



The value of information about solar geoengineering and the two-sided cost of bias

Anthony R. Harding, Mariia Belaia & David W. Keith

To cite this article: Anthony R. Harding, Mariia Belaia & David W. Keith (2022): The value of information about solar geoengineering and the two-sided cost of bias, Climate Policy, DOI: [10.1080/14693062.2022.2091509](https://doi.org/10.1080/14693062.2022.2091509)

To link to this article: <https://doi.org/10.1080/14693062.2022.2091509>



© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 29 Jun 2022.



[Submit your article to this journal](#)



Article views: 639




[View related articles](#)



[View Crossmark data](#)

The value of information about solar geoengineering and the two-sided cost of bias

Anthony R. Harding ^{a,b}, Mariia Belaia^c and David W. Keith^{a,b}

^aThe John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA; ^bThe John F. Kennedy School of Government, Harvard University, Cambridge, MA, USA; ^cForum for Climate Engineering Assessment, American University, Washington, DC, USA

ABSTRACT

Solar geoengineering (SG) might be able to reduce climate risks if used to supplement emissions cuts and carbon removal. Yet, the wisdom of proceeding with research to reduce its uncertainties is disputed. Here, we use an integrated assessment model to estimate that the value of information that reduces uncertainty about SG efficacy. We find the value of reducing uncertainty by one-third by 2030 is around \$4.5 trillion, most of which comes from reduced climate damages rather than reduced mitigation costs. Reducing uncertainty about SG efficacy is similar in value to reducing uncertainty about climate sensitivity. We analyse the cost of over-confidence about SG that causes too little emissions cuts and too much SG. Consistent with concerns about SG's moral hazard problem, we find an over-confident bias is a serious and costly concern; but, we also find under-confidence that prematurely rules out SG can be roughly as costly. Biased judgments are costly in both directions. A coin has two sides. Our analysis quantitatively demonstrates the risk-risk trade-off around SG and reinforces the value of research that can reduce uncertainty.

Key policy insights:

- The value of reducing uncertainty about solar geoengineering is comparable to the value of reducing uncertainty about other key climate factors, such as equilibrium climate sensitivity.
- The benefits of research that reduces uncertainty about solar geoengineering may be more than a thousand times larger than the cost of a large-scale research programme.
- Under-confidence in solar geoengineering's effectiveness can be as costly as over-confidence.
- The majority of the benefits of reduced uncertainty come from reducing climate damages rather than from slowing emissions reductions.

ARTICLE HISTORY



Received 17 September 2021
Accepted 14 June 2022


KEYWORDS

Geoengineering; solar radiation management; value of information; uncertainty

Introduction

Solar geoengineering (SG) could reduce the climatic changes from a given concentration of greenhouse gases. Combined with emissions reductions, adaptation, and carbon dioxide removal (CDR), it could reduce climate impacts in ways not possible with emissions cuts alone. For example, there is evidence that it would be possible to maintain a reasonably uniform spatial distribution of stratospheric aerosols (Dai et al., 2018; Kravitz et al., 2017) at a cost that is very low compared to other climate actions (Smith & Wagner, 2018) to moderate climate hazards, such as changes in water availability, crop yields, and peak temperatures (Irvine & Keith,

CONTACT Anthony R. Harding  tonyharding@hks.harvard.edu  12 Oxford St., Cambridge, MA 02138, USA

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/14693062.2022.2091509>.

© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

2020; National Academies of Sciences, 2021). Yet, uncertainties about SG efficacy and risks are far too large to justify immediate decisions about its deployment. The central question for SG policy today is whether to establish a large-scale research programme to reduce uncertainties and clarify risks (National Academies of Sciences, 2021).

Uncertainties raise the costs of climate policy. Uncertainty about climate sensitivity and the cost of emissions cuts makes it hard for decisionmakers to balance the pace of emissions cuts against climate impacts. Since impacts depend on cumulative emissions, decisions made today have lasting consequences. Thus, the climate's inertia exacerbates the costs of these uncertainties. If SG is included in the policy mix, decisionmakers must also account for uncertainty about SG. The value of research on SG comes from reducing uncertainty.

SG research has many possible outcomes, from spurring discoveries in atmospheric science to the risk that it builds constituencies that lobby for SG implementation against the public interest. Here, we evaluate the benefits of reducing scientific uncertainty about SG from the perspective of global climate policy using an integrated assessment model (IAM). Specifically, we estimate the value of information that eliminates uncertainty in SG efficacy – how well it moderates or exacerbates the climate effects of rising greenhouse gas concentrations (GHGs). The methodology of using an IAM to estimate the value of information has been widely applied to other important climate policy uncertainties (For a review, see Golub et al., 2014).

Concerns about SG predominantly centre on the risks of emissions reduction deterrence today and over reliance on SG in the future. Any number of complex political factors might produce these risks. One prominent example is SG's moral hazard problem – the concern that SG will inhibit emissions mitigation (McLaren & Corry, 2021). This risk may arise from a widespread social bias towards over-confidence about the performance of new technologies (Keith, 2000). Or it may arise if a self-interested minority with an interest in avoiding energy system decarbonization, such as fossil-fuel-rich countries or industries, is able to overcome the majority by exaggerating the efficacy of SG (Keith, 2021). Or it may arise from socio-technical lock-in if SG develops a self-interested constituency committed to advancing its use (Cairns, 2014; McKinnon, 2019).

