OPEN ACCESS
Climate Change Economics, Vol. 12, No. 3 (2021) 2150008 (26 pages)

DOI: 10.1142/S2010007821500081

OPTIMAL CLIMATE POLICY IN 3D: MITIGATION, CARBON REMOVAL, AND SOLAR GEOENGINEERING

MARIIA BELAIA[*,‡], JUAN B. MORENO-CRUZ[†] and DAVID W. KEITH[*,‡,§]

[*]*The John A. Paulson School of Engineering and Applied Sciences Harvard University, Cambridge, MA 02140, USA*

[†]*School of Environment, Enterprise and Development University of Waterloo Waterloo, ON, Canada N2L 3G1*

[‡]*The John F. Kennedy School of Government Harvard University, Cambridge, MA 02140, USA*
[§]*david_keith@harvard.edu*

Received 15 December 2020
Revised 22 May 2021
Accepted 4 June 2021
Published 2 August 2021

We introduce solar geoengineering (SG) and carbon dioxide removal (CDR) into an integrated assessment model to analyze the trade-offs between mitigation, SG, and CDR. We propose a novel empirical parameterization of SG that disentangles its efficacy, calibrated with climate model results, from its direct impacts. We use a simple parameterization of CDR that decouples it from the scale of baseline emissions. We find that (a) SG optimally delays mitigation and lowers the use of CDR, which is distinct from moral hazard; (b) SG is deployed prior to CDR while CDR drives the phasing out of SG in the far future; (c) SG deployment in the short term is relatively independent of discounting and of the long-term trade-off between SG and CDR over time; (d) small amounts of SG sharply reduce the cost of meeting a 2°C target and the costs of climate change, even with a conservative calibration for the efficacy of SG.

Keywords: Climate change; solar geoengineering; carbon dioxide removal; climate policy; integrated assessment.

JEL Code: Q54

1. Introduction

Emissions must be cut to limit climate impacts, but emissions cuts cannot reduce the impact of past emissions. Solar geoengineering (SG) and carbon dioxide removal (CDR) both allow reduction in the climate impacts of historical emissions so both may

[§]Corresponding author.

play important roles in managing the impacts of climate change. Yet, they are very different instruments. SG is a direct intervention in radiative forcing (RF) that acts fast partly breaking the link between greenhouse gas (GHG) concentrations and climate change. CDR breaks the link between emissions and concentrations. Both SG and CDR reduce the inertia between action and outcome in the climate system by enabling humans to reduce the footprint of past emissions on future climates. SG has significant uncertainties and risks but low cost while CDR has high cost, lower risks, and larger inertia.

Optimal economic models can explore the intertemporal trade-offs between emissions mitigation, SG, and CDR. These trade-offs can inform policy even though the quantitative results are strongly determined by the ways in which the impacts and efficacy of the technologies are parametrized. In contrast to the potential importance and growing visibility of SG and CDR, the integrated assessment modeling (IAM) literature seldom considers them.

In this paper, we introduce SG and CDR in an otherwise standard model of climate and economy. Introducing SG into a canonical IAM is not straightforward. First, SG produces different spatial and spectral patterns of RF than do GHGs, so even if SG is used to offset the global average net RF from GHGs, it cannot reverse the effects of GHGs on the climate. Second, SG introduces new forms of damages, which combined with the partial compensation for GHGs climate impacts, means the economic response to SG differs from the response to GHGs. Moreover, most IAMs use temperature as the exclusive driver of climate damages, but temperature lags RF. Naively introducing SG without accounting for these important differences would result in SG dominating any optimal policy. To address these concerns, we develop a systematic and parsimonious approach to representing SG's regional disparities that may be generally applicable to a range of IAMs.

While the effects of CDR and emissions reductions on the climate are the same, introducing CDR is challenging for economic reasons. Namely, the technologies are not yet developed and assumptions regarding costs and scalability are still hard to calibrate. We calibrate this model using the current scientific and economic literature and use it to calculate optimal climate policy under different scenarios and modeling assumptions.

We find that SG and CDR are part of the optimal portfolio that results in lower temperatures and overall lower costs of climate change, relative to a mitigation-only policy. SG peaks when CDR is deployed at a scale sufficient to achieve net-negative emissions. Not only do SG and CDR reduce overall costs if combined with emissions reductions, but we identify synergies between SG and CDR that highlight the importance of considering them jointly as part of an optimal policy. First, the use of CDR allows for the phasing out of SG. Second, SG reduces the need for large amounts of CDR and thus reduces the peak and costs of achieving negative emissions.

Our work contributes to the nascent but increasing literature on the economics of climate geoengineering. Few prior IAM studies look at an interplay between mitigation and SG (Bahn *et al.*, 2015; Bickel and Agrawal, 2013; Emmerling and Tavoni, 2018a; Goes *et al.*, 2011; Gramstad and Tjøtta, 2010; Heutel *et al.*, 2018). A larger strand of IAM studies explores the role of CDR in conjunction with mitigation, with disproportionate focus on bioenergy with carbon capture and storage (BECCS) (van Vuuren *et al.*, 2017; Vinca *et al.*, 2018). This is the first IAM that analyzes the temporal trade-offs between mitigation, SG and CDR. We present the first version of the "napkin diagram", which illustrates interrelationship between mitigation, SG and CDR and has been widely used in discussions of SG policy for about a decade, derived from an optimization model. Only three papers address mitigation, CDR, and SG in an IAM, but in a very limited fashion. Bickel and Lane (2009) estimate the benefits of specific CDR and SG pathways using exogenous trajectories separately and without optimization. Emmerling and Tavoni (2018a, 2018b) explore the relationship between mitigation and SG under uncertainty about SG effectiveness and under alternative global cooperation settings, respectively. Both papers use the WITCH model, which has a detailed energy sector representation and includes one form of CDR–BECCS, albeit they present it as part of mitigation. In other words, both papers neither disentangle the contributions of BECCS and traditional mitigation nor present model results for the case without BECCS. In those papers the trade-offs are between SG and mitigation and not in the three dimensions as we explore in our paper.

The rest of the paper proceeds as follows. First, we introduce our modeling approach in detail, placing emphasis on our treatment of SG and the assumptions regarding the costs of CDR, and present our calibration strategy. Next, we present the results of our simulation and discuss the implications for economic policy. Finally, we draw conclusions and elaborate on the limitations of our approach.

2. Modeling Approach

We build on the Dynamic Integrated model of Climate and Economy (DICE), version 2016-R2 (Nordhaus, 2018). DICE aggregates the world into one unit, coupling an economic growth model, a single time-varying supply curve for emissions control, and a climate represented by a fixed climate sensitivity, a two-box ocean model, and a three-box carbon cycle. It is most commonly used to compute globally optimal emission reduction pathways. DICE has been widely adapted for studies of climate policy including geoengineering (Bickel and Lane, 2009; Heutel *et al.*, 2016). We provide a rational for our treatment of SG and CDR in the remainder of this section.

2.1. *Solar geoengineering*

Global climate impacts are a sum of local impacts that depend on local changes in climatic parameters such as soil moisture or peak temperature. The calibration in DICE assumes that changes in local variables are correlated with changes in global mean temperature (ΔT). More specifically, DICE assumes that climate damages are proportional to the squared deviation of temperature from its preindustrial value. SG creates a RF that can offset some of the positive RF from CO_2. The magnitude of the SG RF is a policy choice as are choices about the spatial and spectral distribution of

SG's RF. SG could eliminate time-averaged ΔT if feedback is used to adjust the RF based on observed temperature (MacMartin *et al.*, 2014). But no adjustment of SG RF can exactly mirror the RF from CO_2. SG is not anti-CO_2. This means that if SG is used to reduce temperature, it will yield a different distribution of local climate than if that same global average temperature was achieved by a reduction in CO_2. Naively introducing SG RF into an IAM without considering this difference grossly exaggerates its effectiveness. In addition, SG does not address direct geochemical impacts of CO_2, such as ocean acidification, CO_2 fertilization, and changes in plant transpiration.

