

CONTRIBUTED PAPER

# On the Utility of Research into Geoengineering Technologies for Risk-Avoidant Agents

Milana Kostić

University of Southern California, Los Angeles, CA, USA  
Email: [mkostic@usc.edu](mailto:mkostic@usc.edu)

(Received 25 January 2023; revised 24 February 2023; accepted 08 March 2023; first published online 11 April 2023)

## Abstract

In a recent paper Winsberg (2021) argued in favor of research into geoengineering by relying on Good's theorem, which states that conducting research maximizes one's expected utility. However, this result sometimes fails for risk-avoidant agents (Buchak 2010). Since risk avoidance captures some of the "precautionary" intuitions that critics of geoengineering share, it is important to see if geoengineering research would maximize one's utility if risk avoidance is taken into account. I show that under some conditions conducting geoengineering research would not maximize risk-weighted expected utility.

## 1. Geoengineering technologies

Geoengineering involves large-scale manipulations of the environment to counteract climate change. Advocates of implementing geoengineering strategies mostly motivate their position by claiming that mitigation strategies (various methods developed to reduce the release of carbon dioxide in the atmosphere, e.g., decreasing the fossil fuel dependence) are insufficient for reversing the effects of climate change (Winsberg 2021).

One of the main types of geoengineering technologies currently under discussion, stratospheric aerosol injection (SAI), consists in injecting very small sulfate particles into the stratosphere similar to those caused by volcanic eruptions, which could result in scattering a small fraction of sunlight back into space and thereby in global cooling. According to Winsberg, this strategy is cheap to implement, achievable with the currently existing technology, and effective (Winsberg 2021, 1113).

Oponents of implementing geoengineering technologies often invoke various forms of the precautionary principle in their arguments. The precautionary principle is a concept for guiding decision-making under risk or uncertainty. Roughly, it is argued that caution should be applied in using geoengineering strategies in order to minimize the (often unknown) negative environmental, geopolitical, economic, and

other consequences of implementing geoengineering strategies. There have been vigorous debates on the right way to operationalize the principle (and even on whether or not such operationalization is possible at all). The framework employed in this paper, Buchak's risk-weighted expected utility theory, has been presented in the literature as an alternative to the precautionary principle that can give similar policy recommendations, while avoiding the aforementioned drawbacks of the precautionary principle (Buchak 2019).

Critics argued not only against the deployment, but also against conducting (or for limiting) *research* into geoengineering (e.g., Gardiner 2010; McKinnon 2019). However, in a recent paper, Winsberg utilized a result in formal epistemology due to Good (1966) as a starting point in developing a "modest defense" of research into geoengineering technologies. In the following section I will present Good's result and the way Winsberg applies it to the particular case of research into geoengineering strategies.

## 2. "Research into geoengineering technologies maximizes expected utility"

I.J. Good (1966) famously proved that a rational agent should never refuse cost-free evidence (see also Blackwell 1953). In other words, Good proves that an agent maximizes their expected utility by gathering some cost-free evidence, conditionalizing on it, and then choosing the option that maximizes expected utility relative to the updated credences, rather than by choosing the option that maximizes expected utility based on their initial credences. This holds on the assumption that the agent would not choose the same option, regardless of what the incoming evidence shows, i.e., on the assumption that there is no dominant option.<sup>1</sup>

Winsberg argues that one can show, based on Good's result, that we would maximize expected utility if we were to make a decision about implementing geoengineering strategies upon conducting more research, barring the failure of the assumptions Good's theorem relies on. Winsberg focuses on the assumption that the credences used in evaluating the expected utility of conducting an experiment must be the same as the credences used when updating on the experimental result: "More crudely, if present-me thinks that future-me might misinterpret the evidence, then present-me might judge the belief-revision that would occur in the light of the misinterpretation of the new evidence to have negative expected utility" (Winsberg 2021, 1115).