These are sensible concerns. However, it is important to consider the other side of the coin. These risks of over-optimism need to be weighed against the risks of over-pessimism (Parson, 2021). The consequences of

		REALITY	
		SG Works	SG Fails
BELIEF	SG Works	Correct	Over-confidence
	SG Fails	Under-confidence	Correct

Figure 1. Risk-Risk Trade-off. A simple illustration of the risk-risk trade-off of SG represented as a table of error types.

these risks are evident in a simplistic binary version of the decision problem (Figure 1). Independent of the cause, the complex political factors that might produce these risks can be represented as a bias towards over-optimistic beliefs about SG's effectiveness. We introduce two parsimonious frameworks that leverage this characterization of risk as an over-optimistic bias about SG to quantitatively estimate and illustrate the risk-risk trade-offs around SG.

Other studies have analysed the value of SG research and its consequences. Goeschl et al. (2013) and Quaas et al. (2017) use economic theory to characterize intergenerational incentives and implications of SG research and technology transfer. Arino et al. (2016) estimate the option value of SG – the value of having it available in the future – with uncertain climate sensitivity. Closest to this work, Moreno-Cruz and Keith (2013) use a two-period model to estimate both the value of SG availability with uncertain climate sensitivity and the value of reducing uncertainty about SG's side-effects. We contribute to this literature by using a calibrated IAM to estimate the value of information about SG efficacy. Our quantitative illustrations of the risk-risk trade-offs that compare the costs of over-optimism about SG's capabilities to the costs of prematurely ruling out SG further contributes to the literature by introducing the role of biased decision-making.

We use a global utilitarian framework to explore the value of reducing uncertainty about SG and illustrate its risk-risk trade-offs. But the world is not – of course – managed by a single benevolent decisionmaker. Any decisions about establishing a large-scale research programme for SG, like any possible future decisions about deployment, are complex and consequential and thus require careful and scientifically-grounded evaluation. By abstracting from the complexities of legal, ethical, and geopolitical uncertainties (Flegal et al., 2019), our analysis provides insight into the benchmark towards which practical policy might strive. Understanding the consequences of reducing uncertainty about SG, or of biased judgments about SG, informs practical decisions about establishing a solar geoengineering research programme.

Model

We model climate policy decision-making using an extended version of the Dynamic Integrated model of Climate and the Economy (DICE) (Nordhaus, 2017) developed by Belaia et al. (2021). DICE integrates a simple climate model and a macroeconomic growth model to optimize climate policy that aims to maximize global welfare by balancing the costs of climate policy against climate damages. The model is extended model to include the choice of using CDR and SG.

In Figure 2, we illustrate how SG is characterized in the model. A key characteristic of SG is that it is imperfect – it cannot perfectly reverse GHG-driven climatic changes and will introduce novel risks. These imperfections and risks are distinctly incorporated in the model through an efficacy parameter (Blue) and a direct risk parameter (Red), respectively. Our analysis focuses on uncertainty in SG efficacy. We neglect uncertainty in direct risks of SG because variations around its calibrated value produce small changes in climate policy in comparison to variations in SG efficacy (Belaia et al., 2021).

While SG can reduce some GHG-induced climatic changes, it does so unequally, so that changes in some regions or climatic variables may be exacerbated even as the global average temperature is reduced (National Academies of Sciences, 2021). This unevenness in climate response is captured by the efficacy of SG, which we parametrize as the angle θ between linearized climate responses to GHGs and SG in a vector space defined by the number of regions (Figure 2) (Moreno-Cruz et al., 2012). The traditional DICE damage function – which measures climate damages as a function of changes in global climate – is scaled by the SG efficacy parameter so that net climate damages are proportional to the global weighted sum of the squared deviation of regional climates (see Belaia et al., 2021 for the derivation).

We calibrate uncertainty in the SG efficacy parameter, θ , by fitting a log-normal distribution to recent climate model simulations (Irvine et al., 2019; Irvine & Keith, 2020; Kravitz et al., 2014; Moreno-Cruz et al., 2012), as described in the Supplementary Materials. Using a long-tailed distribution to capture tail risk, we adopt a calibration of θ that is both more uncertain than justified by the spread of existing empirical evidence and is biased towards underestimating the efficacy of SG. For computational feasibility, we discretize the calibrated distribution to five possible representative states-of-the-world (SOWs) by matching the first three