The essential challenge of incorporating SG into IAMs is finding a parsimonious way to represent the heterogeneity in local climate variables. The challenge goes beyond heterogeneity of changes in local climates, as IAMs with regional climate damage functions still use impact estimates that assume a correlation between climate variables, such as temperature and precipitation, that are calibrated using models of CO_2-driven climate change. SG changes the relationship between local variables, such as temperature and precipitation, introducing biases when using the established climate impact correlations to analyze SG impacts. Yet, econometric estimates of macroeconomic impact of climate changes are largely driven by temperature, not precipitation (Harding *et al.*, 2020). This means that temperature is a reasonable proxy for regional economic impacts of SG-driven climate changes.

We represent the effect of SG in two functions: efficacy and impacts. *Efficacy* captures the limited ability of SG to reduce aggregate climate damages due to the heterogeneity of climate response to SG. *Impacts* captures the human and environmental side-effects and costs associated with producing the SG's RF.

Efficacy. We first consider equilibrium climate response and assume that (a) local climate response is linear in global RF caused by GHG or some form of SG, and that (b) climate impacts are proportional to the global weighted sum of the local squared deviation of climate variables from their reference. Climate model response is, of course, not perfectly linear, the essence of this assumption is that the errors in linear extrapolation are small compared to other errors in estimating climate response (Gillett *et al.*, 2004; Moreno-Cruz *et al.*, 2012). With these assumptions, the climate response to GHG or SG may be represented by vectors in a multidimensional vector space consisting of all relevant climate variables (Fig. 1), a framing we adopt from Moreno-Cruz *et al.* (2012).

We can now show that the magnitude of the climate impact is proportional to the squared length of the residual vector formed by the vector addition of SG and GHGs. In DICE, damages are quadratic in the deviation of global mean temperature from its preindustrial level, ΔT, given by $D^{\mathrm{DICE}}(\Delta T) = \alpha \cdot \Delta T^2$, where α is a proportionality constant, % GWP/K^2. Assuming that climate response is in equilibrium, temperature change is given by $\Delta T = S \cdot C$ where S is climate sensitivity, K/(W m^{-2}), and C is GHG RF, W m^{-2}. Accordingly, damage costs can be represented as follows:

$$D(C) = \alpha(S \cdot C)^2.$$

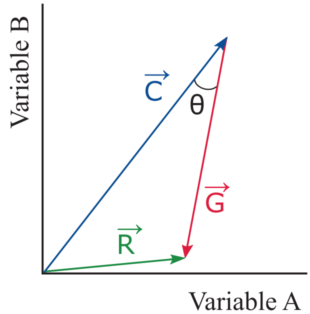

Figure 1. The RVM. Axes show deviations of two climate variables (e.g., peak temperature or average runoff in some region) from a reference climate. The blue vector shows changes in A and B as a linear function of RF from GHGs C, while the red vector shows how the variables change with RF from SG, G. The residual changes, R, are shown as a green vector and θ is the efficacy angle.

To introduce regional impacts, we can express aggregate damages as the sum of its regional components. This is similar to the regionalized version of DICE, known as RICE, in which regional damage costs are proportional to the squared deviation in global mean temperature from the preindustrial level with region-specific coefficients of proportionality (Nordhaus, 2010).

In the present study, we assume that change in regional climate variables X_i is proportional to changes in sensitivity to GHG RF

$$X_i = S_i \cdot C.$$

Representing regional variables in a vector form we have $\mathbf{X} = C \cdot \mathbf{S}$.

Global damage costs are given by a weighted sum of the square deviation in local climate variables:

$$D(C) = \sum_i \alpha_i (X_i)^2 = \sum_i \alpha_i (S_i \cdot C)^2.$$

Equivalently, in a vector form

$$D(C) = \alpha \cdot \mathbf{S} \cdot \mathbf{S} \cdot C^2.$$

Next, we assume that climate response to SG is linear too: $X_i = K_i \cdot G$. Damages of GHG and SG RF are given by a weighted sum of the squared deviation in the local climate variables:

$$D(C, G) = \sum_i \alpha_i (S_i \cdot C - K_i \cdot G)^2.$$

Again, in vector form it reads

$$\begin{aligned} D(C,G) &= \alpha(\mathbf{S}\cdot C - \mathbf{K}\cdot G)\cdot(\mathbf{S}\cdot C - \mathbf{K}\cdot G) \\ &= \alpha(\mathbf{S}\cdot\mathbf{S}\cdot C^2 + \mathbf{K}\cdot\mathbf{K}\cdot G^2 - 2\cdot\mathbf{S}\cdot\mathbf{K}\cdot G\cdot C), \end{aligned}$$

where, for simplicity and without loss of generality, we factor the differences in weights α_i into $\mathbf{S}$ and $\mathbf{G}$. Let us denote g by

$$g = \frac{|K|\cdot G}{|S|\cdot C}.$$

With this expression, and assuming $\frac{|K|}{|S|} = 1$, we can write damages from GHG and SG as follows:

$$D(C,g) = \alpha(S\cdot C)^2(1 + g^2 - 2g\cos(\theta)).$$

Change in global mean temperature with SG is given by $\Delta T = S\cdot C - K\cdot G$. Replacing this expression above delivers

$$E(g) = 1 + g^2 - 2\cdot g\cdot\cos(\theta), \tag{1}$$

where E is the efficacy of SG in reducing climate damages, defined as the magnitude of climate impacts with SG divided by the magnitude of climate impacts without SG.

With this residual vector model (RVM), the efficacy of SG in reducing climate impacts depends only on g, the length of SG vector normalized by the length of GHG vector, and θ the angle between the two vectors in Fig. 1. This model has two immediate consequences. First, if the amount of SG is chosen to offset all RF from GHGs — roughly equivalent to offsetting all ΔT — it will not minimize all impacts. Second, the minimum residual climate damage is $1 - \cos^2(\theta)$ when g is set to $\cos(\theta)$, the value that minimizes impacts.

Next, we need to accommodate the nonequilibrium case with time-varying GHG and SG. In a more sophisticated model, this might be done using linear extrapolations from climate model runs scaled to an ocean model following the example of MAGIC or FAIR (den Elzen and Lucas, 2005; Meinshausen *et al.*, 2011). For DICE, we adopt a simple approach in which the SG RF enters the temperature equation as a perfect offset to the GHG RF to compute time-varying ΔT. We then apply the equilibrium efficacy from Eq. (1) to an estimate of what ΔT would have been without SG, an estimate derived as the product of the instantaneous GHG RF and the climate sensitivity. This yields our overall damage parametrization:

$$D(\Delta T, C, G) = \alpha\cdot(\Delta T + s\cdot G)^2\cdot E\left(\frac{G}{C}\right) + I\cdot G, \tag{2}$$

where C for "carbon" and G are GHG and SG RF, respectively; $\alpha = 0.23\%$ GWP/K^2 and $s = 0.84$ K/(W m^{-2}) are the climate impact coefficient and climate feedback parameter both using values from DICE, and $I = 0.1\%$ GWP/W m^{-2} is a linear coefficient capturing the impacts and costs of SG as discussed below. When no SG is used ($G = 0$), there is no change in climate damages, $E(g) = 1$, and we arrive at the original DICE damage function.

