRM technology we are left open to the possibility that an unsafe, under-researched, untested SRM technology will be deployed in a climate emergency. To protect ourselves against this outcome, we ought to do the research now (Parson *et al.* 2013). This argument is too sanguine about what could make any deployment unsafe. It is far from fanciful to worry that a ruthless political regime could see deployment as an opportunity to perpetuate existing unjust distributions of power and resources. Given this

possibility, the assumption that a deployable technology delivered by a well-governed research programme starting now would ipso facto result in – or even make more likely – safe deployment in the future looks less well supported. A well-researched and tested technology could be more dangerous than an untested technology in the hands of a ruthless regime precisely because it would be more likely to work. Any argument for governance to stimulate research should not take a hidden punt on future political contexts being free from regimes that would seek to use the technology for their own unjust and/or illegitimate purposes.

The argument for precautionary action to stimulate SRM research in the face of future climate impacts requires not only that these impacts are possible, but further that they are above some threshold of probability beyond which the costs involved in taking the action are justified.[10] However, if we do not know the exact probabilities of these impacts, how can we know whether they are above any threshold of probability? The right approach here is to make a distinction between exactitude in our judgements of the probability of uncertain outcomes and accuracy in our judgements about whether these outcomes are above some threshold of probability: judgements of the latter type can be well founded when judgements of the former type are not. Shue (2010, p. 147; 2015) defends this method: 'one can reasonably, and indeed ought to, ignore entirely questions of probability beyond a certain minimal level of likelihood'. In these cases, Shue thinks, we are required to take precautionary action to minimise the dangers. These cases have the following features:[11]

[10] For a more general statement of the precautionary approach to climate change see Gardiner (2006), McKinnon (2009). For general criticism of precautionary approaches see (Sunstein 2005)(Sunstein 2007)(Sunstein 2007). It worth noting that even Sunstein accepts that a precautionary principle is justified in the face of possibly catastrophic climate change (Sunstein 2007). Excellent recent treatments of the precautionary principle are (Steel 2014, Hartzell-Nichols 2017).
[11] Shue identifies an additional third feature not listed here, as follows: '(3) *non-excessive costs*: the costs of prevention are not excessive (a) in light of the magnitude of the possible losses and (b) even considering other important demands on our resources'. (2010, p. 148). I shall take it to be straightforward that prevention costs are non-excessive in both the case of warming above 2C

> *massive loss*: the magnitude of the possible losses is massive;
>
> *threshold likelihood*: the likelihood of the losses is significant, even if no precise probability can be specified, because the mechanism by which the losses would occur is well understood, and the conditions for the functioning of the mechanism are accumulating (Shue 2010, p. 148).

Massive loss and *threshold likelihood* are clearly present in the cases of climate impacts consequent on successes or failures of the Paris Agreement. With respect to massive loss, climate change is already causing enormous damage to the most vulnerable people in the world. Regardless of whether parties meet their NDCs, committed climate change could become runaway climate change with truly catastrophic outcomes for future people. With respect to threshold likelihood (Shue's second condition), we know how and why climate change happens, we know that our continued greenhouse gas (GHG) emissions and massive land use changes are exacerbating climate change, we know with very high levels of confidence how climate change causes devastating impacts, and we know that these impacts are getting worse. Our knowledge of the mechanisms by which warming happens, and the mechanisms by which it impacts on human beings, is not complete; but what we do know is 'well understood'.

Given that climate impacts lie above a threshold for precautionary action, is governance that stimulates SRM research part of a package of precautionary measures required to protect the fundamental interests of future people? This case is not at all obvious. Commenting on this type of argument as it applies to

and the case of governance stimulating SRM research. In the climate case, the costs are those created by mitigation. In the SRM case, the costs are those created by not doing the research. Given the early-days state of research into SRM, these costs are almost entirely opportunity costs. Some SRM researchers make the case that this research could enable us to learn important things that would also benefit mitigation efforts, e.g. with respect to clouds. Granting this, if massive loss and threshold likelihood are satisfied in the case of governance stimulating SRM research, these opportunity costs are far from excessive, especially given that we may be able to learn these important things about clouds by other means.

deployment, Lauren Hartzell-Nichols reflects that, 'SRM strategies that pose new or further threats of catastrophe cannot be considered as an appropriate precautionary measure against climate change since there are other precautionary measures – mitigation and adaptation measures – that could be taken to greatly reduce, if not eliminate, the threats of catastrophe climate change threatens' (Hartzell-Nichols 2012, p. 165). Hartzell-Nichols makes her argument only with respect to the deployment of SRM as a precaution against the possible worst effects of committed climate change and does not extend the reasoning to SRM research and its governance. However, I think the argument can and should be extended in this way. If governance that stimulates enthusiastic SRM research is itself a token of the type of case Shue identifies as liable to precautionary action, then this governance and the research it stimulates cannot itself stand as a precaution against the future impacts of climate change.

