nature climate change

PERSPECTIVE

PUBLISHED ONLINE: 16 FEBRUARY 2015 | DOI: 10.1038/NCLIMATE2493

A temporary, moderate and responsive scenario for solar geoengineering

David W. Keith[1]* and Douglas G. MacMartin[2]

The risks and benefits of solar geoengineering, or solar radiation management (SRM), depend on assumptions about its implementation. Claims that SRM will reduce precipitation, increase ocean acidification or deplete stratospheric ozone, or that it must be continued forever once started, are not inherent features of SRM; rather, they are features of common scenarios for its implementation. Most analyses assume, for example, that SRM would be used to stop the increase in global temperature or restore temperature to pre-industrial values. We argue that these are poor scenario choices on which to base policy-relevant judgements about SRM. As a basis for further analysis, we provide a scenario that is temporary in that its end point is zero SRM, is moderate in that it offsets only half of the growth in anthropogenic climate forcing and is responsive in that it recognizes that the amount of SRM will be adjusted in light of new information.

One cannot meaningfully evaluate solar geoengineering without a scenario for its implementation. It is now common, for example, to assert that more scientific research is needed to assess the balance between the risks and benefits of solar geoengineering, hereafter called solar radiation management (SRM). Yet the balance between risks and benefits depends at least as strongly on how SRM is deployed (for example on technology choice, timing and magnitude of the induced radiative forcing) as it depends on the climate's response to a specified SRM scenario.

Clear language is an essential tool for analysing this messy topic. We use SRM to denote a technology used to deliberately alter radiative forcing at sufficient scale to measurably alter the global climate. Any technology for producing radiative forcing will have a set of technology-specific impacts, such as ozone loss arising from the introduction of aerosol particles in the stratosphere. However the radiative forcing is produced, the efficacy of SRM is inherently limited by the fact that a change in solar radiative forcing cannot perfectly compensate for the radiative forcing caused by increasing greenhouse gases.

SRM has been variously framed as a substitute for cutting emissions (mitigation), as an emergency measure to be used if climate risks are higher than expected, or as a means to restoring surface temperatures to pre-industrial. Explicit or implicit, such scenarios shape any assessment of risk and efficacy of SRM.

Ocean acidification has been listed as a risk of SRM[1], yet acidification depends almost solely on cumulative CO_2 emissions and is unaffected by SRM. Ocean acidification is a risk of SRM only if SRM is used as a substitute for emissions mitigation; and in this case, the risk derives from the increase in emissions not from SRM.

Reduced precipitation is another frequently cited risk of SRM (see Supplementary Information for examples). It is true that if the SRM radiative forcing is large enough to offset all of the change in global mean temperature due to anthropogenic CO_2 — a common assumption — then precipitation will indeed be reduced in most locations[2]. Simple physical arguments demonstrate that it takes a smaller SRM forcing to stop the rise in precipitation as CO_2 concentrations increase than is required to stop the rise in temperature[3]. Reduction in precipitation is, however, a product of the magnitude of SRM used in the scenario. If the SRM radiative forcing was adjusted to maintain global-average precipitation rates at their pre-industrial level then temperatures would be above pre-industrial. The claim that geoengineering will reduce average precipitation thus turns on the assumption that more SRM will be used than is required to stop the increase in precipitation caused by rising CO_2 concentrations.

As these examples illustrate, judgements about whether the use of SRM can be justified are determined by policy assumptions about how it will be used at least as strongly as they are determined by scientific analysis.

We articulate a scenario in sufficient detail to allow quantitative analysis of its physical and social implications, but we do not attempt to describe a political scenario that might result in this physical scenario being implemented. We do not claim that this scenario is likely or optimal, only that it is less suboptimal than the scenarios used most commonly. We adopt the central planner framing common in economic models that underlie much climate policy analysis and assume that decisions about implementation of SRM are made to maximize some measure of global welfare[4]. In practice, the nexus of decisions about SRM will involve nation states which are influenced by many factors, not least public and private transnational organizations, each of which have complex internal politics. Moreover, decisions about SRM take place in an environment in which decision makers face multiple issues and make decisions under substantial uncertainty. In this environment, the worst-case outcomes might include gross misuses of SRM or even war[5].

Although we think it is unrealistic, we adopt the central planner framing for three reasons. First, because it is a common benchmark for climate policy analysis, it is a useful framework in which to compare SRM with other response options such as emissions mitigation and adaptation. Second, there is simply no tractable way to analyse the full decision problem, and our goal is not analysis but rather the construction of a scenario that is useful for further analysis including exploration of the political and institutional implications. Third, and finally, we hope that articulating an outcome that is closer to the social planner's optimum will aid the development of policy that might nudge the world towards a better outcome.