Many technologies could be used to implement SG with the most studied being stratospheric aerosols, marine cloud brightening, and cirrus thinning (National Academies of Sciences, 2021). Deployment of any of these involves choosing a goal for the intended spatial and temporal distribution of RF. Consistent with our use of DICE — a global average model — we analyze the case in which SG is deployed with the goal of producing a globally uniform distribution of RF, an objective which can most plausibly be achieved by stratospheric aerosols (National Academies, 2021).

Several climate model analyses provide a basis for estimating the efficacy of uniform SG as an angle in the RVM:

- Moreno-Cruz *et al.* (2012) calculated angles using the HadCM3L model with SG implemented as zonally uniform variations in stratospheric optical depth finding angles of 4° and 23° for annual average temperature (T) and precipitation (P) when the 22 Giorgi regions (Giorgi, 2006) were weighted by economic output.
- Kravitz *et al.* (2014) evaluated 12 climate models from GeoMIP under equilibrium $4 \times CO_2$ with SG represented by an adjustment to the solar constant evaluated over 22 regions. For T they computed a Pareto optimal value of g (the largest value without making any region worse off), yielding an angle of 18°–32° with a median of 24.5° across the 12 models. For P the Pareto optimal $g = 0$ because P moves away from its preindustrial value for $g > 0$. The angle corresponding to the average of the optimal g's across regions is 53°.
- Irvine *et al.* (2019) and Irvine and Keith (2020) analyzed change over the IPCC Intergovernmental Panel on Climate Change (IPCC) Special Report on Extremes (SREX) regions focusing on four important climate hazards: changes in T, precipitation minus evaporation (PE), extreme temperature (Tx), and extreme precipitation (Px). They focus on PE because with P being reduced along with T, it is precipitation–evaporation which determines water availability. The 2019 study (Irvine *et al.*, 2019) uses the Geophysical Fluid Dynamics Laboratory (GFDL) high resolution tropical cyclone-resolving model with SG represented by an adjustment of the solar constant, while the 2020 study (Irvine and Keith, 2020) uses an ensemble run of the National Center for Atmospheric Research (NCAR) high resolution stratospheric model with explicit representation of sulfate injection stratospheric Aerosol Geoengineering Large Ensemble (GLENS), (Tilmes *et al.*, 2018). The resulting angles were 5.6°, 25°, 6.3°, and 23° for T, PE, Tx, and Px in the GFDL model, while for the NCAR model the associated values were 5.9°, 30°, 13°, and 15°.

Climate impacts will depend on hazards such as T, PE, Tx, and Px along with other outputs, such as crop yields and sea-level rise. We know of no simple and robust way to estimate impacts based on a combination of these values.

There is growing econometric evidence that T and Tx are among the most important determinants of climate impacts (Burke *et al.*, 2015; Dell *et al.*, 2014; Harding *et al.*, 2020) so one could defend choosing a SG efficacy angle of less than 10°. However, we believe that at this initial stage of research on the potential role of SG as part of climate policy, focusing on the worst case SG efficacy value is crucial. In this spirit, in the present paper we focus on the angle of 30°, which is the largest (worse) value for any variable in SREX studies. We refer to this parametrization as *conservative*. We also present the model results for the 10° efficacy case, which we refer to as *empirical*.

Impacts. The second effect of SG are the costs and side-effects of producing the RF. There are three components: (a) the direct cost of deployment such as the delivery vehicles and materials, (b) the cost of observing systems used to monitor deployment, and (c) the side-effects such as ozone loss or air pollution.

The best understood method for delivering SO_2 to the atmosphere is to use aircraft that could inject aerosols at altitudes between 18 and 25 km. An upper bound of roughly 4 Tg/year of SO_2 would be required for $1\,\mathrm{W\,m^{-2}}$ of RF. The estimated costs for delivering this mass flux are in the range \$1.5–8 billion per year. These are annualized costs including aircraft procurement and operations (McClellan *et al.*, 2012; Smith and Wagner, 2018). Delivery of more advanced materials that might have less side-effects might be significantly more expensive. In addition, there would be the cost of maintaining redundant systems to manage the risk interruption. We adopt a conservative estimate of \$25 billion per year for $1\,\mathrm{W\,m^{-2}}$ of RF (Moriyama *et al.*, 2017).

For monitoring costs, we assume a global cost equal to the entire US Global Change Research budget of \$1.9 billion per year (USGCRP, 2021).

Side-effects of stratospheric aerosols include stratospheric ozone depletion (Pitari *et al.*, 2014), increase in ground-level pollutants from descending stratospheric aerosol (Eastham *et al.*, 2018), impact on crop yields (Proctor *et al.*, 2018), and acid rain (Kravitz *et al.*, 2009). Many of these have not been quantified, are small in direct economic impact (stratospheric ozone loss and acid rain) or have uncertain sign (crop yields). To estimate the costs of stratospheric ozone loss and other health impacts, we use the results from an air quality study by Eastham *et al.* (2018), which found 11,000 additional deaths due to the injected aerosol mass descending to the surface at SG RF of $2\,\mathrm{W\,m^{-2}}$ combined with current US value of statistical life to yield \$55 billion per year for $1\,\mathrm{W\,m^{-2}}$. Note that this is a high estimate as the study by Eastham *et al.* (2018) found that the net of the direct and indirect effects of SG was a decrease in net mortality.

To account for other impacts, we round these combined estimates up to 0.1% of GWP (\$130 billion) for $1\,\mathrm{W\,m^{-2}}$ and assume that it would remain proportional to GWP since much of the cost is a monetized mortality impact. By rounding up the values, we emphasize the approximate nature of these calculations.

Simple Quadratic. Models of SG are uncertain, and almost all analyses stop at physical hazards rather than impacts. No analyses yet provides a systematic link between SG and comprehensive assessment of impacts. As a tie to prior literature, we adopt a simple alternative specification for the impacts introduced in Moreno-Cruz and Keith (2013) in which SG perfectly compensates for temperature but causes damages that are quadratic in the amount of SG RF. Impacts are calibrated so that when SG is used to offset all temperature change from $2 \times CO_2$, its damages are as large as the temperature-driven damages due to $2 \times CO_2$. This specification has the virtue of simplicity, and it is conservative in that its estimate of SG's damages is much larger than empirical estimates.

2.2. *Carbon dioxide removal*

From minor adjustments to agricultural practices that increase soil carbon to electrochemically accelerated weathering an extraordinary range of methods could remove carbon dioxide from the atmosphere (House *et al.*, 2007; National Academies of Sciences, Engineering, and Medicine, 2018). For inclusion into simple IAMs it's useful to divide these methods into two broad categories based on the lifetimes of the stored carbon.

Methods in which the carbon is stored as chemically reduced carbon in the biosphere have lifespans of years to centuries, sufficiently short that removal cannot be considered permanent over the duration of long-term climate policy. Salient examples include CDR by changes to forestry or agricultural practices. The lifetime of carbon stored in these systems depends on future policy choices and on climate because warming increases the flux back to the atmosphere. Rather than CDR it might be better to call these methods "carbon banking" or "delayed emissions".

Methods in which the carbon is stored either as CO_2 in geologic storage (BECCS or direct air capture (DAC)–CCS) or in chemically bound form as dissolved salts in the ocean (accelerated weathering) have lifetimes exceeding a millennium. Whatever their costs or environmental impacts, it's reasonable to count the carbon as permanently removed.