Winsberg argues that a possibility that experimental results get misinterpreted exists when the prior probability of the proposition that geoengineering is beneficial is so low that "it is nearly impossible for research ever to, by my lights, significantly raise it" (Winsberg 2021, 1120). Then, this assumption, together with the assumption that "All scientific research has a non-trivial probability of being misinterpreted by decision makers in one direction or the other" (1120), would, according to Winsberg, entail that the research into SAI has a non-trivial probability of being taken to "offer stronger support for implementation than [one] would warrant and a nearly zero probability of doing the opposite" (1120). In other words, such an extremely low prior

<sup>1</sup> In the case that the agent would choose the same option no matter what, the expected utility of making a decision before and after conducting the experiment would be exactly the same.

**Table 1.** The Geoengineering Allais Paradox

	1%	10%	89%
Mitigation I	5	5	5
Geoengineering I	0	10	5
Mitigation II	5	5	0
Geoengineering II	0	10	0

probability in the proposition that geoengineering is beneficial, together with some plausible assumptions, would entail the failure of Good's theorem, according to Winsberg. However, Winsberg argues that having such an extremely low probability in the said proposition is implausible, partly because it will likely be difficult to counteract climate change using solely mitigation strategies. So the harmful effects of unmitigated climate change would likely surpass the harmful effects of implementing SAI, i.e., SAI would be overall beneficial (1122–4). In the sections that follow I will argue that conducting research into geoengineering technologies may not maximize one's utility even if no assumption is made about extremely low probability in the proposition that geoengineering is beneficial.

### 3. Risk-weighted expected utility theory

Limitations of the expected utility theory (EUT), used by Good and Winsberg, have been brought out by appealing to the Allais paradox. Suppose that you were given a choice between implementing mitigation and geoengineering strategies, where the corresponding gains are given in table 1.

It seems intuitive to prefer mitigation I over geoengineering I when the most likely outcome is a middling gain (e.g., corresponding to keeping the rise of global temperature below 2 °C above preindustrial levels) and geoengineering II over mitigation II when the most likely scenario is the catastrophic one (e.g., corresponding to a global temperature increase above 4 °C over the preindustrial level).<sup>2</sup> However, this prediction is not given by EUT: there is no assignment of utilities and probabilities in the EUT framework that predicts that mitigation has a higher expected utility in the first case and a lower expected utility in the second. Buchak's diagnosis of the Allais paradox is that the agents may be sensitive not only to the relevant probabilities and utilities, but also to the global properties of the gambles, i.e., it may matter what the best- and worst-case outcomes are as well as how the values of the gamble are spread out for the agents' decision-making. In the case of the Allais paradox, the reason why it may seem intuitive to opt for mitigation I over geoengineering I and not for mitigation II over geoengineering II could be that in the first case one stands to obtain the middling gain (5) for sure, i.e., without exposing oneself to the risk of obtaining the worst outcome (0). Avoiding the possibility of

<sup>2</sup> The intuition that geoengineering is acceptable under the condition that the most likely scenario we are facing is the catastrophic one (and that otherwise mitigation would be preferable), as mentioned, seems to be shared by Winsberg (2021, 1124) as well as other participants in the debate (see Gardiner et al. 2020, 13).

losing the guaranteed middling gain may matter more than the possibility of obtaining the highest possible outcome (10) that the geoengineering I option provides. As Buchak states, whenever the agents consider what may happen in the worst-case scenarios as more important to their choices than the outcome in the best-case scenarios, their behavior can be described as risk-avoidant (Buchak 2013, 1). Finally, the EUT framework can be modified to incorporate this additional component of people’s decision-making processes.

One way to calculate the expected utility of act *A* is as follows:<sup>3</sup>

$$EU(A) = U(O_1) + \sum_{2 \leq i \leq n} P(S_i)(U(O_2) - U(O_1)) + \sum_{3 \leq i \leq n} P(S_i)(U(O_3) - U(O_2)) + \dots + P(S_n)(U(O_n) - U(O_{n-1})).$$

That is, the expected utility is calculated by taking the utility of the worst-case option (O1) and then adding the difference between the second worst-case option (O2) and the worst option (O1), multiplied by the probability of obtaining at least that difference in utility, and so on for each pair of options *O<sub>i</sub>*, *O<sub>i+1</sub>*.