Our objective is to provide a scenario for implementation of SRM that is specific enough to be assessed and critiqued yet general enough to be used for a wide variety of science and policy

---

[1]Harvard University, One Brattle Square, Office 483, Cambridge, Massachusetts 02138, USA. [2]California Institute of Technology, 1200 E. California Blvd, M/C 305-16, Pasadena, California 91125, USA. *e-mail: david_keith@harvard.edu

© 2015 Macmillan Publishers Limited. All rights reserved

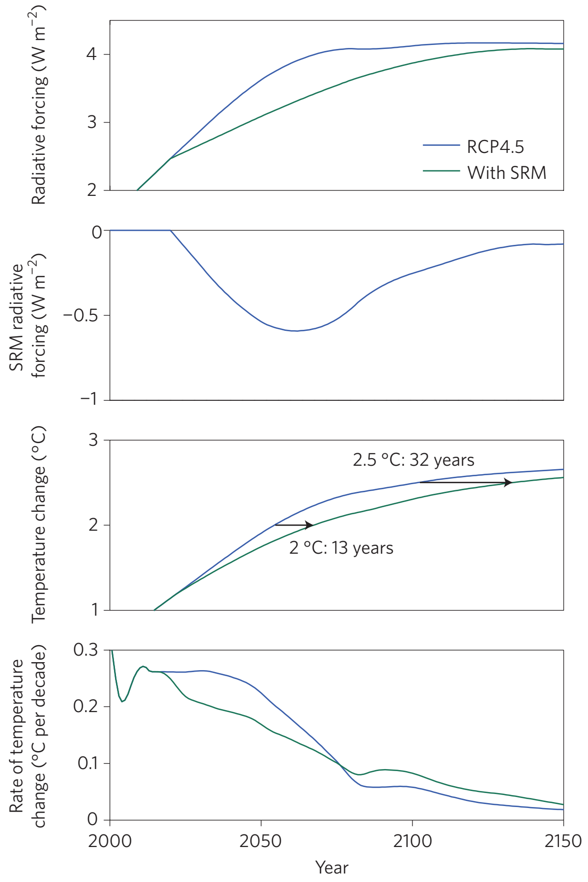

Figure 1 | Illustration of the SRM scenario for an RCP4.5 emissions profile. The top panel shows the total radiative forcing for RCP4.5, and a radiative forcing profile in which the rate of change is halved starting in 2020; that is, for year k, $\mathrm{RF}_{\mathrm{new}}(2020 + k) = \mathrm{RF}_{\mathrm{RCP4.5}}(2020 + k/2)$. The difference between these gives the suggested initial SRM profile in the second panel. The effect on global mean temperature as predicted by MAGICC (with a 3 °C climate sensitivity) is shown in the third panel, and the corresponding decadal rates of change in the final panel. Note that temperature and its rate of change would depend on climate sensitivity, but the amount of SRM would not. If rate-independent climate impacts increase superlinearly, then the benefits will be larger than is evident in the third panel. If impacts are quadratic in temperature, then impacts will be reduced by 20% in 2070 (roughly the time when SRM radiative forcing peaks) although the temperature increase ΔT is only reduced by 10%. (See ref. 16 for a climate damage function that depends on both the magnitude and rate of change of temperature.)

analysis. We define the scenario in the next section while deferring the considerations that motivate our choice of scenario to the section following that. Next we explore a specific choice of scenario including technological details as a worked example. The final section provides a concluding summary.

Scenario

Our scenario combines three elements: a specific method of altering solar forcing, an initial trajectory for SRM radiative forcing over time, and a plan for altering the trajectory based on new information. We aim to provide a scenario that is articulated in sufficient detail to allow quantitative evaluation of risk and efficacy.

Further, our scenario is chosen to meet the following criteria: (i) it is temporary in that the end point is zero SRM; (ii) it is moderate in that it does not offset all of the global mean temperature change due to increased greenhouse gases; and (iii) it is responsive in that it explicitly recognizes that the amount of SRM will be adjusted in light of new information. We elaborate the motivation behind each criterion in 'Guiding principles' below.

We link the amount of SRM to the amount of mitigation, in that slower growth in greenhouse gas forcing means a slower growth in SRM, but we do not make the converse linkage. We suggest that the risks and benefits of SRM be evaluated by comparing scenarios with and without SRM that use the same radiative forcing trajectory, although we recognize that the choice to use SRM may itself influence the amount of mitigation in one direction or the other. The scenario is defined as follows:

Radiative forcing trajectory. Beginning in 2020, adjust the global SRM radiative forcing so as to halve the rate of growth of net non-SRM anthropogenic radiative forcing. The top panel of Fig. 1 provides an example for a specific radiative forcing scenario.

Technology. Use stratospheric aerosol SRM with as even as possible a global distribution of radiative forcing. As a possible example (elaborated in the section 'A specific example'), one might begin using direct injection of SO_2 gas and transition to H_2SO_4 vapour (to improve aerosol size distribution) by 2030. One might begin efforts to develop — and where appropriate test — more advanced scatterers that offer lower ozone impact, lower overall health impact or less diffuse light scattering with the intention of transitioning to advanced particles by 2050.