Figure 2. Model Schematic. Illustration of how SG is integrated into the extended DICE model. The left-hand side shows how the DICE damage function is augmented to include direct risks (Red) and limits to the efficacy of SG in moderating GHG-induced changes in climate (Blue). Direct risks are proportional to the amount of SG and include the costs of deployment and monitoring along with the damages directly dependent on the SG mechanism, such as ozone loss or health impacts of stratospheric aerosol. SG may moderate or exacerbate local GHG-induced climate changes. This is captured by an efficacy parameter (Blue), illustrated in the right-hand side of the figure. Regional climate responses to GHG and SG are linearized and represented as vectors in a space defined by the number of regions. The SG response vector can be decomposed into a moderating component, which exactly counteracts the effects of GHGs, and an orthogonal component which represents the imperfections in SG’s ability to restore climate changes. The angle, θ , between these linearized response vectors determines our parameterization of SG’s efficacy. Climate damages are multiplicatively scaled by the efficacy value, E , for the level of GHG-induced change in RF restored by SG, given by g . The bottom panel illustrates the efficacy of SG for different efficacy angles and levels of SG deployment. At an angle of $\theta = 0^\circ$, for example, SG that restores all GHG-induced RF change eliminates all climate damages; but, at an angle of $\theta = 45^\circ$, SG exacerbates some local climate changes and moderates others so it will only reduce aggregate climate damages by around 40%.

moments of the continuous distribution. Ranging from 2° to 47° , these representative values imply that SG has sufficient efficacy to reduce global climate damages by 99% to 45%. Again, note that SG that reduces global climate damages may moderate climate damages in some regions while exacerbating damages in others.

The direct risks of SG capture the risks of sulphate aerosol injection, such as ozone loss or increases in ground-level air pollution, and the costs of deployment and monitoring. These are parameterized separately in the damage function as a linear function of the SG radiative forcing. Calibration based on empirical evidence gives a cost of 0.1% of Gross World Product (GWP) per W/m^2 of SG.

The expected value of perfect information (EVPI)

We estimate the value of information about SG as the expected value of perfect information (EVPI). The EVPI represents the difference between expected outcomes given perfect information that eliminates all uncertainty (EV|PI) and the expected outcomes for unresolved uncertainty (EV). We use the expected net present value (NPV) of consumption as the outcome of interest. The difference in expected consumption between the EV|PI and EV scenarios captures the decisionmaker’s willingness-to-pay today for perfect information. This gives

the following equation for EVPI:

$$EVPI = EV|PI - EV = \sum_{s=1}^S p_s \left(\sum_{t=1}^T \frac{C_{s,t}^{PI}}{(1+r)^t} \right) - \sum_{s=1}^S p_s \left(\sum_{t=1}^T \frac{C_{s,t}^{NL}}{(1+r)^t} \right) \quad (1)$$

where p_s is the probability of state s , t is the time period, r is the discount rate, and C^{PI} and C^{NL} is consumption for the perfect information (PI) and no learning (NL) scenarios, respectively.

We assume that perfect information is anticipated by the decisionmaker and that it is exogenously provided in a specific year of learning. For our benchmark analysis, we use 2050 as the year of learning. For our NPV calculations throughout the analysis, we use a constant exogenous discount rate of $r = 3\%$ and only include consumption from 2020 to 2200. In the Supplementary Materials we perform a sensitivity analysis for the discount rate and other key parameters.

As a benchmark, we display outcomes for the perfect information and no learning scenarios when the year of learning is 2050 (Figure 3(a-d)). With no learning, the decisionmaker must set climate policy under the assumption that SG efficacy can take any of the five possible SOWs. Consistent with intuition, if the decisionmaker learns efficacy is high (i.e. theta is low), then optimal policy uses less mitigation and CDR than in the no learning case and the converse if efficacy is low (i.e. theta is high) (Figure 3(a-b)). Arrival of perfect information leads to abrupt and widely varying adjustments in climate policy. For example, SG implementation ranges from 2 to 6 Wm^{-2} in 2100 and the year of net-zero emissions ranges over 200 years. This illustrates the potential influence of SG efficacy in climate policy.

We find that the EVPI for SG efficacy is around \$9.3 trillion when learning occurs in 2020 (Figure 3(e)). This value declines approximately linearly with the year of learning to around \$5 trillion for learning in 2120. With later learning, there is less opportunity to reap the benefits of information. Potentially avoidable climate damages and policy costs that have been incurred cannot be reversed. This indicates that there are significant gains to learning about SG early.

The EVPI can be decomposed into the value of reductions in climate damages and the value of reductions in carbon policy costs – both mitigation and CDR costs (Figure 3(e)). We find that the majority of the value of information about SG efficacy stems from reductions in climate damages rather than money saved by slowing emissions reduction efforts, contrary to common expectations. This is because variation in climate damages across possible SG efficacy values dominates variation in policy cost (Figure 3(c-d)).

Receiving information that resolves *all* uncertainty around SG efficacy is an unrealistic assumption. Research is slow and it cannot – of course – eliminate all uncertainty. While our perfect information estimates provide an upper-bound estimate on the decisionmakers willingness-to-pay to reduce uncertainty, we also consider more policy-relevant values. We estimate the expected value of partial information as the value of sample information that *partially* reduces uncertainty around SG efficacy (see Supplementary Materials). We find that the expected value of partial information in 2030 is slightly concave in its fractional reduction of the width of uncertainty (Figure 3(f)). For example, receiving information in 2030 that reduces uncertainty about SG efficacy by a third has an expected value of around \$4.5 trillion – about a half the EVPI. This indicates that there is significant value to research about SG even if it is unable to resolve all uncertainty.