Buchak proposes the following way to calculate the risk-weighted expected utility of an act:

$$REU(A) = U(O_1) + r \left( \sum_{2 \leq i \leq n} P(S_i) \right) (U(O_2) - U(O_1)) + r \left( \sum_{3 \leq i \leq n} P(S_i) \right) (U(O_3) - U(O_2)) + \dots + r(P(S_n))(U(O_n) - U(O_{n-1})).$$

The difference from evaluating the expected utility of an act in EUT is in introducing the term *r*, which is used to capture the agent’s risk attitude. That is, in the case of risk-weighted expected utility theory (REUT), it is not only the case that the differences in utility are weighted by the likelihood of obtaining such difference. On the contrary, they are weighted by an additional factor: the risk of obtaining any improvement over the lowest-ranked option. If, for example,  $r(x) = x^2$ , then the weight given to the potential improvements over the lower-ranked options will be less than it is in the expected utility calculation. One can see the differences in values assigned by EUT and REUT to the geoengineering I option in figure 1.

The EU of the geoengineering I option calculated in the usual way ( $EU(\text{geoengineering I}) = 0.89 \times 5 + 0.1 \times 10$ ) is equal to the EU of geoengineering I calculated in the alternative way discussed in this section ( $EU(\text{geoengineering I}) = 0.99 \times 5 + 0.1(10 - 5)$ ), as one can see by comparing figures 1(a) and 1(b). One can also see the REU of the geoengineering I option in figure 1(c). Note that the additional gain matters much less for the REU than for the EU maximizer.

#### 4. Evaluating (risk-weighted) expected utility of research

In this section I will present Good and Buchak’s proposed procedure for evaluating the (risk-weighted) expected utility of research. Suppose that the agent is choosing

<sup>3</sup> This is equivalent to a more commonly used formula:  $EU(A) = \sum_{1 \leq i \leq n} P(S_i)U(O_i)$ .

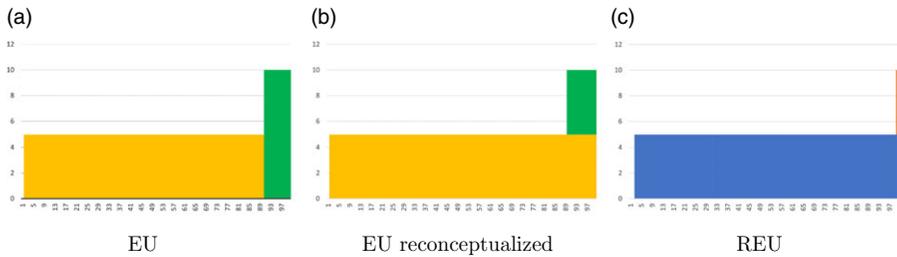


Figure 1. REU vs. EU of the geoengineering I option.

between two acts,  $A_1$  and  $A_2$ , and that the relevant states of the world are  $S_1$  and  $S_2$ . Suppose further that the agent is in a position to conduct an experiment with two possible experimental outcomes:  $E_1$  and  $E_2$ . Let  $\text{argmax}_{A_i}(\text{R})\text{EU}_{E_j}(A_i)$  be the act that would maximize the agent’s (R)EU once the agent’s credences are updated with the experimental result  $E_j$ . For instance, suppose that both  $\text{REU}_{E_1}(A_1) > \text{REU}_{E_1}(A_2)$  and  $\text{EU}_{E_1}(A_1) > \text{EU}_{E_1}(A_2)$ . Then,  $\text{argmax}_{A_i}(\text{R})\text{EU}_{E_1}(A_i)$  is simply  $A_1$ . Suppose also that  $(\text{R})\text{EU}_{E_2}(A_2) > (\text{R})\text{EU}_{E_2}(A_1)$ , so  $\text{argmax}_{A_i}(\text{R})\text{EU}_{E_2}(A_i)$  is  $A_2$ . Intuitively, this corresponds to the situation in which it is rational for an agent to choose to perform act  $A_1$  if they learn  $E_1$  and to perform act  $A_2$  upon learning  $E_2$ . The utilities of such an act can be represented in table 2.

Finally, the (risk-weighted) expected utility of an experiment is simply the (risk-weighted) expected utility of the act that amounts to opting for act  $A_1$  if one learns  $E_1$  and for  $A_2$  if one learns  $E_2$ .