Responsiveness. Adjust the amount of forcing relative to the initial trajectory defined above based on any evidence that the effects of using SRM differ from expectations in ways that affect the assessment of benefits or harms. Examples include evidence that the effect on depletion of ozone is significantly larger than expected, evidence that the regional climate response (temperature, precipitation and so on) to forcing differs from model-based predictions, or evidence of unexpected impacts of climate change such as larger than expected rates of Arctic methane release.

Monitoring. Observe the climate system as required to allow policy-relevant improvements in prior estimates of the efficacy, benefits and harms of SRM. Examples include (i) current weather and climate observation systems, (ii) new global observation systems focused on the stratosphere and upper troposphere to improve measurement of atmospheric chemistry and aerosols, including instruments such as high-spectral resolution limb-sounders and new lidar instruments, and (iii) a systematic programme of *in situ* stratospheric observations.

For any stabilizing emissions pathway for greenhouse gases, the above scenario leads to a finite time deployment of SRM. This scenario is illustrated in Fig. 1 for an RCP4.5 emissions profile, with the corresponding global mean temperature and its rate of change predicted using MAGICC[6]. Note that the amount of SRM used under this scenario depends on the evolution of all other anthropogenic forcings, but it does not depend on climate sensitivity.

Further, while Fig. 1 maintains half the growth rate indefinitely, we explicitly include in our definition of this scenario the assumption that the amount of SRM would be adjusted in one direction or the other as time went on, based on what was learned about the impacts and risks either of uncompensated climate change or from SRM. While this analysis is driven by a fixed emissions trajectory, the intensity of efforts to restrain emissions — and thus the emissions trajectory — will also respond to new information.

We do not claim that this scenario is optimal. Rather we claim that good-quality policy-motivated scientific analysis requires an explicit scenario, and that this scenario is less obviously suboptimal than some scenarios employed in the literature. The GeoMIP[7]

© 2015 Macmillan Publishers Limited. All rights reserved

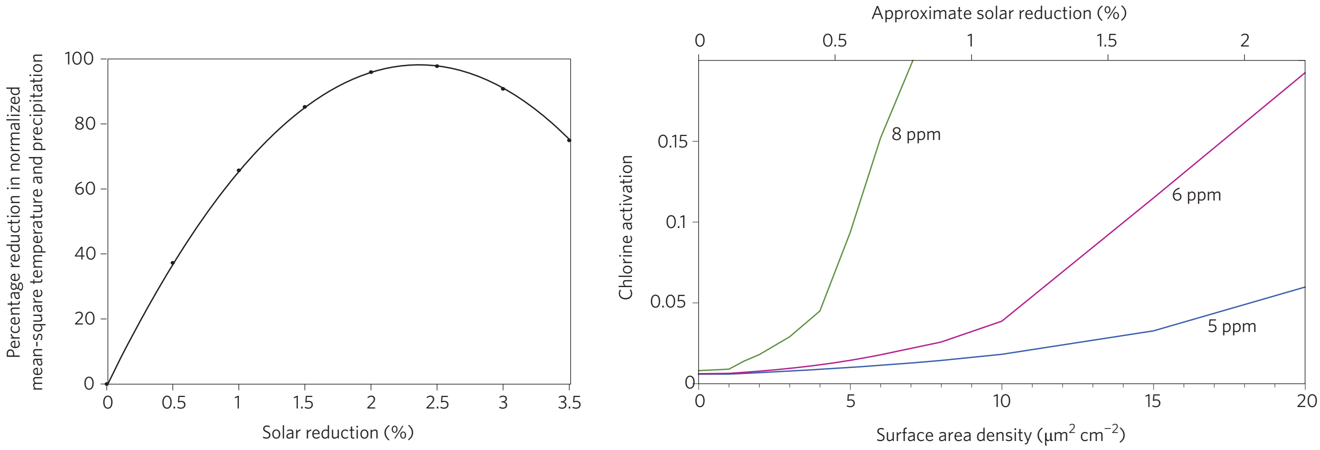

Figure 2 | The rationale for moderate SRM. Examples of the benefits and harms of SRM (left and right panels respectively) illustrating that benefits increase sublinearly and harms increase superlinearly. Whatever weighting is used to aggregate benefits and harms, the amount of SRM that maximizes the sum of benefits and harms will be less — perhaps much less — than the amount of SRM that maximizes benefits. The left panel shows a rough proxy for local climate damages, specifically the *y*-axis is the reduction in the sum of the mean of the quadratic deviation of temperature and precipitation across all climate model grid-points, where the deviation of each variable from its pre-industrial value is normalized by its interannual standard deviation. The climate model is HadCM3L (ref. 31) used with the methods and assumptions from ref. 8. The right panel shows chlorine activation as a function of surface area density (SAD) of sulphate aerosol computed using the AER model[32] under mid-latitude lower-stratosphere conditions. Chlorine activation, a crucial determinant of ozone loss, is strongly determined by water vapour concentration. Anderson *et al.*[17] provide a rationale for our choice of parameters. The secondary *x*-axis (top) shows an illustrative calculation of the corresponding solar reduction assuming that 0.5-μm radius sulphate aerosols were evenly dispersed over a 5-km height of altitude.