A better way to evaluate governance of SRM research

I shall examine some conditions under which governance to stimulate SRM research meets Shue's *massive loss* and *threshold likelihood* conditions.[12] The point of identifying these conditions is not to assert emphatically that governance to stimulate SRM research must, or does, satisfy them: this is an empirical question. Rather, the point of identifying the conditions is to enable specification of some demanding requirements on stimulating governance of research if it is to be justified as a precautionary action against future climate impacts.

With respect to *massive loss*, we must start by recognising that there are many things that could cause the abrupt suspension of SRM deployment: for example, war, sabotage, or natural disaster. If deployment were to be abruptly suspended,

[12] This discussion is not exhaustive. For example, if concerns about moral hazard turn out to be empirically well grounded, shutdown provisions in research governance should be activated because a failure to curb emissions as a result of moral hazard would make climate impacts far worse than they would have been otherwise.

global average temperatures would be likely to rise very quickly to pre-deployment levels. Some studies suggest a similar, although less pronounced, effect on precipitation (See Figure 1. See also Jones *et al.* 2013, Keller *et al.* 2014, Zhang *et al.* 2015). This 'termination shock' would be damaging to future people both as a result of the impacts of the temperature rises (and other changes) previously masked by deployment, but also because of the speed at which these impacts would occur (Goes *et al.* 2011, Svoboda *et al.* 2011, Reynolds *et al.* 2016). In a termination shock, climate impacts that would have taken decades or longer to materialise as a result of cumulative emissions would be felt within a matter of years, causing great damage to people affected by the shock, especially the most vulnerable (Baum *et al.* 2013, MacMartin *et al.* 2014, Keith and MacMartin 2015). The severity of the termination shock will depend on the length of the deployment, the degree of mitigation that has happened under deployment, and the resilience of natural and social systems that have developed under deployment. Multi-century deployments with little mitigation could mean that natural and human systems have next to no chance of adapting to these changes in the short term, with devastating consequences for medium to long term mitigation and adaptation action. Even advocates of research into temporary and limited deployment used only to constrain the rate of change of global mean temperature, rather than to stabilize temperature at some acceptable level (MacMartin *et al.* 2014, Keith and MacMartin 2015), accept that abrupt termination would cause damage (Cairns 2014).[13]

The possibility of massive losses caused by a termination shock is relevant to thinking about the governance of research if there are well-understood mechanisms that could lock-in a path from research to deployment. If such mechanisms exist, the first part of what is needed for *threshold likelihood* is

[13] It should also be noted in this context that the range of scenarios in which SRM deployment is imagined are all framed in ways that have significant consequences for assessment of the acceptability of deployment in those scenarios. In their review of the literature assessing geo-engineering proposals, Bellamy et al. provide evidence that a number of frames that could support far less favourable assessments of deployment are routinely excluded from scenarios (Bellamy *et al.* 2012, p. 605, Fragniere and Gardiner 2016).

present in the SRM research case, namely a well-understood mechanism by which massive losses can be brought about. Broadly speaking, a society has become locked in to a particular technology, system, form of organisation, or framing of an issue when the reasons for not abandoning the locked-in feature are a function of the high costs of doing this, rather than the benefits and advantages of the locked-in feature itself (Jamieson 1996). Features of a society that are locked-in can persist and dominate even when there are many better ways to pursue the ends, or arrange the affairs, in question.[14] Lock-in of SRM research programmes, and lock-in from any programme to deployment, could happen as a result of positive feedbacks in four categories.[15]

First, economies of scale: any particular SRM research programme is likely to be expensive to get started; and once started, the continued funding of (and any progress made by) that programme makes it harder for alternative programmes to launch, even when they have clear potential to be better than the already existing programme. At a certain point in any programme, the only way to field test the technology is to actually deploy on a global scale, even if outdoor experiments start on a smaller scale (Hamilton 2013, Klein 2014). And once research-through-deployment is happening the space for that type of field experiment for any other SRM research programme is – quite literally – already filled with the existing technology.