For context, we also estimate the EVPI about equilibrium climate sensitivity (ECS), which is treated as one of the most important uncertainties in climate change and is studied extensively (IPCC, 2021; Sherwood et al., 2020). For example, in a review, Knutti et al. (2017) cite over 100 estimates published in the preceding 10 years. We construct a probability distribution for ECS following historical evidence as shown in Figure 5 of Knutti et al. (2017). The distribution is again discretized into five possible SOWs using the same moment matching process as used for SG efficacy. We assume that SG efficacy takes its median value of 15° .

For learning in 2020, we estimate that the EVPI about ECS is around \$18 trillion when SG is not available and around \$6 trillion when SG is available (Figure 3(e)). As with SG, the EVPI for ECS declines as the year of learning is delayed. This is exacerbated in the scenario without SG because of the climate's inertia. Comparing these to our estimate for EVPI about SG of \$9.3 trillion suggests that, even if there are unitary uncertainties in the estimation of the EVPI, the value of learning about SG is comparable to the value of learning about ECS.

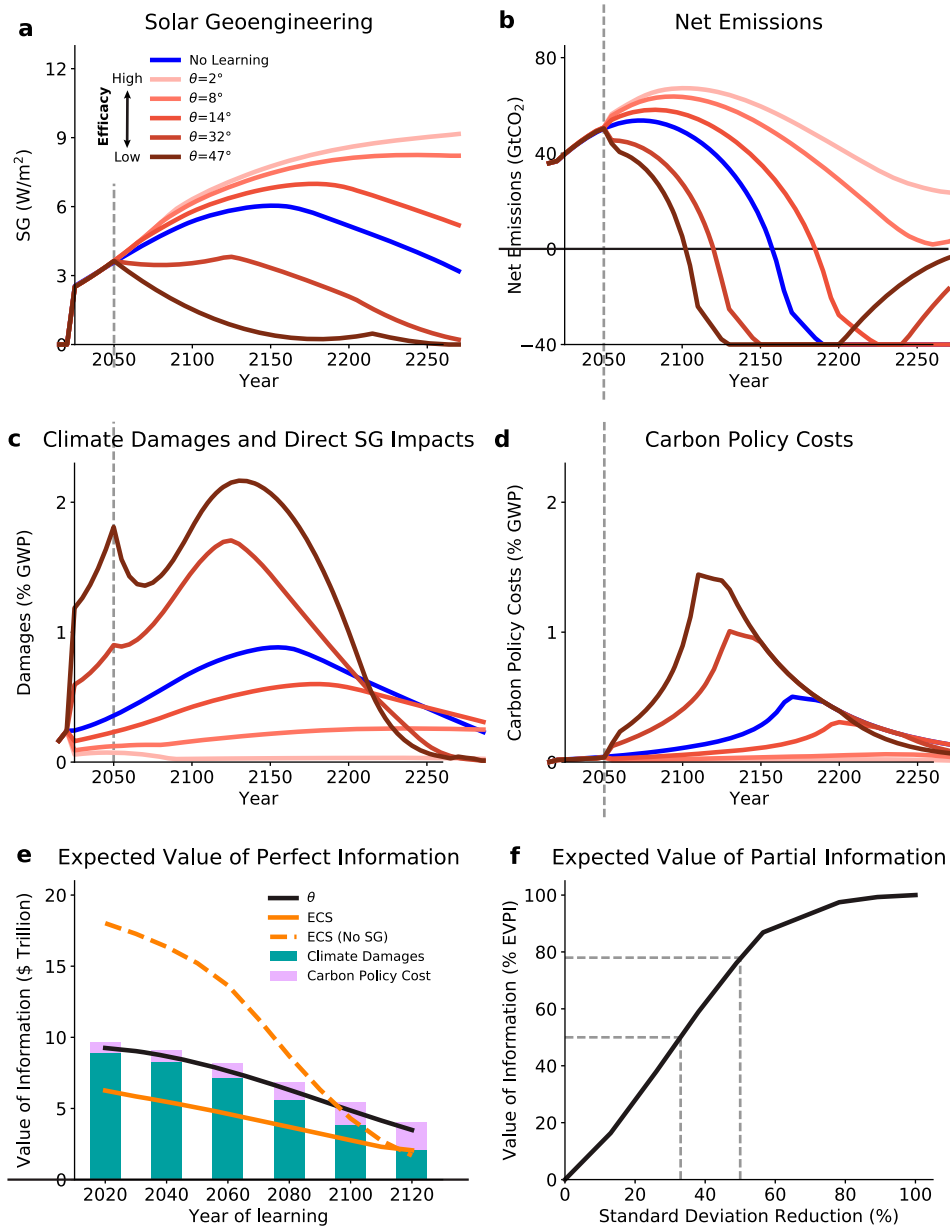


Figure 3. Value of learning about SG efficacy. (a) and (b) display the cost-benefit optimized climate instrument choices over time in the case of no learning (blue) and perfect learning in 2050 (red). Net emissions represent business-as-usual emissions net of emissions reductions and CDR. (c) and (d) display damages and carbon policy instrument costs – both mitigation and CDR costs – for no learning (blue) and perfect learning in 2050 (red). For no learning, mean damages and costs are shown. (e) displays the value of information about SG efficacy (black), ECS (orange solid), and ECS with no SG available (orange dashed). Green and pink bars show contributions to the value of information about SG efficacy from reduced climate damages (green) and reduced carbon policy costs (pink). (f) displays the expected value of partial information about SG efficacy as a function of the percent reduction in the standard deviation of priors for learning in 2030. Grey lines denote 33% and 50% reduction.