The short-term methods are generally less expensive. For methods that involve economically induced alterations to current farming or forestry practices, the marginal costs start at zero if one assumes that existing markets are economically efficient. But the potential for safe and economical scale-up is limited (National Academies of Sciences, Engineering, and Medicine, 2018). Technologies such as DAC–CCS, BECCS, and accelerated weathering are far less developed so it's difficult to estimate costs, but it seems likely that near-term marginal costs exceed 100 \$/t-CO_2. For at least some of these methods, there do not seem to be physical bounds on their maximum annual capacity.

In DICE-2016, CDR is assumed to be unavailable until the year 2150. And, when it does enter, its cost is tied to the scale of the economy because CDR costs are defined in

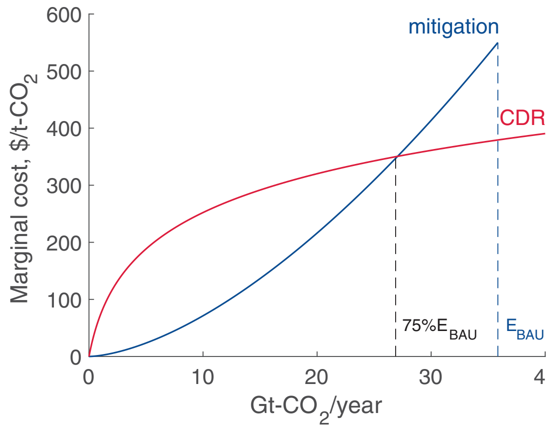

Figure 2. Marginal mitigation and CDR cost curves. Mitigation is limited by BAU emissions (blue dashed line), CDR can go up to maximum of 40 Gt-CO_2/year. Our CDR curve is calibrated based on the assumption that it crosses the mitigation curve at 75% of current emissions so that CDR is cheaper for the last 25% of emissions. Values shown are for 2015, both costs decline with time as specified for mitigation in DICE.

relation to business-as-usual (BAU) emissions. In contrast to DICE, we assume a supply curve that starts at zero to capture short-term low-cost methods for which markets are already active, and which is strongly concave to reflect the large-capacity high-cost methods. In other words, our CDR cost curve can be thought of as a sum of two curves, representing two general forms of CDR: (i) low-cost approaches such as soil carbon management with a limited environmental and economic scalability; and (ii) high-cost approaches that have marginal costs that are nearly scale-independent such as addition of alkalinity to the open ocean. The fact that CDR combines low-cost scale-limited technologies with high-cost roughly scale-independent technologies means that the marginal cost curve must be concave. We chose a logarithmic function for marginal CDR costs

$$d \cdot \ln(R + 1), \tag{3}$$

where d is the cost scaling parameter, and R is the CDR uptake rate in Gt-CO_2 per year.

We calibrate d with one simple parameter: the fraction of current emissions for which marginal cost of mitigation exceeds marginal cost of CDR (Fig. 2). That is, the point when CDR becomes a dominant approach to counteracting emissions.

In rough accord with recent supply curve studies (Davis *et al.*, 2018; Goldman Sachs, 2019), we set this parameter at 25%. With this value, marginal mitigation and CDR cost curves intersect when mitigation reaches 75% of emissions under control. For our reference year, 2015, this value corresponds to 26.9 Gt-CO_2, at which mitigation costs are 350 \$/t-CO_2. Thus, parameter d is given by

$$d = \frac{350}{\ln(26.9 + 1)},$$

from where we arrive at $d = 105$ \$/t-CO_2. This implies that large-scale CDR will be more cost-effective than mitigation-only when the economy begins to reduce emissions in the most stubborn sectors. We assume that future costs decline from these 2015 values at the same rate as DICE uses to reduce mitigation costs (i.e., rate of backstop cost decline).

Finally, we impose an upper limit $R_{\max}$ on the amount of annual carbon removal R. We set $R_{\max}$ at 40 Gt-CO_2, which is about the level of baseline emissions in the reference year. At the end of the following section we demonstrate the robustness of our results to this value.

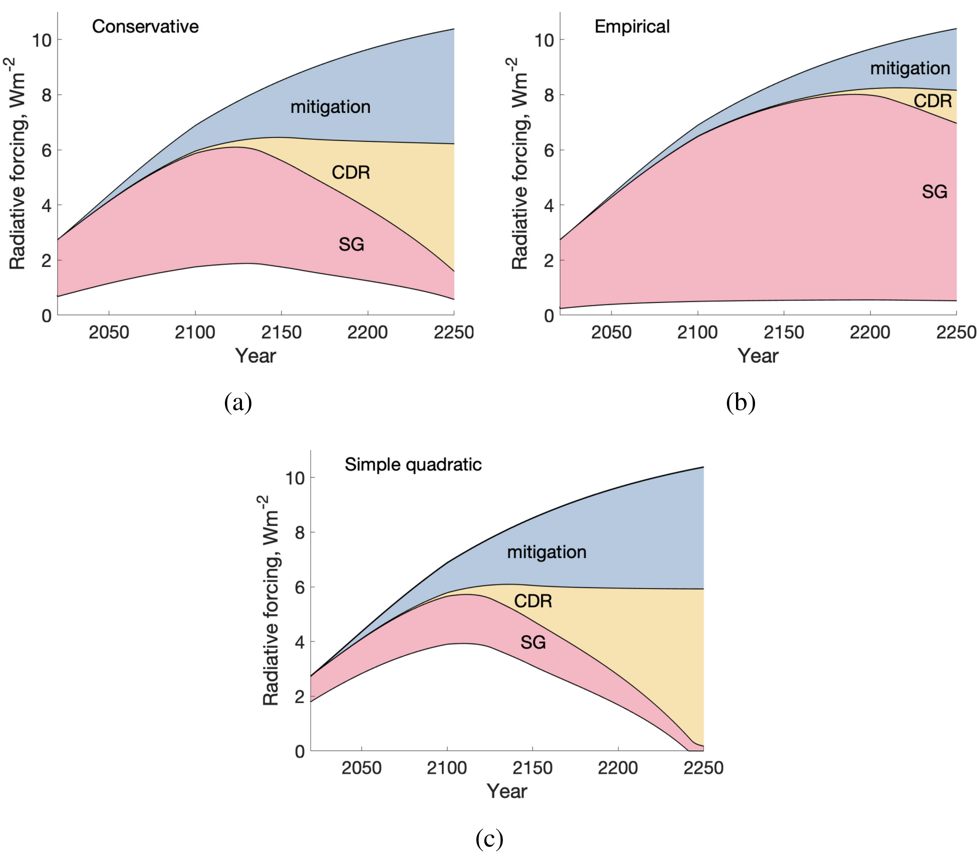

Figure 3. Optimal climate policy portfolio with three instruments. Contributions of policy instruments to a reduction in BAU RF (uppermost black line). Results for conservative, 30° (a), and empirical, 10° (b), parametrization of SG efficacy in RVM, and for simple quadratic SG modeling approach (c).