Second, learning effects: as any research programme advances, perhaps to the point of research-through-deployment, the opportunities for learning in ways that enable improvement and fine tuning of the technology accumulate. Alternative programmes cannot match this, which progressively weakens the case for their initiation and funding.

---

[14] The meaning of 'better' in this context depends on the nature of the locked in feature. For example, 'better' for technologies could mean more effective and cheaper; for frames for social issues, it could mean 'fairer'.
[15] These four sites of positive feedback causing lock-in of SRM are adapted from Foxon (2007).

Third, adaptive expectations: it need not take long for a 'new normal' to take root. Human beings have a well-documented status quo bias which could operate quite quickly to adjust their preferences to any SRM research programme actually in existence, even if there are potential alternatives that are better (Kahneman *et al.* 1991). Furthermore, 'environmental generational amnesia' (Kahn 2011) – whereby the baseline for judging environmental degradation becomes less demanding over generations as degradation actually happens – could make the expectations of children born within a period of research-through-deployment radically different from those of their parents by the time they grow up. Perhaps only nostalgic 'old fools' with 'lighted rooms' inside their heads would mourn for lost blue skies and starry nights (Larkin 1974).

Fourth, co-ordination effects: the more research programmes in other countries are initiated and developed, the more it is to any country's advantage to accelerate their own research programme. Given that any research programme must eventually field test through actual global deployment, the existence of any research programme in a national content makes deployment much more likely overall. The obvious analogue here is the nuclear arms race.

These four sites of positive feedback, and various combinations of them, can operate both to lock-in a particular token SRM research programme (when a less dangerous token is available) and/or to lock-in to deployment as part of the research process for any SRM technology (in other words, to lock-in to deployment from SRM as a type of research programme). The nature of the investment and infrastructure needed to scale up innovations in technologically well-developed societies probably makes some lock-in – of tokens and of types – inevitable. But some locked-in technologies are more dangerous than others. Indeed, we are in a climate crisis in large part because our predecessors locked-in to a carbon-intensive global economic structure. If we are able to identify the potential for great danger in technological research programmes (as tokens or types) that might be in the process of locking-in, we should act in ways that thwart the locking-in and so avoid unnecessary endangerment of future people.

We know that in the process of socio-technical lock-in, earlier events in the sequence matter more than later ones. Earlier events capture technical, cognitive, political and psychological ground in ways that can crowd out features potentially in competition with what is getting locked-in, with the result that outcomes are certain to be not wholly - and are sometimes not at all – as good as they could be for the people, organisations, and societies living with the outcomes. The well-understood temporality of the process of lock-in makes it particularly germane to where we are now with SRM research. Given that we are right at the beginning of the possible development of SRM technologies, the governance choices we make now could make a disproportionately large contribution to the locked-in status of the technology in the future, along with the governance, frames, and expectations that will accompany it (Jamieson 1996).

To address this, governance proposals must contain strenuous flexibility mechanisms by which research programmes can be shackled, or even quickly shutdown, at minimal cost if the signs and signals of lock-in start to show. The massive losses that could follow for future people without these strenuous mechanisms justify their inclusion.[16] The details of what shackling or shutdown means will vary according to the case. Shackling or shutdown could be achieved by withholding funding, or refusal of ethical approval for continuing research activities in the institutional setting of the research. If the research is much further progressed, treaties banning large-scale field experiments might be appropriate. There are obvious parallels here with various nuclear test ban treaties (Morrow, Kopp and Oppenheimer 2013).