experiments G1 to G3, for example, assume that SRM is used at a level sufficient to compensate for all of the increase in anthropogenic forcings; yet, as we argue in 'Guiding principles' below, this cannot be an optimal balancing of benefits and risks. Although GeoMIP and other similar simulations are intended primarily to understand the climate response to SRM, the results are often interpreted as if they applied to SRM in general rather than as contingent on a particular — and we argue inappropriate — implementation scenario. Note that this is not a criticism of GeoMIP (indeed, our own previous work has used similar assumptions[8–10]), but rather of the overreaching interpretations of the results. Finally, our proposal builds on work of Smith and Rasch[11], who used peak-and-decline greenhouse gas concentration scenarios with a moderate amount of SRM.

Guiding principles

Three considerations shape our choice of scenario: moderation, responsiveness and impermanence.

Moderation (half measures). We define the benefits of SRM as the reduction in the magnitude or rate of climatic change due to anthropogenic greenhouse gas forcing; that is, the reduction in climate impacts. This definition is not trivial. Among other things, it ignores the fact that some regions or industries may benefit from anthropogenic climate change, and from their perspective a reduction in that climate change may therefore count as a harm rather than a benefit; and it ignores the extent to which the benefits (and harms) of climate change will be mediated by social, political, cultural and economic factors that themselves change as a consequence of social responses to climate change.

The impacts of climate change are primarily felt locally; that is, they depend on the local changes in variables such as temperature, precipitation and soil moisture. Analysis of the global utility of SRM therefore depends on how local benefits and harms are aggregated.

At one extreme, one can adopt the global optimal framing common in climate policy analysis. Under this assumption, the benefits of SRM first increase, then saturate and decline with increasing global radiative forcing. This holds true whatever weighting function is used to aggregate benefits across regions.

When the maximization of aggregate utility guides policy there is, usually, an implicit assumption that the winners will compensate the losers. The other extreme is Pareto's constraint, which states that a policy should be used to increase aggregate utility only so long as it makes no region worse off. For some choice of impact metric there will be regions that are worse off with any amount of SRM, so the Pareto-improving amount of SRM is zero. The same argument applies to mitigation: there are impact measures in which there are some winners as greenhouse gases increases, so the Pareto-improving amount of mitigation is zero. This extreme example serves as a warning against rigid application of Pareto's constraint. Almost any real-world public policy makes someone worse off. Rules that lie between global and Pareto optimality serve as better guides to policy than do either extreme. Analysis that demonstrates that the regional effectiveness of SRM is limited under Pareto's constraint should therefore be interpreted with caution[12,13].

Impacts from climate change are typically assumed to increase faster than linearly with the magnitude of the change (for example quadratic in Nordhaus[4], Weitzman[14], Goes *et al.*[15], cubic in Lempert *et al.*[16]). An example of the consequence of this assumption is shown in Fig. 2, where for illustration we have chosen a damage function that is quadratic in the local deviations of temperature and precipitation relative to a pre-industrial baseline. The specific choice of damage function is not critical to the argument, only the assumption that climate damages always increase faster than linearly, so that the 'benefits' of SRM (due to the intended reduction in climate changes) increase more slowly than linear, with the highest marginal benefit accruing initially.

Many of the technology-specific impacts of SRM are uncertain, but it seems plausible that many will increase faster than linearly. Mid-latitude ozone loss, for example, can be nonlinear owing to the threshold for heterogeneous chlorine activation as a function of water vapour, temperature and surface area density[17].

The benefits and costs in Fig. 2 are not in comparable units, so one cannot add them and find the amount of SRM that maximizes

© 2015 Macmillan Publishers Limited. All rights reserved

net benefits. Whatever the weighting of these benefits and cost functions, the ratio of benefits to costs will be largest for very small amounts of SRM. We assume that the broad features of these benefit and damage functions are typical features of SRM. This observation motivates our choice to cut the rate of growth of radiative forcing in half.

The most general — and strongest — statement about the need for moderation in the use of SRM is the following: assuming (i) the benefits of SRM are a concave function, or equivalently that the marginal benefits decrease with increasing radiative forcing, (ii) there is a maximum benefit beyond which increasing SRM causes harm even when the technology-specific impacts are ignored, and finally (iii) that technology-specific impacts increase faster than linearly with increasing SRM. Given these three assumptions one must conclude that the optimal amount of SRM will always be less than the amount of SRM that maximizes benefits.