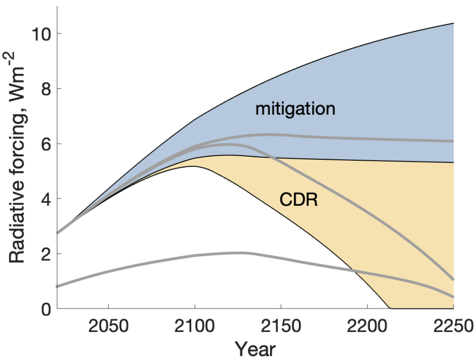

(a)

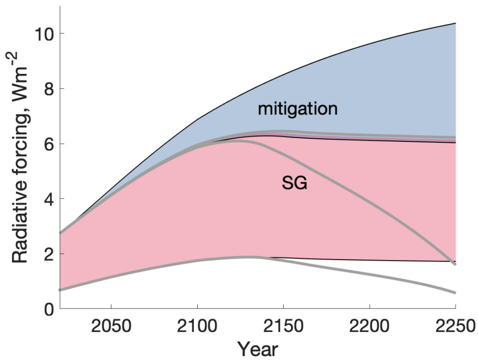

(b)

Figure 4. Policy portfolio with only two instruments. Contributions of policy instruments to a reduction in BAU RF under portfolios limited to mitigation and CDR (a) and mitigation and SG (b). In both panels gray lines show reference results for complete three-instrument portfolio from Fig. 3(a).

3. Results

The contributions of different policy instruments are shown in Fig. 3. We use RF as an indicator in our figures because instantaneous contributions to reducing RF are separable whereas contributions to reduced temperature are hard to disentangle given the ocean thermal inertia. We show temperature response and economic damages in Appendix B (Figs. B.1 and B.2) and in Table 1. Panels (a) and (b) show the results for the RVM-based approach to SG representation for conservative (our base case) and empirical SG efficacy parametrizations, respectively. Panel (c) shows results for the case where SG enters in the form of an additive quadratic term in the damage function.

SG peaks when RF from CO_2 reaches its maximum which is roughly coincident with peak temperatures and net-zero emissions (Figs. 3 and B.1). SG deployment begins immediately and ramps up well before peak RF, whereas large-scale use of CDR occurs after peak RF.

Deployment of SG precedes large-scale CDR in the optimal policy under all the alternative assumptions we examined. This is in sharp contrast with current climate policy debates which are beginning to contemplate large-scale near-term use of CDR, but which generally ignore SG or implicitly assume that its deployment would occur after a ramp-up in CDR.

We analyze the contributions of each of the instruments to the overall policy in Fig. 4. Removing SG increases deployment of mitigation and CDR but does little to change their overall trajectory producing a peak RF of 5.2 $\mathrm{W\,m^{-2}}$ in 2099 compared to the baseline (Fig. 3(a)) of 2.02 $\mathrm{W\,m^{-2}}$ in 2125. Removing CDR means that use of SG never decreases, we do not see a reduction in RF, and concentrations never return to preindustrial. We find that SG and CDR supplement mitigation in sharply distinct ways. CDR reduces concentrations of carbon in the atmosphere, reversing the impacts

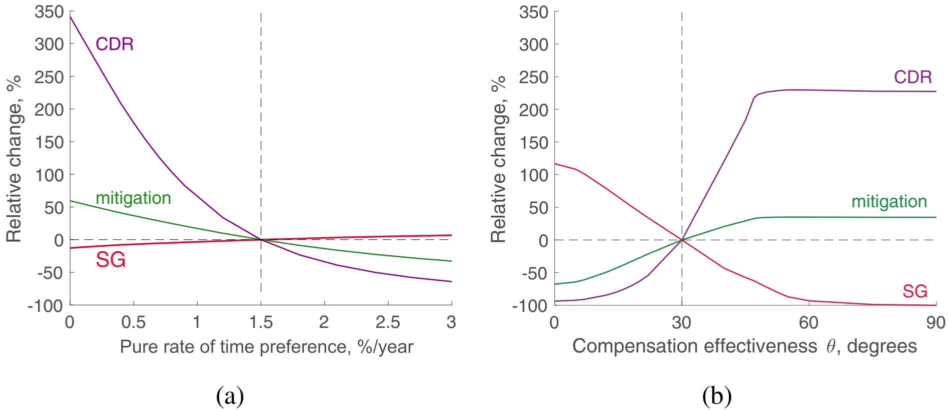

Figure 5. Impact of changes in discounting or SG efficacy. Change in use of instruments with welfare discount rate (a), or with the SG efficacy angle. Vertical dashed lines show default values as a reference. Changes in Mitigation and CDR are relative to their cumulative use over the next century, 2020–2120; while for SG, changes are relative to its peak level. Our choice of metrics arises from the distinct nature of the two quantities: atmospheric carbon is a *stock*, while SG RF is a *flow*.

of historical emissions, and it could eventually restore the climate to close to preindustrial (Fig. 4(a)). SG is less costly than mitigation or CDR and it acts faster, but it cannot eliminate the climate risks from historical emissions (Fig. 4(b)).

An important trade-off in climate policy is between permanent and costly options like mitigation and CDR, and immediate and low costs strategies like SG. This trade-off is governed by the discount rate. We explore the change in cumulative (2020–2120) mitigation and CDR along with peak SG level as we change the pure rate of time preference in Fig. 5(a) (see also Figs. B.3 and B.4). Because the costs of SG are low

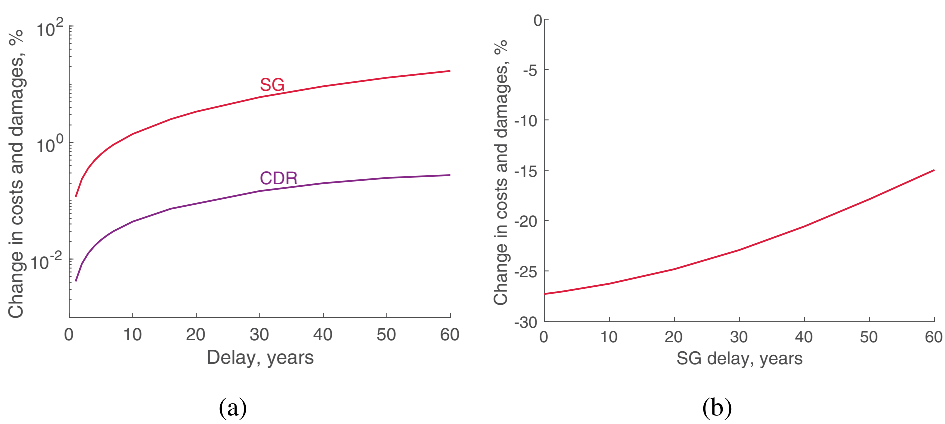

Figure 6. Cost of delaying implementation of CDR or SG in terms of percentage change in NPV of climate policy costs and climate damages over the period 2020–2120 discounted at the rate of 3% per year. Change relative to the no-delay policy (a), and relative to the no-SG policy (b).

and its effects are almost instant, SG responds far less to changes in time preference than do mitigation and CDR. Decisions about SG primarily involve the current generation balancing its costs and impacts against its efficacy in reducing current risks, while trade-offs between SG, mitigation, and CDR over time is a second-order effect.

Adding CDR and SG to the policy mix reduces the overall cost of climate policy compared to mitigation alone. This is unsurprising because adding options cannot increase costs of an optimal solution. The surprising and important result is that CDR and SG contribute across different dimensions, as is evident from the fact that reducing the discount rate sharply increases deployment of CDR with a relatively small impacts on SG, while changing the efficacy of SG makes sharp changes in the deployment of SG with relatively small impacts on CDR and mitigation.