This leaves the second part of threshold likelihood, namely that the conditions for the functioning of the mechanism (through which massive losses could occur) are accumulating. In the case of governance stimulating SRM research, a

[16] Some existing governance reports have highlighted flexibility as a desirable feature of research governance regimes, but they defend this in the name of enabling research that changes over time with changing circumstances, whereas my case is for flexibility as a way of transforming stimulating governance into shackling governance (Morrow 2017).

significant part of this mechanism is lock-in.[17] Assuming that governance proposals to stimulate research include strenuous shackling and shutdown provisions, is it the case that the conditions for the functioning of lock-in mechanisms are accumulating? Without knowing this, we do not know whether SRM research, and governance that aims to stimulate it, is a live danger calling for precaution itself (and thus without standing as a precaution against future climate impacts), or a latent danger that could be used as a precaution against future climate impacts because conditions in the world are not right for it to become dangerous. Establishing the answer to this question for SRM research and its governance – indeed, for any case we assess using Shue's criteria – is an entirely empirical question. While the question is very far from settled, there is already some evidence that the framing of SRM could be generating cognitive lock-in to a more limited range of options than is justified once other frames are admitted into assessments of different scenarios for research and eventual deployment (Bellamy *et al.* 2012).

Identifying the significance of knowing whether conditions are accumulating in a way that is likely to generate premature lock-in to particular SRM research programmes, and from these programmes to deployment, enables us to identify a further aspect of the bar of justification that governance proposals for SRM research must meet. Governance must include powerful oversight provisions for the continuous assessment of whether flexibility mechanisms must be activated to shackle or shutdown the research programmes stimulated by the governance frameworks laid out in the proposals.

There is a variety of ways in which such oversight could be achieved. For example, Lin acknowledges the need for governance of SRM research that takes account of far more than the physical risks created by any research programme, and identifies lock-in as a 'systemic concerns' that must be addressed by an

[17] There may be additional mechanisms we should consider. In particular, the form of moral hazard involving mitigation deferral is worth considering because of how the operation of this mechanism could make the massive losses of climate change, and/or a termination effect, so much worse at some point in the future (Baatz 2016).

adequate governance framework for research (Lin 2015). A further example is Dilling and Hauser's suggestion that project proposals be legally required to state how the research will address concerns that go beyond physical risks to include 'decision-making power/responsibility for outcomes, and the societal meaning of the research' (Stirling 2008, Stirling *et al.* 2008, Corner *et al.* 2012a, Preston 2012). What unites specific proposals in this area, regardless of their differences of detail, is a commitment to broadening out assessment of SRM research programmes through public participation and inclusion of more marginalised perspectives (MacMartin *et al.* 2014).

While such proposals are a good start with respect to oversight, they rarely contain explicit and formal provisions for the protection of future people. Very often they are entirely bound to existing political institutions, so it is no surprise that they do not transcend the limits of these institutions with respect to the creation of robust governance requirements fit to offer future people genuine protection against the termination effect as a result of lock-in to a dangerous SRM technology that we, their predecessors, failed to avert. Whatever specific oversight provisions are adopted, they ought to be extended in this way, so as to provide an alert that existing governance that is stimulating SRM research should transform quickly into a shackle, perhaps to the point of shutting down the relevant research programmes altogether, if lock-in appears to be serious. The fact that none of our present global and transnational governance institutions and regimes adequately protect the interests of future people is no excuse for not seeking such governance in the SRM case. On the contrary, the governance of SRM offers an unusual and promising context of a governance vacuum in which new institutions and regimes could be created with ethically required components offering protection to future people that other governance structures could emulate. Creating new forms of governance for SRM research could be an opportunity for us to do better than our predecessors in creating institutions enabling us to meet our intergenerational obligations.

Conclusion

I am not optimistic that shackling and shutdown mechanisms - along with ethically adequate oversight provisions - will be included in new governance frameworks for SRM research. It is more likely that new governance for SRM research will be grafted onto already existing institutions, hardly any of which offer ethically adequate protection to future people.18 Suppose that this is how things unfold over the coming decades and that future people in, say, 2300 nevertheless find themselves in a best-case scenario world, as follows. In 2020, governance of SRM research stimulated research that progressed at speed and locked-in deployment, without oversight or management of the lock-in. SRM was deployed in 2060 for a relatively short time to slow the rate of change of global mean temperature, which enabled aggressive mitigation to take place. SRM deployment was gradually phased out from 2160, in a well-managed and responsible process(McKinnon 2017)(McKinnon 2017)(McKinnon 2017). In 2300 global mean temperature has stabilised in accordance with the goals of the Paris Agreement.