Although this conclusion might appear trivial, a substantial fraction of the literature on SRM implicitly or explicitly assumes that SRM will be used to restore global temperatures to pre-industrial or to maintain radiative forcing at a fixed value.

Feedback (responsiveness). Both the climate system response and the performance of SRM technology are uncertain. It is not plausible that a planner would decide a century-long scenario for SRM implementation and stick to it independent of outcomes. All models are imperfect, and all observations suffer from errors and measure unforced climate variability ('noise') along with the response to forcing ('signal'). Whether or not it is initially planned, feedback is inevitable as the planner reacts to observations and discovers model errors[9,10].

If injection of aerosols into the stratosphere is causing an unexpected impact such as an increase in the opacity of upper tropospheric cirrus clouds then a rational planner will either: (i) abandon injection immediately if the impact is sufficiently large, (ii) phase out the SRM at a pace that balances the magnitude of the newly discovered impact against climate damages resulting from the rapid decrease in SRM radiative forcing, or (iii) continue injection while developing an alternative SRM technology that can be phased in as the original one is phased out so that the desired radiative forcing trajectory is maintained.

Our scenario includes feedback explicitly in that the use of SRM depends on the growth of other radiative forcing. If greenhouse gas emissions decline, so will the growth of SRM, whereas if tropospheric sulphate aerosol forcing declines, the growth of SRM will accelerate.

Feedback requires monitoring. Some variables such as radiative forcing and change in stratospheric chemistry that drive ozone loss could be detected with high signal-to-noise ratio (SNR), but detection and attribution of changes in regional climate caused by SRM can take decades because of low SNR[18]. With appropriate design, feedback control can nevertheless be effective in the face of climate variability and model errors[9,10], although it will take longer to detect and hence respond to some changes than others. Of course, a large unexpected impact is by definition detectable on rapid timescales. Our concept of responsiveness includes both tracking and responding to changes in radiative forcing, as well as changing the plan in response to unexpected impacts.

Most analysis of SRM is presumably intended to inform decisions about its use. A maxim of decision analysis is that variables over which the decision-maker has a choice are treated differently from exogenous variables or outcomes. The decision about how much SRM to implement is necessarily an iterated choice. If emissions decline or a risk of SRM is found to be larger than anticipated by available models, then a rational decision-maker will reduce the amount of SRM. Analysis that ignores feedback and treats the amount of SRM as exogenous or pre-determined may produce unrealistic conclusions about its risks and efficacy.

Temporary SRM. Without the use of carbon removal technologies, the climate impact of carbon emissions lasts for millennia, far longer than the inherent timescale for solar geoengineering defined either by the lifetime of injected aerosols (of the order of a week in the troposphere and a year in the stratosphere), or by the lifetime of the deployment hardware and support systems (years to decades).

A scenario in which SRM is maintained for millennia is, in our view, almost pure speculation as there is little basis for forecasting social and technological trends over such a timespan. We therefore consider scenarios in which SRM is temporary and for which the timescale of deployment is of the order of a century.

Even century-scale deployment raises concerns about the rapid climate change that would arise from a sudden termination of SRM. We distinguish two broad reasons for termination: new information or the inability to continue. We have already considered the implications of new information in our discussion of feedback. Inability to continue SRM might arise from war or some other social collapse. While not discounting the possibility of social collapse, we note that humanity has operated technologies such as trans-oceanic communication links and electric power grids for more than a century in spite of horrific wars. Moreover, in considering the implications of a possible social collapse on the public policy of SRM, one must set the risks of termination against the (likely) greater human suffering that would arise directly from the collapse itself[5].

Temporary deployment does not reduce long-term climate change. Warming in 2300, for example, is almost completely determined by cumulative carbon emissions and is unaffected by SRM that ends in 2200. Some commentators conclude that such temporary SRM offers no benefits, suggesting that it must be maintained forever (see Supplementary Information).

These claims ignore two facts. First, many climate change impacts depend on the rate of change. Recent climate changes are far more rapid than past changes (see Figs 3 and S1 in ref. 19), and projected changes are more rapid still (also shown in the same figures). The rate of change is important for the ecosystem's ability to adapt[19,20], as well as for human adaptation costs[15,16]; these latter citations suggest damage metrics that incorporate both absolute temperature changes and rate-dependent effects. Rate-dependent impacts of climate change have also been noted in studies concerned with the rapid climate change that might occur if solar geoengineering were suddenly terminated (for example refs 15, 21, 22). Second, some climate thresholds such as a possible shutdown of the thermohaline circulation — thought to be an important driver of climate system nonlinearity — depend on the rate of climate change[23].

It is clear that this scenario does not directly address thresholds that are a function only of the magnitude of the change rather than the rate, although it does delay reaching these thresholds, giving more time both to learn about the system and develop alternate strategies. For example in Fig. 1, the time to reach a temperature rise of 2 °C above pre-industrial increases from 2055 to 2068, while the time to reach a 2.5 °C rise increases by 32 years. Using SRM to limit the rate of warming has also been discussed by Eliseev *et al.*[24] and MacMartin *et al.*[10].