We evaluate the cost of uncertainty for SG and ECS in isolation, but in reality these costs coexist and interact. If SG works, it is a quick but imperfect substitute for emissions cuts and reduces the value of information about ECS by reducing damages when ECS is high. Conversely, if SG does not work, this raises the value of information

about ECS (Figure 3(e)). Other sources of uncertainty not considered here, such as in the damage function or emissions reduction costs, will also interact with these uncertainties (Ricke et al., 2012).

The cost of bias

Our estimates of the EVPI about SG assumes that the decisionmaker's knowledge of SG is uncertain and also unbiased. What if their beliefs are biased? Our assumption of a single, global central planner also neglects salient concerns about SG, such as moral hazard and socio-technical lock-in. While these concerns are distinct and complex, in the simple model we use here they amount to a bias in favour of deployment. Here, we explore the cost of bias using two parsimonious frameworks.

In the first framework, we remove all uncertainty around SG efficacy but introduce a bias between the decisionmaker's belief about SG efficacy, which they use to set climate policy, and reality. That is, the decisionmaker could set policy with the belief that SG works when in reality it does not, or vice versa (Figure 1). To flesh out this simple binary model of SG risk, we consider a range of SG efficacy values for both the decisionmakers belief and the true value. The cost of bias is calculated as the difference in the NPV of consumption for policy set by a decisionmaker with biased belief about SG efficacy and policy set by a decisionmaker who knows the true value. This captures the decisionmaker's willingness-to-pay to avoid bias. We assume the decisionmaker never learns about their bias.

Along the diagonal, the decisionmaker is unbiased and thus the costs of bias are zero (Figure 4(a)). Comparing across the unbiased diagonal, we find that over-optimistic bias about SG – believing SG is more efficacious than reality – can be more costly than over-pessimistic bias. However, both directions of bias can have large costs. For example, the cost of over-optimistically believing SG efficacy is 0° when it is really 45° is around \$115 trillion and the cost of over-pessimistically believing SG efficacy is 45° when it is really 0° is lower but still large at around \$95 trillion. This quantitative illustration indicates that ruling out SG when it may work could be as costly as relying on SG when it may not work.

Figure 4(b-c) decompose the costs of bias into their climate damage and carbon policy cost components, respectively. Over-optimistic bias can be beneficial by slowing emissions reduction and CDR efforts, thus lowering carbon policy costs; however, these benefits are outweighed by increased climate damages. This finding is consistent with salient concerns about the moral hazard problem and techno-optimism. For over-pessimistic bias, small differences between beliefs and reality increase emissions cuts and CDR which reduces climate damages; however, this benefit is outweighed by the corresponding higher costs of climate policy. For large differences between beliefs and reality, an over-pessimistic bias leads to both higher policy costs and higher climate damages.

In our second framework, we simultaneously consider both uncertainty and bias about SG. Using our calibrated probability distribution for SG efficacy from above, we represent bias as a multiplicative scaling factor, s , ranging from 0 to 1 (Figure SM.2). The parameter s biases beliefs by multiplicatively scaling values in the calibrated probability distribution by $1-s$. At $s=0$, there is no scaling so the decisionmaker's beliefs are unbiased and are consistent with the true probability distribution of SG efficacy. At $s=1$, the decisionmaker's beliefs are biased such that they believe, with certainty, that SG works perfectly as a substitute for emissions reductions, i.e. $\theta=0$. For values between 0 and 1, bias in decision-making causes too little use of carbon control policies and an over-deployment of SG, capturing many concerns about SG.

In response to concerns about SG, some have argued for a ban on SG deployment. For comparison, we also estimate the cost of a temporary global moratorium on SG deployment. For both the bias and global moratorium scenarios, we assume that in the year 2050, the decisionmaker receives perfect information about the true value of SG efficacy. This information resolves all uncertainty, bias, and moratoriums. To measure the costs of bias, we calculate the difference in NPV of consumption for policy set under biased beliefs and policy under unbiased beliefs. We measure the cost of a moratorium as the difference in NPV of consumption for policy set under a moratorium and policy without a moratorium.

As the decisionmaker's beliefs become more biased, that is as s increases from 0 to 1, the cost of bias increases from \$0 to \$1.9 trillion (Figure 4(d)). The majority of these costs are attributable to the costs of over-confidence when SG efficacy is low. The costs of bias are greatest when SG does not work. In