In this scenario, future people get lucky and are not damaged by our present failure to incorporate shackling and shutdown mechanisms, and oversight provisions, in SRM research governance. Even in this scenario, these future people have at least one well-grounded and serious moral complaint against us. This is that we knowingly imposed on them an impermissible risk of termination shock. We did this by failing to govern research in a way that would have prevented or minimised this via a provision to oversee, and potentially to shackle or shutdown parts of the research fit to create the danger. We cannot read back from the fact that 'in the end, it turned out for the best' to the conclusion that we acted permissibly in governing SRM at an earlier point in time. Our judgements of when and why prosecution of drunk drivers is justified shows the deep moral roots of this claim. We do not think that the fact that the driver hurt no-one lets them off the hook, morally speaking. Instead, we think that their dangerous behaviour makes them culpable (and makes prosecution justified) regardless of the fact that no-one was hurt. Indeed, we think this even if

18 Things could be otherwise, without the world becoming a utopia (González-Ricoy and Gosseries 2016).

the drunk driving not only caused no harm but actually did some good (for example, by enabling the drunk driver to attend the birth of her child, and so foster a lasting parental bond).

In response, an SRM researcher might say: SRM research can be a precaution against climate impacts if there is no better precaution. If mitigation efforts are unlikely to work, the researcher might say, then there is in fact no better precaution than SRM research. But this reply is too coarse-grained. If it is true that we ought to take seriously SRM research as a precaution against climate impacts, there are still questions about how that research should be governed, and we can (at least in principle) choose between differently governed SRM research programmes in considering the case for SRM research as a precaution. My argument here is that SRM research governed with shackling or shutdown provisions, and the right kind of oversight, is a better precaution against climate impacts (if it is any precaution at all) than SRM research governed without these provisions. So taking seriously the principle that SRM research can be a precaution against climate impacts, if there is no better precaution, we should adopt my suggestions.

Here, I have identified two desiderata of governance of SRM research that are generated by taking seriously our obligations not to wrong future people. My aim here has been to provide part of a toolkit for assessing the quality of our present choices with respect to governance for SRM research, despite deep uncertainty about the impact of climate change, and about the effects of SRM deployment, on future people. We have some knowledge about how and why technologies can lock-in, and that in some cases this can have seriously damaging impacts on those living with locked in technology. And we know that SRM research is at a stage at which the mechanisms of lock-in could start to operate. In this early stage of SRM research, lock-in could be slowed or stopped with the right kind of governance. If we do not achieve this governance we will wrong future people by recklessly endangering them, even if in fact these dangers end up not ripening into harms.

Acknowledgements

References

Baatz, C., 2016. Can we have it both ways? On potential trade-offs between mitigation and solar radiation management. *Environmental Values*, 25 (1), 29–49.

Barrett, S., Lenton, T.M., Millner, A., Tavoni, A., Carpenter, S., Anderies, J.M., Chapin, F.S., Crépin, A.S., Daily, G., Ehrlich, P., Folke, C., Galaz, V., Hughes, T., Kautsky, N., Lambin, E.F., Naylor, R., Nyborg, K., Polasky, S., Scheffer, M., Wilen, J., Xepapadeas, A., and De Zeeuw, A., 2014. Climate engineering reconsidered. *Nature Climate Change*, 4 (7), 527–529.

Baum, S.D., Maher, T.M., and Haqq-Mistra, J., 2013. Double Catastrophe: Intermittent Stratoshperic Geoengineering Induced by Societal Collapse. *Environment, Systems and Decisions*, 33 (1), 168–80.

Bellamy, R., Chilvers, J., Vaughan, N.E., and Lenton, T.M., 2012. A review of climate geoengineering appraisals. *Wiley Interdisciplinary Reviews: Climate Change*, 3 (6), 597–615.

Cairns, R.C., 2014. Climate geoengineering: Issues of path-dependence and socio-technical lock-in. *Wiley Interdisciplinary Reviews: Climate Change*, 5 (5), 649–661.

Caney, S., 2014. Climate change, intergenerational equity and the social discount rate. *Politics, Philosophy & Economics*, 13 (4), 320–342.

CAT, 2016. Global Temperatures - Climate Action Tracker [online].

Cicerone, R.J., 2006. Geoengineering: Encouraging research and overseeing implementation. *Climatic Change*, 77 (3–4), 221–226.

Corner, A., Pidgeon, N., and Parkhill, K., 2012a. Perceptions of geoengineering: public attitudes, stakeholder perspectives, and the challenge of ?upstream? engagement. *Wiley Interdisciplinary Reviews: Climate Change*, 3 (5), 451–466.