Next, we explore the costs of delaying deployment of SG or CDR. As we see in Fig. 6(a), the cost of delaying SG is roughly 50 times larger than cost of delaying CDR. A delay in CDR brings an immediate benefit — a decrease in policy spending — yet at a price of larger damages, which lag behind the policy implementation (Appendix B, Fig. B.5). These damages are, in part, compensated for by a slightly elevated level of SG. Yet, the delay in CDR results in a net loss. In the absence of SG, delayed CDR would call for an increase in mitigation, because while net emissions are positive, mitigation and CDR can be substituted with a small impact on total costs (Fig. 2). The cost of delaying CDR is comparatively low because it is not deployed at a large scale until mitigation has reduced most of the emissions. A delay in SG brings immediate costs in the form of climate damages as well as increased spending on mitigation and CDR (Appendix B, Fig. B.5). With this, earlier deployment facilitates larger benefits of SG, as measured in the change in net present value (NPV) of policy costs and climate damages relative to the no-SG case (Fig. 6(b)).

Temperature targets are salient in climate policy. With SG temperature remains below the Paris Agreement 1.5°C stretch target in our baseline case (Fig. 3(a)). There are many reasons to limit SG deployment that are not captured by our cost-benefit analysis, including strategic interactions between states, and use of a precautionary framework that adopts a risk-averse weighting to evaluating new technologies. Figure 7 explores the consequences of imposing a 2°C temperature target with a constraint on the maximum deployment of SG which is varied parametrically. Allowing SG to exceed 2.9 $\mathrm{W\,m^{-2}}$, optimal temperature is below 2°C threshold.

While CDR is costly, its reduction in concentrations is permanent, whereas SG is less costly and inherently fleeting. Independent of costs and discounting we find that CDR always brings concentrations back towards preindustrial levels, with a timing that depends on discounting, the effectiveness of SG, and the limits on SG (Figs. 5 and 7). Without CDR damages remain constant at the same level throughout the simulation. With CDR, damages towards zero by the end of the simulation. Because SG is immediate, it is virtually unresponsive to discounting (Fig. 5(a) and Table 1). And, because CDR is permanent, SG is not significantly replacing CDR, but is supplementing its contribution.

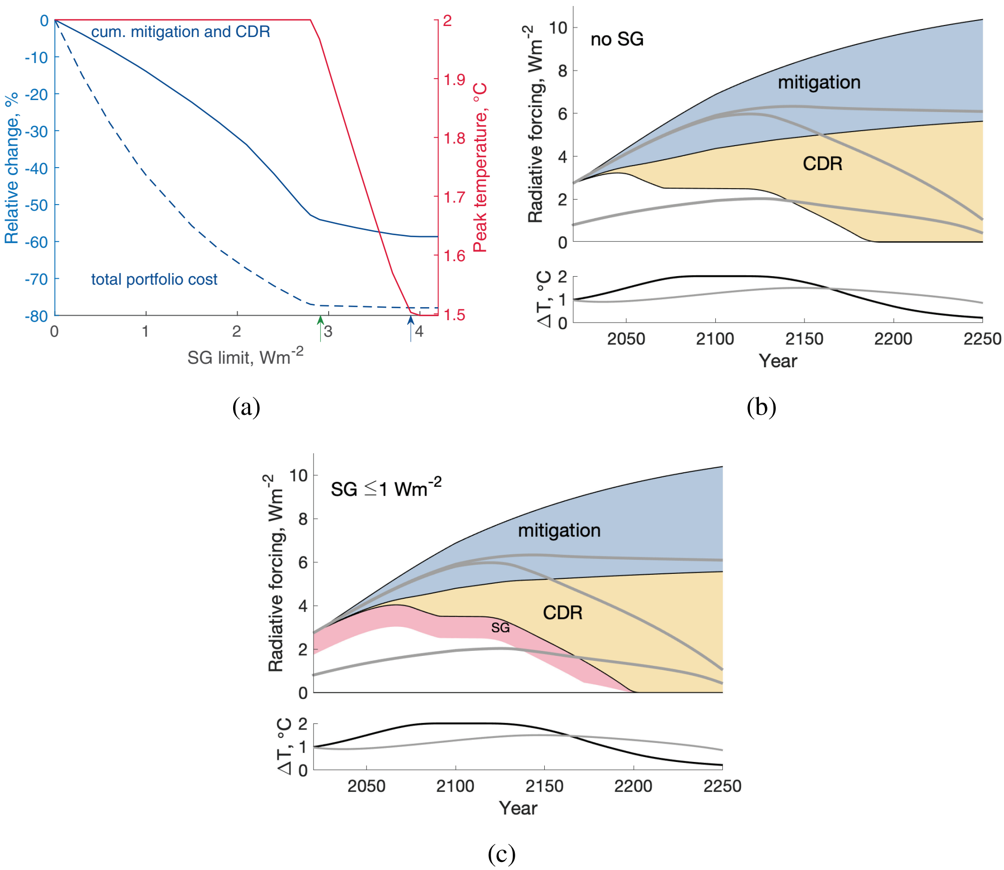

Figure 7. Meeting the 2°C target with varying amount of SG. Trade-offs between SG and cumulative mitigation by CDR and mitigation in meeting the 2°C target (a) as a function of an imposed limit on the amount of SG than can be deployed. Temperatures stay below 2°C when SG constraint $> 2.9\,\mathrm{W\,m^{-2}}$ (green arrow). Unconstrained peak SG is $3.95\,\mathrm{W\,m^{-2}}$ (blue arrow). Change in cumulative mitigation and CDR along with NPV portfolio cost (discounted at 3% per year) are computed over 2020–2120. Use of instruments and temperature response are shown for (b) no SG and (c) SG constrained to $< 1.0\,\mathrm{W\,m^{-2}}$. In both panels gray lines show reference results for complete three-instrument portfolio from Fig. 3(a). Note that small amounts of SG produce a sharp drop in total cost but a small reduction in use of CDR and mitigation.

Finally, we address the question of sensitivity of the results to the upper limit on annual CDR deployment. Lowering the limit would imply stronger longer-term optimal mitigation and slower SG phase-out as shown in Appendix B (Fig. B.8) for the limit of 20 Gt-CO_2/year. In the case of no limit, CDR is deployed at a larger scale, substituting the last 15% of mitigation (Appendix B, Fig. B.8). Worth noting, these last 15% of emissions decline over time even in the absence of a policy due to the autonomous energy efficiency increase exogenously specified in DICE. As a result, the need for CDR in this case declines over time too. In a nutshell, the limit on CDR has no implications for the near-term policy. It does not alter the year of reaching net-zero

Table 1. Summary of the model results for the optimal policy portfolio that includes mitigation, CDR, and SG.

DICE parametrization	SG modeling approach[a]	Peak emissions		Net-zero emissions	Peak SG		Peak temperature		Peak annual damage costs[e]		Peak annual mitigation and CDR costs		Change in damages relative to mitigation-only[f]	Change in policy costs relative to mitigation-only[f]
		$Gt\text{-}CO_2$	Year	Year	$W\ m^{-2}$	Year	°C	Year	% GWP	Year	% GWP	Year	%	%
Default[b]	No SG[c]	40	2050	2100	n/a	n/a	3.6	2123	3.0	2123	1.5	2110	−2	+17
	RVM, conservative	45	2062	2126[e]	4	2116	1.5	2147	1.9	2124	0.94	2140	−17	−56
	RVM, conservative Double SG impacts	43	2059	2116	3.3	2106	1.8	2133	2.2	2119	0.94	2140	0	−40
	RVM, empirical	59	2089	2199	7.5	2191	1	2019	1.1	2191	0.3	2211	−58	−95
	Simple quadratic	44	2058	2117	1.8	2116	2.8	2135	2.3	2131	1.1	2132	−22	−40.1
The 2° target[d]	Simple quadratic	41	2050	2105	2.7	2100	2	2091	2.1	2105	1.2	2125	n/a[h]	n/a[h]
Zero welfare discount rate	RVM, conservative	36	2047	2094[g]	3.4	2086	1.1	2125	1.2	2100	1.6	2111	+17	−70
	Simple quadratic	34	2042	2086	1.7	2080	1.9	2109	1.3	2106	1.9	2102	+5	−60
Double welfare discount rate	RVM, conservative	50	2078	2154	4.3	2152	1.9	2175	2.5	2158	0.57	2170	−19	−30
	Simple quadratic	50	2071	2150	1.8	2152	3.7	2166	3.6	2169	0.6	2165	−24	−30