For simplicity, we have shown that temporary SRM can be useful in a scenario without carbon removal technologies — that is, even in a scenario where SRM can only slow the rate of warming but cannot limit its ultimate magnitude. If carbon removal is used, then temporary SRM can limit both the rate and peak magnitude of climate change. The use of carbon removal also eases — but does not remove — the risk inherent in the commitment to continue temporary SRM or suffer the climate changes associated with a rapid phase out.

A specific example

The above scenario and its justification are specific in terms of how to define the radiative forcing trajectory for SRM but not

© 2015 Macmillan Publishers Limited. All rights reserved

about how to produce it. To provide more context that might help in understanding this scenario, it is useful to consider some specifics about a particular way that it might be implemented, keeping in mind that this is only one possible approach, and there are other ways that would have different technology-specific impacts. Providing this level of detail on one possible approach serves to illustrate that: (i) the direct economic cost of initial SRM deployment would be so low that it is unlikely to play an important role in decisions by governments, and (ii) the technology development timeframe for initial deployment could be as short as a few years. This claim refers only to direct deployment costs and technological barriers; the costs of science and monitoring might exceed the cost of deployment for at least a decade, and the indirect benefits and harms of SRM are expected to be orders of magnitude larger.

Of the various approaches that have been suggested for SRM, the best understood is to introduce sulphate aerosol into the stratosphere. It is the only method that could be applied without substantial further technical development to generate global radiative forcing of a similar magnitude to greenhouse gas forcing.

The amount of sulphate aerosol required as a function of time depends on the forcing scenario and on the radiative forcing per unit of sulphate. For radiative forcing less than about 0.5 W m^{-2} the radiative forcing efficacy is about 0.6–0.8 W m^{-2} for an injection rate of one million tons of sulphur (MtS) per year for most proposed methods of introducing sulphate[25]. So in the first decade of the scenario shown in Fig. 1 the rate at which the sulphur addition would increase — starting from zero in 2020 — would be 0.035 MtS yr^{-2}. That is, at the end of the first year the injection rate would be 0.035 MtS yr^{-1} and after a decade it would be 0.35 MtS yr^{-1}.

Feedback control — responsiveness — could be used to ensure that the global radiative forcing increased as intended even if the efficacy per unit sulphur is uncertain. Measurement of radiative forcing from tropospheric aerosols is difficult because of their complex indirect effects on clouds, but stratospheric radiative forcing can probably be estimated with far greater accuracy. The aerosol distribution could, for example, be estimated from a combination of orbiting limb-sounders and lidars corroborated using *in situ* observations from which the radiative forcing could then be accurately estimated using optimal estimation methods. This approach rests on the fact that the main uncertainty in *a priori* estimates of radiative forcing for a given sulphate injection rate is in predicting the aerosol distribution, while prediction of radiative forcing given an aerosol distribution is far less uncertain.

Although delivery mechanisms have not been designed in detail, analysis suggests that the most cost-effective approach uses aircraft[26]. Initially this might involve retrofitting business jets with off-the-shelf low-bypass ratio engines to allow them to fly at higher altitudes. Using McClellan and co-workers' analysis of re-engined G650 aircraft that include industry standard estimates of aircraft availability and flight rates, and assuming that the payload is liquid sulphur that is oxidized *in situ,* about two aircraft would be required in the first year, rising to 30 by 2040. The capital cost of purchasing and modifying these 30 aircraft would be roughly US$2.2 billion.

Deployment could begin with SO_2 but as the aerosol concentration increases, more of the added sulphate simply adds to the mass of existing aerosol, increasing aerosol size and so reducing the efficacy per unit sulphate[27]. This problem can be avoided by direct release of H_2SO_4 from an aircraft as proposed by Pierce *et al.*[25]. (Note that while the work of English *et al.*[28] appears to contradict this result, it simulates a process in which H_2SO_4 is perfectly mixed at the grid scale of the general circulation model, which does not — and would not be expected to — produce a result significantly different from the SO_2 oxidation case simulated by Heckendorn[27].)

If a decision were made to deploy SRM, we assume that efforts to develop new technologies would be pursued more actively so they might be available if problems with sulphate were worse than expected. This might include particles with less ozone impact, or particles with more efficient back-scattering (thus requiring fewer of them), or possibly space-based systems.

Critical to any plausible implementation scenario is monitoring of its effects so that either the amount or the implementation technology can be modified based on new information. As noted, it is much easier to detect the radiative forcing (or ozone chemistry impacts) than it is to detect the impact on regional climate variables such as temperature or precipitation because the former will have a much higher SNR. Any response that is too small to detect in the presence of natural variability should also not result in significant negative consequences, although any unusual weather extremes may be blamed on the SRM deployment nonetheless.