Corner, A., Pidgeon, N., and Parkhill, K., 2012b. Perceptions of geoengineering: public attitudes, stakeholder perspectives, and the challenge of 'upstream' engagement. *Wiley Interdisciplinary Reviews: Climate Change*, 3 (5), 451–466.

Crutzen, P.J., 2006. Albedo enhancement by stratospheric sulfur injections: A contribution to resolve a policy dilemma? *Climatic Change*, 77 (3–4), 211–219.

Elliott, K., 2010. Geoengineering and the Precautionary Principle. *International Journal of Applied Philosophy*, 24 (2), 237–253.

Foxon, T.J., 2007. Technological lock-in and the role of innovation. *In*: G. Atkinson, S. Dietz, and E. Neumayer, eds. *Handbook of Sustainable Development: Second Edition*. Cheltenham: Edward Elgar, 140–154.

Fragniere, A. and Gardiner, S.M., 2016. Why Geoengineering is Not 'Plan B'. *Climate Justice and Geoengineering: Ethics and Policy in the Atmospheric Anthropocene*, (October), 15–32.

Friedrich, T., Timmermann, A., Tigchelaar, M., Elison Timm, O., and Ganopolski, A., 2016. Nonlinear climate sensitivity and its implications for future greenhouse warming. *Science Advances*, 2 (11), e1501923–e1501923.

Gardiner, S.M., 2006. A Core Precautionary Principle*. *Journal of Political Philosophy*, 14 (1), 33–60.

Gardiner, S.M., 2010. 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, and H. Shue, eds. *Climate Ethics: Essential Readings*. Oxford: Oxford University Press, 284–314.

Gardiner, S.M., 2011. *A perfect moral storm : the ethical tragedy of climate change*. Oxford University Press.

Gardiner, S.M., 2014. A Call for a Global Constitutional Convention Focused on

Future Generations. *Ethics & International Affairs*, 28 (3), 299–315.

Geoengineering Research Governance Project (GRGP) Oxford Workshop on a Code of Conduct for Responsible Geoengineering Research – Anna-Maria Hubert | FCEA [online], n.d.

Goes, M., Tuana, N., and Keller, K., 2011. The economics (or lack thereof) of aerosol geoengineering. *Climatic Change*, 109 (3–4), 719–744.

Goldberg, S., 2016. UN climate science chief: it's not too late to avoid dangerous temperature rise | Environment | The Guardian [online]. *The Guardian*.

González-Ricoy, I. and Gosseries, A., eds., 2016. *Institutions For Future Generations*. Oxford University Press.

Gosseries, A. and Meyer, L.H., 2009. *Intergenerational justice*. Oxford University Press.

Hale, B., 2012. Geoengineering and Moral Hazards. *Engineering the Climate: The Ethics of Solar Radiation Management*.

Hamilton, C., 2013. *Earthmasters: The Dawn of the Age of Climate Engineering*. Yale University Press.

Hartzell-Nichols, L., 2012. Precaution and Solar Radiation Management. *Ethics, Policy and Environment*, 15 (2), 158–171.

Hartzell-Nichols, L., 2017. *A Climate of Risk Precautionary Principles, Catastrophes, and Climate Change.* Taylor and Francis.

Harvard's Solar Geoengineering Research Program [online], n.d.

Hauer, M.E., Evans, J.M., and Mishra, D.R., 2016. Millions projected to be at risk from sea-level rise in the continental United?States. *Nature Climate Change*, 6 (7), 691–695.

Heyward, C., 2015. Is There Anything New Under the Sun? *In*: A. Maltais and C. Mckinnon, eds. *The Ethics of Climate Governance*. London: Rowman and Littlefield International, 133–154.

Horton, J. and Keith, D., 2016. Solar Geoengineering and Obligations to the Global Poor. *In*: *Climate Justice and Geoengineering: Ethics and Policy in the Atmospheric Anthropocene*. Rowman and Littlefield International, 79–92.

Hubert, A.-M., Kruger, T., and Rayner, S., 2016. Geoengineering: Code of conduct for geoengineering. *Nature*, 537 (7621), 488–488.

Jamieson, D., 1996. Ethics and intentional climate change. *Climatic Change*, 33

(3), 323–336.