Good shows that the EU of conducting an experiment is always higher than the EU of not performing an experiment, which is understood as opting for the act with the highest EU given the agent’s initial credences. However, under some conditions, the REU of an experiment may be lower than the REU of not conducting an experiment (Buchak 2010). In the next section I will discuss what those conditions are and whether they plausibly obtain in the case of research into geoengineering.

### 5. Research into geoengineering need not maximize risk-weighted expected utility

In the context under consideration in this paper, there are plausibly two acts whose risk-weighted expected utility we could evaluate: geoengineering and mitigation strategies. The former is often seen as “the risky strategy,” whereas the latter is seen as “the safe strategy.” Some potential risks of implementing SAI include aggravating ozone depletion, a globally or regionally changed precipitation regime, and regional temperature imbalances (Winsberg 2021, 1119–20). It is projected that such changes in precipitation patterns would disproportionately affect people in the global South (Robock 2015), which could further exacerbate global inequalities. Furthermore, deployment of SAI could worsen the negative effects of climate change. Namely, the injection of the sulfate aerosols would have to continue for decades. If, at any point, the injection of the sulfate aerosol were halted, a rapid temperature increase would ensue, and such a sudden halt would possibly have worse consequences than the unmitigated climate change (Winsberg 2021, 1120). However, it is also plausible that

**Table 2.** Utilities of conducting an experiment

$S_1 \wedge E_1$	$S_1 \wedge E_2$	$S_2 \wedge E_1$	$S_2 \wedge E_2$
$U(\text{argmax}_A (R)EU_{E_1}(A_i), S_1)$	$U(\text{argmax}_A (R)EU_{E_2}(A_i), S_1)$	$U(\text{argmax}_A (R)EU_{E_1}(A_i), S_2)$	$U(\text{argmax}_A (R)EU_{E_2}(A_i), S_2)$
$U(A_1, S_1)$	$U(A_2, S_1)$	$U(A_1, S_2)$	$U(A_2, S_2)$

**Table 3.** Utilities of implementing geoengineering and mitigation strategies

	$S_1$ geoengineering beneficial	$S_2$ geoengineering not beneficial
Geoengineering	10	0
Mitigation	5	5

implementing geoengineering could turn out well under some conditions. Proponents of geoengineering argue that this technology could help reach the desired target of preventing the global temperature from exceeding 2 °C from the preindustrial phase (Rogelj et al. 2019).

In summary, given such potential outcomes of implementing geoengineering, classifying geoengineering as a risky strategy seems to be plausible. That is, geoengineering seems to be a risky strategy in that it could be potentially very beneficial or potentially very harmful (depending on the different states of the world). Let us represent the two relevant states of the world as follows:  $S_1$ : geoengineering is beneficial (i.e., it does not lead to any of the previously mentioned harmful consequences);  $S_2$ : geoengineering is not beneficial (i.e., it does lead to some of the harmful consequences). The utilities of implementing geoengineering and mitigation can then plausibly be represented as in table 3.

The utility of implementing mitigation strategies is the same in both states of the world,  $S_1$  and  $S_2$ . That is, implementing mitigation would have the same utility regardless of whether geoengineering is beneficial or not. Furthermore, it is plausible to assume that implementing mitigation would be more beneficial than implementing geoengineering in the state of the world in which geoengineering is harmful but worse than implementing geoengineering in the state of the world in which geoengineering is beneficial.

Finally, suppose that we can conduct an experiment in which we can obtain evidence ( $E_1$ ) that confirms the hypothesis that geoengineering is beneficial, or evidence ( $E_2$ ) that disconfirms the hypothesis that geoengineering is beneficial. As described in the previous section, the procedure for evaluating the REU of research involves assigning initial utilities (in this case, values 10, 5, 0) to the states as in table 4 (assuming that the agent would implement geoengineering upon seeing experimental result  $E_1$  that confirms that geoengineering is beneficial and mitigation upon seeing the experimental result  $E_2$  that confirms that geoengineering is not beneficial).