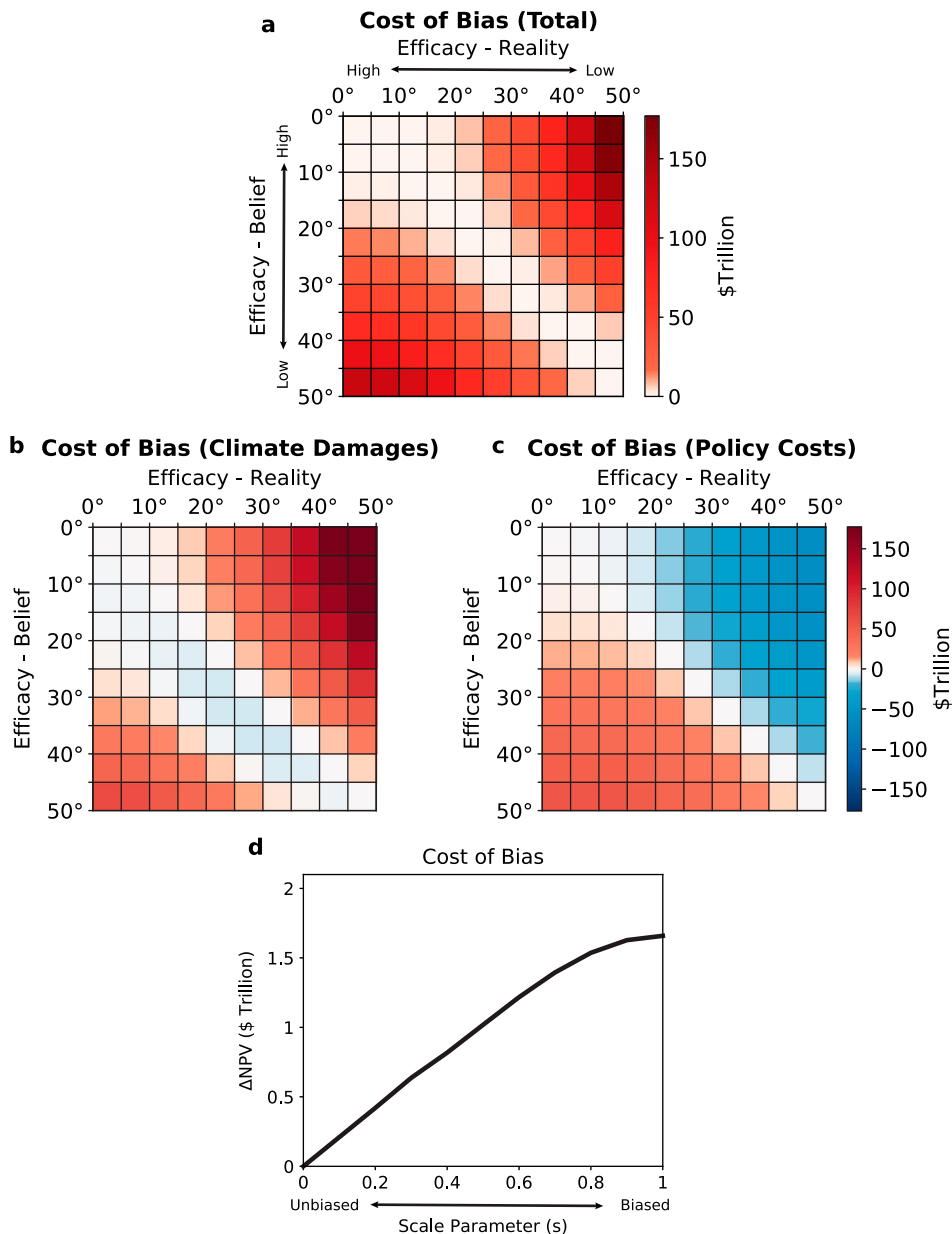


Figure 4. The cost of bias. (a), (b), and (c) show the cost of bias about SG efficacy without uncertainty. This is a quantitative version of the trade-off shown in Figure 1. Possible beliefs and realities are represented by point values. The cost of bias in SG efficacy measured by (a) reduced consumption, (b) increased climate damages, and (c) increased carbon policy costs. (d) displays the cost of bias with uncertainty. Bias is measured by parameter s which multiplicatively scales the decisionmaker's uncertain beliefs by $1-s$.

comparison, we find that the cost of a temporary global moratorium on SG deployment is about \$10 trillion, more than a factor of 5 larger than the cost of bias. The difference in costs is driven by higher temperatures under a temporary moratorium which cause higher climate damages later in the twenty-first century. This comparison suggests that while the costs of biased SG policy under uncertainty can be costly, it is a fraction of the costs of ruling out SG through a moratorium until perfect information arrives. The large costs of both bias and a global moratorium again indicate the importance of research to quickly resolve uncertainty about SG.

Conclusion

We estimate the value of information that eliminates uncertainty about SG efficacy is up to roughly \$10 trillion, comparable to the value of eliminating uncertainty about climate sensitivity. Recent calls for developing formal SG research programmes by the National Academy of Sciences (2021), and based on judgments by randomly selected climate experts in the US and China (Dai et al., 2021), suggest a NPV cost on the order of \$2 billion (see Supplementary Material). Our model has many potential order-unity uncertainties, so our estimates could be uncertain by a factor of three or more. But, even with this uncertainty, and even if research reduces uncertainty by only one-fifth, the value of reduced uncertainty about SG is about a thousand-fold larger than the cost of a SG research programme. Additionally, such a research programme could – if designed appropriately – manage risk-risk trade-offs and address other legal, ethical, and geopolitical uncertainties (Flegal et al., 2019; Grieger et al., 2019; Honegger, 2020).