One approach that might help in evaluating the climate response due to SRM is to introduce some time-varying modulation of the SRM radiative forcing. The response due to SRM can then be estimated by looking for the correlated signal in any climate variable. It is easier to distinguish a time-varying response from background variability than it is to distinguish a secular change in SRM radiative forcing from similar changes in greenhouse gas concentrations. Even with modulation, it could take decades to be confident in attributing regional impacts to SRM[18]; however, modulation could allow earlier and more accurate detection of impacts on the chemistry and dynamics of the stratosphere where the signal-to-noise ratios will be much larger.

Discussion

First and most simply, this scenario demonstrates that there may be value to temporary SRM. Humanity is not committed forever once SRM begins; rather, there is an implied commitment to a measured wind down rather than an abrupt termination.

We have explored SRM as a complement to mitigation in that we assume that SRM is used to reduce climate risks while mitigation proceeds. Here we compare this to other framings.

SRM is often considered as a substitute for mitigation. In the extreme case this means the use of SRM without any reduction of emissions. This could be effective in the short run but would be totally ineffective in the long run as greenhouse gas concentrations would rise without limit. More plausibly, SRM could be a partial substitute, although this entails risks due to the increased greenhouse gas concentrations. These risks are linked through the choice of policy although they are physically unrelated to SRM.

Alternately, SRM might be used only in case of a climate emergency[29]. Our view is that if SRM is seriously contemplated (developed, governed and incorporated into climate policy) as an emergency measure, then it arguably makes more sense to begin some gradual and moderate SRM as a precursor. The reasons are primarily about providing time for learning. Reasons include: (i) starting early gives more time to learn about SRM effects, and how to do SRM better, as well as more time to learn about mitigation; (ii) starting at a small forcing amplitude provides a better environment for finding bad unknown-unknowns, as the consequences will be less severe at a small amplitude than at large; (iii) if there is a 'tipping point' beyond which climate impacts increase steeply, then the scenario described here would delay reaching it (assuming we are not already beyond it) and thus give more time to learn about it; and finally, (iv) moderate and gradual use of SRM provides a basis to develop governance mechanisms, whereas a 'climate emergency' might well be the worst circumstance for developing methods to govern a new technology like SRM.

The converse argument is that if SRM is intended only for emergencies then there is less chance it will be used. This is a preferred outcome if (i) the risks of SRM prove so large that even for partial temporary SRM they outweigh its benefits outside an emergency; or if (ii) socio-technical lock-in[30] is sufficiently strong that starting SRM amounts to a *de facto* commitment to use it at large

© 2015 Macmillan Publishers Limited. All rights reserved

scale. The central planner framing we have adopted here ignores the institutional factors that create strong lock-in.

Finally, we have shown that temporary SRM can be useful without use of carbon removal; but the technologies are complementary in that carbon removal can allow temporary SRM to limit both the rate and absolute magnitude of climate change.

The central message of this paper is not that the proposed scenario is likely or optimal, it is simply that analysis of SRM that is intended to inform policy should — at a minimum — be explicit about the implementation scenario that drives the analysis and about the way that conclusions are dependent on the scenario choice.

Received 18 February 2014; accepted 9 December 2014; published online 16 February 2015