Note that it is not required that the agent necessarily acts upon receiving evidence. What is assumed is that in this context, it would be rational for an agent to implement geoengineering whenever the evidence shows that geoengineering is beneficial (i.e.,

**Table 4.** Utility of conducting an experiment (research into geoengineering technologies)

	$S_1 \wedge E_1$	$S_1 \wedge E_2$	$S_2 \wedge E_1$	$S_2 \wedge E_2$
Experiment	10	5	0	5

$REU_{E_1}(A_1) > REU_{E_2}(A_2)$ ) and mitigation if the evidence shows that geoengineering is not beneficial (i.e.,  $REU_{E_2}(A_2) > REU_{E_2}(A_1)$ ). If mitigation is the action currently preferred, then the only additional assumption needed to ensure this is that  $P(S_1|E_1)^2 > M/H$  (where  $M$  refers to the middling and  $H$  to the highest utility value in the matrix). Roughly, the assumption needed is that the evidence we gather is strong enough to have the potential to convince us to choose an action that is different from the one we initially chose. For instance, we might be interested in determining whether geoengineering would worsen ozone depletion, in part because we think we would be better off implementing geoengineering on the condition that it would not worsen ozone depletion. This seems to be a plausible assumption in the context under discussion (otherwise, one may ask why we would be interested in collecting the evidence that would recommend the same action no matter what the experiment shows). However, if this were the case, then the (R)EU of such an experiment would be equal to the (R)EU of the initially preferred act. So, the agent would be indifferent between conducting such an experiment and implementing the act they opted for initially. The same holds for an agent who collects evidence and then refuses to use it but presumably opts for the act that was preferred given their initial credences. However, as Buchak (2010, 101) and Good (1966, 321) point out, such an agent would not be an (R)EU utility maximizer in general.

Finally, suppose that mitigation is currently preferred to geoengineering (i.e.,  $REU(\text{mitigation}) > REU(\text{geoengineering})$ ). This may be the dominant position in the literature. (For an overview of the literature, see Pamplany et al. (2020)). If mitigation is preferred, it can be shown that a risk-avoidant agent would not maximize their risk-weighted expected utility by making a decision upon conducting an experiment if the following holds:<sup>4</sup>

$$\frac{P(S_1 \wedge E_1)^2}{1 - (P(S_1 \wedge E_1) + P(E_2))^2} > \frac{M - L}{H - M},$$

where  $H$  refers to the highest utility value in the matrix (10 in the example under consideration),  $M$  to the middling value (5 in the example under consideration), and  $L$  to the lowest utility value in the matrix (0 in the example under consideration). That is, if

$$\frac{P(S_1 \wedge E_1)^2}{1 - (P(S_1 \wedge E_1) + P(E_2))^2} > 1$$

in the example under consideration, a risk-avoidant agent would maximize their risk-

<sup>4</sup> Since  $REU(\text{exp}) = L + (P(S_1, E_1) + P(S_1, E_2) + P(S_2, E_2))^2[M - L] + P(S_1, E_1)^2[H - M]$  and  $REU(\text{no exp}) = M$ , then  $REU(\text{no exp}) > REU(\text{exp})$  whenever  $M > L + (P(S_1, E_1) + P(S_1, E_2) + P(S_2, E_2))^2[M - L] + P(S_1, E_1)^2[H - M]$ , which reduces to the inequality  $\frac{P(S_1 \wedge E_1)^2}{1 - (P(S_1 \wedge E_1) + P(E_2))^2} > \frac{M - L}{H - M}$ .

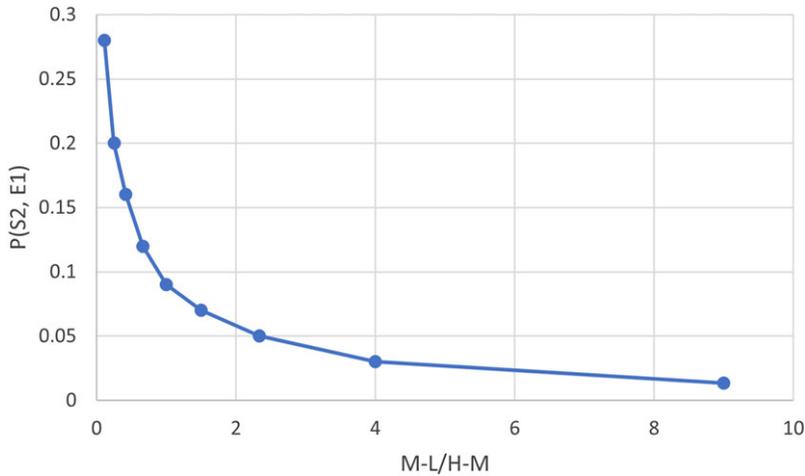
weighted expected utility by not conducting an experiment. Here's a probability distribution that would satisfy this condition:<sup>5</sup>  $P(S_1, E_1) = 0.4$ ,  $P(S_1, E_2) = 0.2$ ,  $P(S_2, E_1) = 0.1$ ,  $P(S_2, E_2) = 0.3$ . This result suffices to call Winsberg's argument into question. That is, pace Winsberg, a utility maximizer would not maximize their (risk-weighted) expected utility by conducting an experiment even if it is not the case that the prior probability of the proposition that geoengineering is beneficial is "extremely low." In the example probability distribution, it is more likely than not that geoengineering is beneficial ( $P(S_1) = 0.6$ ); nevertheless, the agent would prefer not to conduct research into geoengineering technologies.