[a] RVM stands for residual vector model, conservative and empirical SG efficacy parametrizations are 30° and 10°, respectively.
[b] Under default we refer to standard parametrization of DICE-2016.
[c] "No SG" scenario refers to climate policy portfolio that includes mitigation and CDR, but not SG.
[d] The results for the 2° target for residual vector modeling approach are not shown because under the corresponding cost-benefit policy temperature is 2°C.
[e] Recall: total damage costs include climate damages and SG impacts. Peak annual costs are estimated as largest percentage value.
[f] Change in total damage costs (climate damages and SG impacts) and policy costs (mitigation and CDR) are estimated as percentage change in NPV discounted at 3% under three-instruments policy portfolio relative to mitigation-only policy for the corresponding parametrization.
[g] Figure B.6 illustrates net emissions over time.
[h] The 2° target is not feasible in the case of mitigation-only.

emissions, nor the year or level of SG peak. Limiting the annual CDR implies sooner zero emissions and slower SG phase-out.

4. Discussion and Conclusions

In this paper we introduce SG and CDR into DICE, a canonical IAM. We find that the optimal use of SG is in parallel with mitigation and before CDR is deployed at a scale to reach net-zero emissions. This is in contrast to what we perceive to be a widely accepted order of preference, where SG comes after mitigation and CDR. To this end, this order of preference was driving conversations around SG towards the idea that SG should be considered (if at all) after mitigation is exhausted and CDR is used at a large scale. We not only demonstrate that this order should be different if we were to maximize global welfare, but also that the large-scale CDR is what drives SG phase-out after net-zero emissions have been reached.

Overall, we find that SG substantially decreases the overall cost of responding to climate change. It reduces peak economic damages, for example, from 3% of GWP for mitigation and CDR alone to 1.9% when SG is allowed. This specific result depends on the SG parameterization (Fig. 5(b)), but given that our choice of a 30° efficacy angle is larger than the angle found in any climate model's estimate of the efficacy angle for temperature or extreme temperature, and given the growing evidence that temperature is a major determinant of climate impacts (Burke *et al.*, 2015; Harding *et al.*, 2020), our model may underestimate the benefits of SG. Similarly, even if the impacts term is doubled, SG still substantially reduces overall damages (Table 1). These results suggest that SG confers substantial benefits if decisions about its use are made by a global and benevolent decision-maker. With this, SG has the potential to make a strong climate policy more feasible, filling the gap between the desired climate policy outcome, as indicated in the Paris Agreement, and economically feasible options for CDR and mitigation. As the international environmental agreements' literature indicates, reducing the costs of action increases the number of countries willing to contribute on climate coalitions, thus it is possible for SG to reduce free riding on mitigation. Yet, it is possible that harms from SG's misuse might overwhelm its benefits in a world with multiple nonbenevolent decision-makers (Ricke *et al.*, 2013). In addition, a more realistic representation of SG deployment would include a limit on its ramp-up. Thus, despite the fast-acting nature of SG, temperature decrease would be slower than in the results presented here.

The benefits in baseline model parametrization call for SG that peaks at $4\,\mathrm{W\,m^{-2}}$. It might be better to use less. We explore explicit limits to SG in the context of the 2°C target and find that even small amounts of SG provide sharp reductions in the cost of meeting the 2°C target with comparatively small reductions in mitigation (Fig. 7).

Limiting SG to $1\,\mathrm{W\,m^{-2}}$ cuts the total policy costs by 43% while reducing deployment of mitigation and CDR by only 14% and reducing the duration with temperatures at 2°C by 20 years.

Moral hazard is among the most important political concerns about SG. Allowing SG or CDR does postpone mitigation. Compared to a mitigation-only case, cumulative 2020–2050 mitigation is reduced 26% and 4% by allowing SG and CDR. In our model this delay is a benefit — it increases welfare. This does not resolve concerns about moral hazard because there is simply no room for moral hazard in a single-actor optimization model. Rather, this suggests that the moral hazard of SG or CDR should be defined as reductions in mitigation relative to the optimal policy that includes SG and CDR.

The long-term policy is dominated by CDR, which enables SG phase-out. As with most applications of DICE, we are not interested in the very long run, but there is an outcome worth highlighting. In our model the world always returns to preindustrial conditions. This result arises from the nature of CDR as a technology that permanently removes carbon from the atmospheric stock. Once net emissions are zero, there is some rate of CDR at which its marginal cost is equal to the integral to infinity of the discounted marginal damages from SG and climate change. CDR will continue going down its marginal cost curve towards zero when the last ton of postindustrial CO_2 is removed. Worth noting, an incentive to return to a preindustrial climate comes from the shape of the damage function native to DICE, where damage costs are proportional to squared deviation of temperature from its preindustrial level. In reality, SG may be designed to meet one or multiple objectives beyond global welfare maximization. For example, to preserve biodiversity, address specific elements of climate system (e.g., Arctic), facilitate achieving sustainable development goals, reduce global inequality, etc.

In our approach to parametrization of SG efficacy and impacts, we take a conservative aiming to underestimate the efficacy and overestimate the impacts. More specifically, with temperature as a dominant driver of regional economic impacts of climate changes, an efficacy of less than 10° can be justified based on modeling studies cited in Sec. 2.1. Yet, we chose to focus on a larger (worse) value of 30° for the following reason. While the best-guess estimate case is often of interest for practical reasons, the performance of a policy instrument in the worst case can be seen as an important initial step in accessing its functionality. For the reference, we offer results for the 10° case in Fig. 3 and in Appendix B (Fig. B.7). Larger efficacy of SG calls for more SG, thereby reducing climate damage costs even further.

Our modeling approach comes with the usual limitations attached to IAMs in general and DICE in particular. Our centralized, benevolent decision-maker is a fiction, and we ignore large uncertainties associated with climate interventions. Some of climate damages may be irreversible, which would make an early action in every policy dimension even more important, while rendering long-term CDR less relevant. Nevertheless, we offer a first calibration of SG that accounts for regional damages in an IAM, and together with conservative modeling of CDR, we find an optimal policy that limits costs to less than 2% GWP/year, uses SG as a peak temperature shaving strategy while using CDR to compensate for past emissions in the long run. We hope this framework provides some insights for policy makers, climate negotiators, and practitioners as they develop policies to manage climate change in the real world.

Acknowledgments

We are grateful to Gernot Wagner for collaborating on the initial conceptual design of this study, and to Joseph Aldy, Elizabeth Burns, Zhen Dai, Tyler Felgenhauer, Khara Grieger, Klaus Keller, Wilfried Rickels, Dale Rothman, and Massimo Tavoni for their insightful comments. JBM-C acknowledges support by the National Science Foundation grant NSF #1948154.