Jamieson, D., 2014. *Reason in a Dark Time*. Oxford University Press.

Jones, A., Haywood, J.M., Alterskjaer, K., Boucher, O., Cole, J.N.S., Curry, C.L., Irvine, P.J., Ji, D., Kravitz, B., Egill Kristjánsson, J., Moore, J.C., Niemeier, U., Robock, A., Schmidt, H., Singh, B., Tilmes, S., Watanabe, S., and Yoon, J.-H., 2013. The impact of abrupt suspension of solar radiation management (termination effect) in experiment G2 of the Geoengineering Model Intercomparison Project (GeoMIP). *Journal of Geophysical Research: Atmospheres*, 118 (17), 9743–9752.

Kahn, P.H., 2011. *Technological nature : adaptation and the future of human life*. MIT Press.

Kahneman, D., Knetsch, J.L., and Thaler, R.H., 1991. Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *The Journal of Economic Perspectives*, 5, 193–206.

Keith, D.W., 2013. *A case for climate engineering*. Boston: MIT Press.

Keith, D.W. and Irvine, P.J., 2016. Solar geoengineering could substantially reduce climate risks-A research hypothesis for the next decade. *Earth's Future*, 4 (11), 549–559.

Keith, D.W. and MacMartin, D.G., 2015. A temporary, moderate and responsive scenario for solar geoengineering. *Nature Climate Change*, 5 (3), 201–206.

Keller, D.P., Feng, E.Y., and Oschlies, A., 2014. Potential climate engineering effectiveness and side effects during a high carbon dioxide-emission scenario.

Klein, N., 2014. *This Changes Everything*. Penguin.

Larkin, P., 1974. The Old Fools. *In*: *High Windows*. London: Faber and Faber, 13–14.

Lin, A.C., 2015. The missing pieces of geoengineering research governance. *Minnesota Law Review*, 152 (434), 2509–2576.

MacMartin, D.G., Caldeira, K., and Keith, D.W., 2014. Solar geoengineering to limit the rate of temperature change. *Philosophical Transactions of the Royal Society A*, 372, 20140134.

McKinnon, C., 2009. Runaway Climate Change: A Justice-Based Case for Precautions. *Journal of Social Philosophy*, 40 (2), 187–203.

McKinnon, C., 2012. *Climate change and future justice : precaution, compensation, and triage*. Routledge.

McKinnon, C., 2014. Climate change: Against despair. *Ethics and the Environment*, 19 (1), 31–48.

McKinnon, C., 2017. Endangering humanity: an international crime? *Canadian Journal of Philosophy*, 47 (2–3), 395–415.

Mora, C., Dousset, B., Caldwell, I.R., Powell, F.E., Geronimo, R.C., Bielecki, C., Counsell, C.W.W., Dietrich, B.S., Johnston, E.T., Louis, L. V., Lucas, M.P., McKenzie, M.M., Shea, A.G., Tseng, H., Giambelluca, T.W., Leon, L.R., Hawkins, E., and Trauernicht, C., 2017. Global risk of deadly heat. *Nature Climate Change*, 7 (7), 501–506.

Morrow, D., 2017. International Governance of Climate Engineering: A Survey of Reports on Climate Engineering, 2009-2015.

Morrow, D.R., 2014. Ethical aspects of the mitigation obstruction argument against climate engineering research Author for correspondence : *Philosophical Transactions of the Royal Society*, 372, 1–14.

Morrow, D.R., Kopp, R.E., and Oppenheimer, M., 2013. Political legitimacy in decisions about experiments in solar radiation management. *In*: W.C.G. Burns and A. Strauss, eds. *Climate Change Geoengineering: Philosophical Perspectives, Legal Issues, and Governance Frameworks*. Cambridge: Cambridge University Press, 146–167.

Nuccitelli, D., 2016. Scientists Debate Experimenting With Climate Hacking to Prevent Catastrophe. *The Guardian*, Jun.

Parson, E.A., Keith, D.W., Teller, E., Hyde, R., Wood, L., Levitt, S., Dubner, S., Robock, A., Hamilton, C., Keith, D., Keith, D.W., Parson, E., Morgan, M.G., Parson, E., Ernst, L., Lukacs, M., and Biello, D., 2013. Science and regulation. End the deadlock on governance of geoengineering research. *Science (New York, N.Y.)*.