References

1. Robock, A. 20 reasons why geoengineering may be a bad idea. *Bull. Atom. Sci.* **64,** 14–18 (2008).
2. Bala, G., Duffy, P. B. & Taylor, K. E. Impact of geoengineering schemes on the global hydrological cycle. *Proc. Natl Acad. Sci. USA* **105,** 7664–7669 (2008).
3. Kleidon, A. & Renner, M. A simple explanation for the sensitivity of the hydrological cycle to surface temperature and solar radiation and its implications for global climate change. *Earth Syst. Dynam.* **4,** 455–465 (2013).
4. Nordhaus, W. *A Question of Balance: Weighing the Options on Global Warming Policies* (Yale Univ. Press, 2008).
5. Keith, D. W. & Parker, A. The fate of an engineered planet. *Sci. Am.* **308,** 34–36 (2013).
6. Meinshausen, M., Raper, S. C. B. & Wigley, T. M. L. Emulating coupled atmosphere–ocean and carbon cycle models with a simpler model, MAGICC6. Part 1: Model description and calibration. *Atmos. Chem. Phys.* **11,** 1417–1456 (2011).
7. Kravitz, B. *et al.* The Geoengineering Model Intercomparison Project (GeoMIP). *Atmos. Sci. Lett.* **12,** 162–167 (2011).
8. MacMartin, D. G., Keith, D. W., Kravitz, B. & Caldeira, K. Management of trade-offs in geoengineering through optimal choice of non-uniform radiative forcing. *Nature Clim. Change* **3,** 365–368 (2013).
9. MacMartin, D. G., Kravitz, B., Keith, D. W. & Jarvis A. J. Dynamics of the couple human-climate system resulting from closed-loop control of solar geoengineering. *Clim. Dyn.* **43,** 243–258 (2014).
10. MacMartin, D. G., Caldeira, K. & Keith, D. W. Solar geoengineering to limit rates of change, *Phil. Trans. R. Soc. A* **372,** 20140134 (2014).
11. Smith, S. J. & Rasch, P. J. The long-term policy context for solar radiation management. *Clim. Change* **121,** 487–497 (2013).
12. Moreno-Cruz, J., Ricke, K. & Keith, D. W. A simple model to account for regional inequalities in the effectiveness of solar radiation management. *Clim. Change* **110,** 649–668 (2011).
13. Kravitz, B. *et al.* A multi-model assessment of regional climate disparities caused by solar geoengineering. *Environ. Res. Lett.* **9,** http://dx.doi.org/10.1088/1748-9326/9/7/074013 (2014).
14. Weitzman, M. What is the damages function for global warming—and what difference might it make? *Clim. Change Econ.* **1,** 57–69 (2010).
15. Goes, M., Tuana, N. & Keller, K. The economics (or lack thereof) of aerosol geoengineering. *Clim. Change* **109,** 719–744 (2011).
16. Lempert, R. J., Schlesinger, M. E., Bankes, S. C. & Andronova, N. G. The impacts of climate variability on near-term policy choices and the value of information. *Clim. Change* **45,** 129–161 (2000).
17. Anderson, J. G., Wilmouth, D. M., Smith, J. B. & Sayres, D. S. UV dosage levels in summer: Increased risk of ozone loss from convectively injected water vapor. *Science* **337,** 835–839 (2012).
18. MacMynowski, D. G., Keith, D. W., Caldeira, K. & Shin, H-J. Can we test geoengineering? *Energy Environ. Sci.* **4,** 5044–5052 (2011).
19. Diffenbaugh, N. S. & Field, C. B. Changes in ecologically critical terrestrial climate conditions. *Science* **341,** 486–492 (2013).
20. Schloss, C. A., Nuñez, T. A. & Lawler, J. J. Dispersal will limit ability of mammals to track climate change in the Western Hemisphere. *Proc. Natl Acad. Sci. USA* **109,** 8606–8611 (2012).
21. Wigley, T. M. L. A combined mitigation/geoengineering approach to climate stabilization. *Science* **314,** 452–454 (2006).
22. Matthews, H. D. & Caldeira, K. Transient climate-carbon simulations of planetary geoengineering. *Proc. Natl Acad. Sci. USA* **104,** 9949–9954 (2007).
23. Stocker, T. F. & Schmittner, A. Influence of CO_2 emission rates on the stability of the thermohaline circulation. *Nature* **388,** 862–865 (1997).
24. Eliseev, A. V., Chernokulsky, A. V., Karpenko, A. A. & Mokhov, I. I. Global warming mitigation by sulphur loading in the stratosphere: dependence of required emissions on allowable residual warming rate. *Theor. Appl. Climatol.* **101,** 67–81 (2010).
25. Pierce, J. R., Weisenstein, D. K., Heckendorn, P., Peter, T. & Keith, D. W. Efficient formation of stratospheric aerosol for climate engineering by emission of condensable vapor from aircraft. *Geophys. Res. Lett.* **37,** L18805 (2010).
26. McClellan, J., Keith, D. W. & Apt, J. Cost analysis of stratospheric albedo modification delivery systems. *Environ. Res. Lett.* **7,** 034019 (2012).
27. Heckendorn, P. *et al.* The impact of geoengineering aerosols on stratospheric temperature and ozone. *Environ. Res. Lett.* **4,** 045108 (2009).
28. English, J. M., Toon, O. B. & Mills, M. J. Microphysical simulations of sulfur burdens from stratospheric sulfur geoengineering. *Atmos. Chem. Phys.* **12,** 4775–4793 (2012).
29. Blackstock, J. J. *et al. Climate Engineering Responses to Climate Emergencies* (Novim, 2009); http://arxiv.org/pdf/0907.5140
30. Geels, F. W. From sectoral systems of innovation to socio-technical systems: Insights about dynamics and change from sociology and institutional theory. *Res. Policy* **33,** 897–920 (2004).
31. Jones, C. A fast ocean GCM without flux adjustments. *J. Atm. Ocean Technol.* **20,** 1857–1868 (2003).
32. Rinsland, C. P. *et al.* Post-Mount Pinatubo eruption ground-based infrared stratospheric column measurements of HNO_3, NO, and NO_2 and their comparison with model calculations. *J. Geophys. Res.* **108,** 4437 (2003).

Additional information

Supplementary information is available in the online version of the paper. Reprints and permissions information is available online at www.nature.com/reprints. Correspondence and requests for materials should be addressed to D.W.K.

Competing financial interests

The authors declare no competing financial interests.

© 2015 Macmillan Publishers Limited. All rights reserved