Yet, our analysis cannot resolve disagreements about the wisdom of a SG research programme. First, our analysis narrowly focuses on a global, benevolent central planner and condenses global benefits of information over 200-years into a single value. We expect the benefits of research would hold in a more realistic world multi-agent framework; but, it is possible, though paradoxical, that reducing uncertainty may reduce welfare in a multi-agent framework (Heyen, 2019). This could occur if reducing uncertainty about SG reduces incentives to negotiate in a cooperative manner, or it could facilitate deployment by a free-driver.¹ With increased confidence about SG efficacy and risks, unilateral deployment may be more likely, but unilateral action could also be dangerous by triggering conflicts that outweigh any benefit of deployment.

Second, there are distinct but related concerns about moral hazard (McLaren & Corry, 2021), which, as we argue in the Introduction, can be captured as a bias towards deployment. Consistent with these concerns, we find that the costs of over-confidence in SG can be extreme – climate damages are very large if we rely on SG when it does not work. But we also find that under-confidence in SG, or ruling it out when it works, can have comparably high costs. These arise in the form of forgone opportunities to limit costly climate impacts. When considering a commitment to a large-scale SG research programme, decisionmakers should rightly consider the risks of SG's moral hazard problem or over-confidence in SG, but these must be weighed against the corresponding risk of under-confidence. Biased judgments are costly, whatever the direction of the bias. A coin has two sides.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data availability

All data generated and used in this analysis will be made publicly available upon publication.

ORCID

Anthony R. Harding  <http://orcid.org/0000-0002-6289-8253>

References

- Arino, Y., Akimoto, K., Sano, F., Homma, T., Oda, J., & Tomoda, T. (2016). Estimating option values of solar radiation management assuming that climate sensitivity is uncertain. *Proceedings of the National Academy of Sciences*, 113(21), 5886–5891. <https://doi.org/10.1073/pnas.1520795113>
- Belaïa, M., Moreno-Cruz, J., & Keith, D. (2021). Optimal climate policy in 3D: Mitigation, carbon removal, and solar geoengineering. *Climate Change Economics*, 12(3), 1–26. <https://doi.org/10.1142/S2010007821500081>

¹The term 'free-driver', coined by Weitzman (2015), refers to the challenge that because SG is inexpensive, it can be deployed unilaterally in the self-interest of a jurisdiction taking action, albeit with (potentially negative) welfare implications for other jurisdictions. This is the inverse of the free-rider problem posed by expensive emissions reductions.