Furthermore, assuming that according to the agent's current credences it is equally likely for the experimental procedure to show that geoengineering is beneficial as it is to show that geoengineering is not beneficial ( $P(E_1) = P(E_2) = 0.5$ ), as long as  $P(S_2 \wedge E_1) > 0.09$  (as long as the probability of the evidence showing that geoengineering is beneficial while in fact geoengineering is not beneficial is not negligible), the risk-avoidant agent would prefer not to conduct research into geoengineering. Intuitively, a risk-avoidant agent would prefer to implement the safe strategy rather than expose themselves to the risks of obtaining evidence that may misleadingly encourage them to implement geoengineering technologies and suffer some of the aforementioned harmful consequences.

How plausible is it that we may obtain misleadingly encouraging evidence when conducting research into geoengineering technologies? Critics of geoengineering research have warned about the possibility of institutional and technological "lock-in." Namely, they have argued that a "research program often creates a community of researchers that functions as an interest group promoting the development of the technology that they are investigating" (Jamieson 1996, 333). Hence, critics of geoengineering research may argue that obtaining misleadingly encouraging evidence in the institutional framework that may be created to support a geoengineering research program is a salient possibility.

Furthermore, in figure 2 one can see how the value of the minimum probability of the state in which geoengineering is not beneficial while the evidence shows that geoengineering is beneficial ( $P(S_2 \wedge E_1)$ ) that is needed to predict that the risk-avoidant agent would refuse to conduct research changes once we vary the payoff structure, i.e., once we vary the value  $(M - L)/(H - M)$ . One can see that the more we stand to gain by mitigating, i.e., the closer the utility of mitigating to the highest possible utility (i.e., the utility of geoengineering in the state of the world in which geoengineering is beneficial), the smaller the probability of misleading evidence would be needed for a risk-avoidant agent to rationally refuse to conduct research. Only in the case that the relevant utilities were  $L = 0$ ,  $M = 1$ ,  $H = 10$ , that is, only if we were to gain much more by implementing geoengineering than by implementing mitigation strategies, would a risk-avoidant agent prefer not to conduct research on the condition that obtaining misleading evidence is more likely than not ( $P(S_2 \wedge E_1) > 0.28$ ). Since geoengineering is mostly proposed as a supplementary strategy that could not address any other effects of the increased CO<sub>2</sub> emissions apart from the increasing global temperature, even by its fiercest proponents (e.g., Keith

<sup>5</sup> As well as the constraints that  $REU(\text{mitigation}) > REU(\text{geoengineering})$ ,  $REU_{E_1}(\text{geoengineering}) > REU_{E_1}(\text{mitigation})$ , and  $REU_{E_2}(\text{mitigation}) > REU_{E_2}(\text{geoengineering})$ .



**Figure 2.** Dependence of the minimum probability required for risk-avoidant agent to refuse research on the payoff structure.

2017), one may argue that the payoff structure in which we stand to gain much more by implementing geoengineering (in the state in which geoengineering is beneficial) over implementing mitigation does not adequately represent our choice situation.