Preston, C.J., 2012. Solar Radiation Management and Vulnerable Populations. *In*: C.J. Preston, ed. *Engineering the Climate: The Ethics of Solar Radiation Management*. Lanham: Lexington Books, 77–94.

Rayner, S., Heyward, C., Kruger, T., Pidgeon, N., Redgwell, C., and Savulescu, J., 2013. The Oxford Principles. *Climatic Change*, 121 (3), 499–512.

Reynolds, J. and Fleurke, F., 2013. Climate Engineering Research: A Precautionary Response to Climate Change. *Cclr*, (January 2011), 101–108.

Reynolds, J.L., Parker, A., and Irvine, P., 2016. Five solar geoengineering tropes that have outstayed their welcome. *Earth's Future*, 4 (12), 562–568.

Robock, A., 2008. 20 Reasons Why Geoengineering May Be a Bad Idea. *Bulletin of the Atomic Scientists*, 64 (2), 14–18.

Rogelj, J., Den Elzen, M., Höhne, N., Fransen, T., Fekete, H., Winkler, H., Schaeffer, R., Sha, F., Riahi, K., and Meinshausen, M., 2016. Paris Agreement climate proposals need a boost to keep warming well below 2 °c. *Nature*, 534 (7609), 631–639.

Rogelj, J., Luderer, G., Pietzcker, R.C., Kriegler, E., Schaeffer, M., Krey, V., and Riahi, K., 2015. Energy system transformations for limiting end-of-century warming to below 1.5 ?C. *Nature Climate Change*, 5 (6), 519–527.

Schleussner, C.-F., Rogelj, J., Schaeffer, M., Lissner, T., Licker, R., Fischer, E.M., Knutti, R., Levermann, A., Frieler, K., and Hare, W., 2016. Science and policy characteristics of the Paris Agreement temperature goal. *Nature Climate Change*, 6 (9), 827–835.

Shue, H., 2010. Deadly Delays, Saving Opportunities: Creating a More Dangerous World? *In*: S.M. Gardiner, S. Caney, D. Jamieson, and H. Shue, eds. *Climate Ethics: Essential Readings*. Oxford University Press, 146–162.

Shue, H., 2015. Uncertainty as the Reason for Action: Last Opportunity and Future Climate Disaster. *Global Justice: Theory Practice Rhetoric*, 8 (2), 86–103.

Steel, D., 2014. *Philosophy and the Precautionary Principle*. Cambridge: Cambridge University Press.

Stirling, A., 2008. Science, precaution, and the politics of technological risk: Converging implications in evolutionary and social scientific perspectives. *Annals of the New York Academy of Sciences*, 1128, 95–110.

Stirling, A., Smith, A., Leach, M., Pellizzoni, L., Levidow, L., Hendriks, C., Owens, S., von Tunzelmann, N., Wynne, B., and Centre, F., 2008. " Opening Up" and "Closing Down": Power, Participation, and Pluralism in the Social Appraisal of Technology. *Science, Technology, & Human Values*, 33 (2), 262–294.

Sunstein, C.R., 2005. *Laws of fear : beyond the precautionary principle*. Cambridge University Press.

Sunstein, C.R., 2007. The Catastrophic Harm Precautionary Principle. *Issues in Legal Scholarship*, 6 (3).

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

Tan, I., Storelvmo, T., and Zelinka, M.D., 2016. Observational constraints on mixed-phase clouds imply higher climate sensitivity. *Science*, 352 (6282).

Tedsen, E. and Homann, G., 2013. Implementing the Precautionary Principle for Climate Engineering. *Carbon and Climate Law Review*, 1 (2), 90–100.

The Royal Society, 2009. *Geoengineering the climate: science, governance and uncertainty*. Clean Technologies and Environmental Policy.

Wibeck, V., Hansson, A., Anshelm, J., Asayama, S., Dilling, L., Feetham, P.M., Hauser, R., Ishii, A., and Sugiyama, M., 2017. Making sense of climate engineering: a focus group study of lay publics in four countries. *Climatic Change*, 145 (1–2), 1–14.

Zhang, Z., Moore, J.C., Huisingh, D., and Zhao, Y., 2014. Review of Geoengineering Approaches to Mitigating Climate Change. *Journal of Cleaner Production Zhang Z Journal of Cleaner Production*.