- Cairns, R. C. (2014). Climate geoengineering: Issues of path-dependence and socio-technical lock-in. *WIREs Climate Change*, 5(5), 649–661. <https://doi.org/10.1002/wcc.296>
- Dai, Z., Burns, E. T., Irvine, P. J., Tingley, D. H., Xu, J., & Keith, D. W. (2021). Elicitation of US and Chinese expert judgments show consistent views on solar geoengineering. *Humanities and Social Sciences Communications*, 8(1), 1–9. <https://doi.org/10.1057/s41599-020-00694-6>
- Dai, Z., Weisenstein, D. K., & Keith, D. W. (2018). Tailoring meridional and seasonal radiative forcing by sulfate aerosol solar geoengineering. *Geophysical Research Letters*, 45(2), 1030–1039. <https://doi.org/10.1002/2017GL076472>
- Flegal, J. A., Hubert, A.-M., Morrow, D. R., & Moreno-Cruz, J. B. (2019). Solar geoengineering: Social science, legal, ethical, and economic frameworks. *Annual Review of Environment and Resources*, 44(1), 399–423. <https://doi.org/10.1146/annurev-environ-102017-030032>
- Goeschl, T., Heyen, D., & Moreno-Cruz, J. (2013). The intergenerational transfer of solar radiation management capabilities and atmospheric carbon stocks. *Environmental and Resource Economics*, 56(1), 85–104. <https://doi.org/10.1007/s10640-013-9647-x>
- Golub, A., Narita, D., & Schmidt, M. G. (2014). Uncertainty in integrated assessment models of climate change: Alternative analytical approaches. *Environmental Modeling & Assessment*, 19(2), 99–109. <https://doi.org/10.1007/s10666-013-9386-y>
- Grieger, K. D., Felgenhauer, T., Renn, O., Wiener, J., & Borsuk, M. (2019). Emerging risk governance for stratospheric aerosol injection as a climate management technology. *Environment Systems and Decisions*, 39(4), 371–382. <https://doi.org/10.1007/s10669-019-09730-6>
- Heyen, D. (2019). Risk Governance and the strategic role of uncertainty. In R. N. Stavins, & R. C. Stowe (Eds.) (pp. 91–94). Governance of the Deployment of Solar Geoengineering.
- Honegger, M. (2020). Addressing risks and trade-offs in governance. In M.-V. Florin (Ed.), *International governance issues on climate engineering: Information for policymakers* (pp. 72–90). Lausanne: EPFL International Risk Governance Center (IRGC).
- IPCC. (2021). *Climate Change 2021: The physical science basis. Contribution of working group I to the sixth assessment report of the intergovernmental panel on Climate Change* (V. Masson-Delmotte, P. Zhai, A. Pirani, S. L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M. I. Gomis, M. Huang, K. Keitzell, E. Lonnoy, J. B. R. Matthews, T. K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, & B. Zhou, Eds.). Cambridge University Press.
- Irvine, P., Emanuel, K., He, J., Horowitz, L. W., Vecchi, G., & Keith, D. (2019). Halving warming with idealized solar geoengineering moderates key climate hazards. *Nature Climate Change*, 1(9), 295–299. <https://doi.org/10.1038/s41558-019-0398-8>
- Irvine, P. J., & Keith, D. W. (2020). Halving warming with stratospheric aerosol geoengineering moderates policy-relevant climate hazards. *Environmental Research Letters*, 15(4), 044011. <https://doi.org/10.1088/1748-9326/ab76de>
- Keith, D. W. (2000). Geoengineering the climate: History and prospect. *Annual Review of Energy and the Environment*, 25(1), 245–284. <https://doi.org/10.1146/annurev.energy.25.1.245>
- Keith, D. W. (2021). Toward constructive disagreement about geoengineering. *Science*, 374(6569), 812–815. <https://doi.org/10.1126/science.abj1587>
- Knutti, R., Rugenstein, M. A. A., & Hegerl, G. C. (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10(10), 727–736. <https://doi.org/10.1038/ngeo3017>
- Kravitz, B., MacMartin, D. G., Mills, M. J., Richter, J. H., Tilmes, S., Lamarque, J.-F., Tribbia, J. J., & Vitt, F. (2017). First simulations of designing stratospheric sulfate aerosol geoengineering to meet Multiple simultaneous climate objectives. *Journal of Geophysical Research: Atmospheres*, 122(23), 12616–12634. <https://doi.org/10.1002/2017JD026874>
- Kravitz, B., MacMartin, D. G., Robock, A., Rasch, P. J., Ricke, K. L., Cole, J. N., Curry, C. L., Irvine, P. J., Ji, D., Keith, D. W., Kristjánsson, J. E., Moore, J. E., Muri, J. C., Singh, B., Tilmes, S., Watanabe, S., Yang, S., Yoon, J.-H. (2014). A multi-model assessment of regional climate disparities caused by solar geoengineering. *Environmental Research Letters*, 9(7), 074013. <https://doi.org/10.1088/1748-9326/9/7/074013>
- McKinnon, C. (2019). Sleepwalking into lock-in? Avoiding wrongs to future people in the governance of solar radiation management research. *Environmental Politics*, 28(3), 441–459. <https://doi.org/10.1080/09644016.2018.1450344>
- McLaren, D., & Corry, O. (2021). The politics and governance of research into solar geoengineering. *WIREs Climate Change*, 12(3), e707. <https://doi.org/10.1002/wcc.707>
- Moreno-Cruz, J. B., & Keith, D. W. (2013). Climate policy under uncertainty: A case for solar geoengineering. *Climatic Change*, 121(3), 431–444. <https://doi.org/10.1007/s10584-012-0487-4>
- Moreno-Cruz, J. B., Ricke, K. L., & Keith, D. W. (2012). A simple model to account for regional inequalities in the effectiveness of solar radiation management. *Climatic Change*, 110(3), 649–668. <https://doi.org/10.1007/s10584-011-0103-z>
- National Academies of Sciences. (2021). *Reflecting Sunlight: Recommendations for Solar Geoengineering Research and Research Governance*. <https://doi.org/10.17226/25762>
- Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7), 1518–1523. <https://doi.org/10.1073/pnas.1609244114>
- Parson, E. A. (2021). Geoengineering: Symmetric precaution. *Science*, 374(6569), 795–795. <https://doi.org/10.1126/science.abm8462>
- Quaas, M. F., Quaas, J., Rickels, W., & Boucher, O. (2017). Are there reasons against open-ended research into solar radiation management? A model of intergenerational decision-making under uncertainty. *Journal of Environmental Economics and Management*, 84, 1–17. <https://doi.org/10.1016/j.jeem.2017.02.002>
- Ricke, K. L., Rowlands, D. J., Ingram, W. J., Keith, D. W., & Granger Morgan, M. (2012). Effectiveness of stratospheric solar-radiation management as a function of climate sensitivity. *Nature Climate Change*, 2(2), 92–96. <https://doi.org/10.1038/nclimate1328>

- Sherwood, S. C., Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C., Hegerl, G., Klein, S. A., Marvel, K. D., Rohling, E. J., Watanabe, M., Andrews, T., Braconnot, P., Bretherton, C. S., Foster, G. L., Hausfather, Z., von der Heydt, A. S., Knutti, R., Mauritsen, T., ... Zelinke, M. D. (2020). An assessment of earth's climate sensitivity using Multiple lines of evidence. *Reviews of Geophysics*, 58(4), e2019RG000678. <https://doi.org/10.1029/2019RG000678>
- Smith, W., & Wagner, G. (2018). Stratospheric aerosol injection tactics and costs in the first 15 years of deployment. *Environmental Research Letters*, 13(12), 124001. <https://doi.org/10.1088/1748-9326/aae98d>
- Weitzman, M. L. (2015). A voting architecture for the governance of free-driver externalities, with application to geoengineering. *The Scandinavian Journal of Economics*, 117(4), 1049–1068. <https://doi.org/10.1111/sjoe.12120>