Finally, even if we assumed that the agent initially prefers to implement geoengineering, they would not always maximize their REU by conducting an experiment. Buchak shows that the REU of conducting an experiment will be lower than the REU of not conducting an experiment (i.e., lower than the initial REU of implementing geoengineering) approximately whenever  $P(S_1)$  is high and  $P(S_1|E_2)$  is “low but still significant” (Buchak 2010, 100). This result indicates that even if implementing geoengineering strategies would maximize the agent’s risk-weighted expected utility (given their current credences), this does not entail that the risk-weighted expected utility would be maximized by conducting geoengineering research.

## 6. Conclusion

In this paper I called into question a recent argument due to Winsberg (2021), who developed “a modest defense” of research into geoengineering strategies by utilizing Good’s theorem. I pointed out that this result does not hold under some conditions for risk-avoidant agents. In particular, a risk-avoidant agent would choose not to conduct an experiment in the case that conducting an experiment could expose them to the risk of obtaining misleading evidence. The results presented in this paper suffice to call into question Winsberg’s attempt to argue in favor of research into geoengineering strategies by utilizing expected utility theory and Good’s theorem. As I have shown, risk-weighted expected utility theory might plausibly dictate not gathering, but avoiding the evidence, i.e., refusing to conduct the research, just like the advocates of imposing the moratorium on research into geoengineering strategies have proposed.

**Acknowledgments.** Thanks to Craig Callender, Jeff Russell, Eric Winsberg, and the audience at the PSA2022 meeting in Pittsburgh for their helpful comments. Financial support of The Institute for Practical Ethics at UC San Diego is gratefully acknowledged.

## References

- Blackwell, David. 1953. "Equivalent Comparisons of Experiments". *The Annals of Mathematical Statistics* 24 (2):265–72.
- Buchak, Lara. 2010. "Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering". *Philosophical Perspectives* 24 (1):85–120.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford: Oxford University Press.
- Buchak, Lara. 2019. "Weighing the Risks of Climate Change". *The Monist* 102 (1):66–83.
- Gardiner, Stephen. 2010. "Is 'Arming the Future' with Geoengineering Really the Lesser Evil? Some Doubts about the Ethics of Intentionally Manipulating the Climate System". In *Climate Ethics: Essential Readings*, edited by Stephen Gardiner, Simon Caney, Dale Jamieson, and Henry Shue, 284–312. Oxford: Oxford University Press.
- Gardiner, Stephen M., Catriona McKinnon, and Augustin Fragnière. 2020. "Introduction: Geoengineering, Political Legitimacy and Justice". In *The Ethics of "Geoengineering" the Global Climate*, edited by Stephen M. Gardiner, Catriona McKinnon, and Augustin Fragnière, 1–8. Abingdon: Routledge.
- Good, I. J. 1966. "On the Principle of Total Evidence". *British Journal for the Philosophy of Science* 17 (4): 319–21.
- Jamieson, Dale. 1996. "Ethics and Intentional Climate Change". *Climatic Change* 33:323–36.
- Keith, David W. 2017. "Toward a Responsible Solar Geoengineering Research Program". *Issues in Science and Technology* 33 (3):71–7.
- McKinnon, Catriona. 2019. "Sleepwalking into Lock-In? Avoiding Wrongs to Future People in the Governance of Solar Radiation Management Research". *Environmental Politics* 28 (3):441–59.
- Pamplany, Augustine, Bert Gordijn, and Patrick Brereton. 2020. "The Ethics of Geoengineering: A Literature Review". *Science and Engineering Ethics* 26 (6):3069–119.
- Robock, Alan. 2015. "Stratospheric Aerosol Geoengineering". *AIP Conference Proceedings* 1652 (1):183–97.
- Rogelj, Joeri, Piers M. Forster, Elmar Kriegler, Christopher J. Smith, and Roland Séférian. 2019. "Estimating and Tracking the Remaining Carbon Budget for Stringent Climate Targets". *Nature* 571 (7765):335–42.
- Winsberg, Eric. 2021. "A Modest Defense of Geoengineering Research: A Case Study in the Cost of Learning". *Philosophy and Technology* 34 (4):1109–34.

---

**Cite this article:** Kostić, Milana. 2023. "On the Utility of Research into Geoengineering Technologies for Risk-Avoidant Agents." *Philosophy of Science* 90 (5):1140–1149. <https://doi.org/10.1017/psa.2023.57>