Appendix A. Contributions of Policy Instruments to a Reduction in BAU RF

In Sec. 3, we present contributions of mitigation, CDR, and SG to a reduction in BAU RF. While the contribution of SG RF is straightforward, we needed to disentangle contributions of mitigation and CDR and specify BAU RF. For this, we track BAU emissions along with emissions avoided and associated carbon concentration pathways. To account for the carbon cycle dynamics, we add a separate three-equations carbon cycle models for BAU emissions and for mitigation. Using this approach, we estimate RF dynamics that is specific to mitigation-only and no-controls CO_2.

Appendix B. Figures

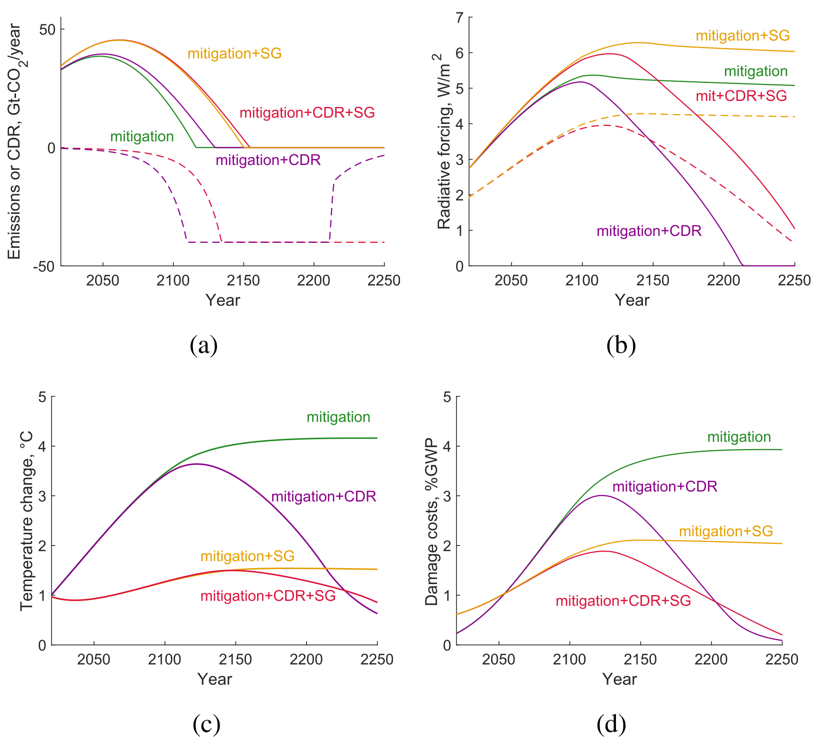

Figure B.1. Optimal policy and its implications under alternative policy portfolio compositions using RVM approach to SG.

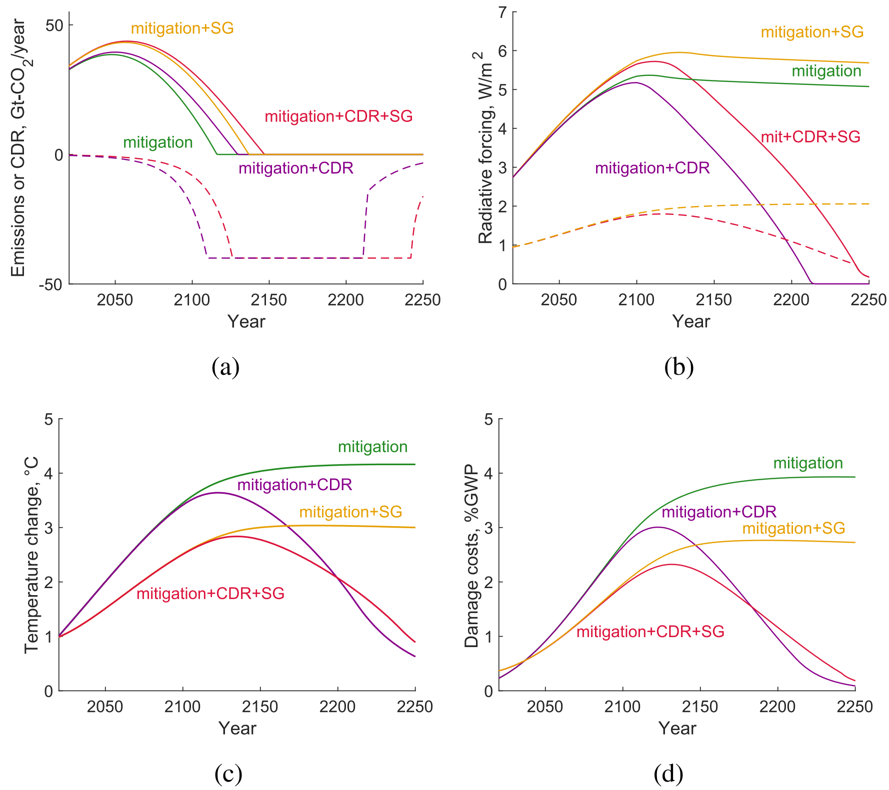

Figure B.2. Optimal policy and its implications under alternative policy portfolio compositions using simple quadratic SG modeling approach.

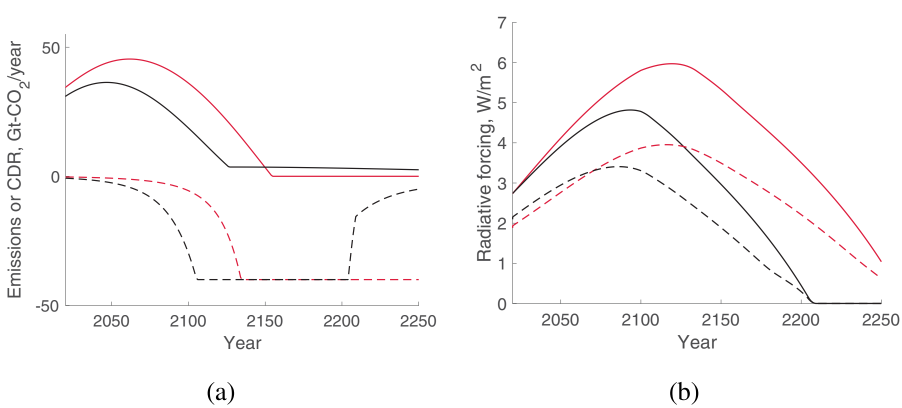

Figure B.3. The optimal three-instruments policy portfolio under zero pure rate of time preference (in black) and under default rate of 1.5% per year (in red). Annual industrial emissions and CDR (panel (a)), GHG and SG RF (panel (b)), associated temperature change from preindustrial level (panel (c)), and total damage costs (panel (d)).

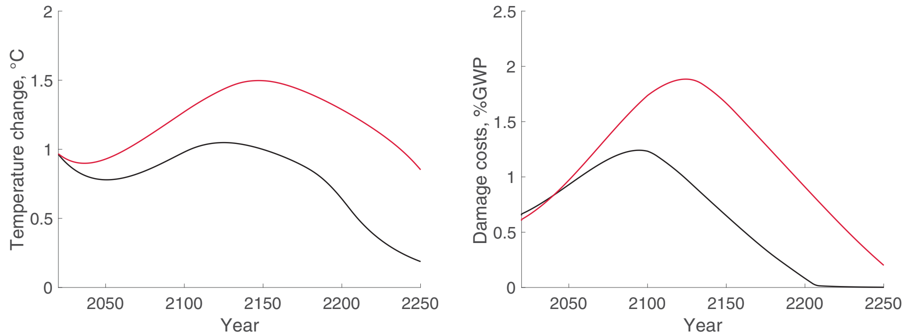

(c) (d)

Figure B.3. (*Continued*)

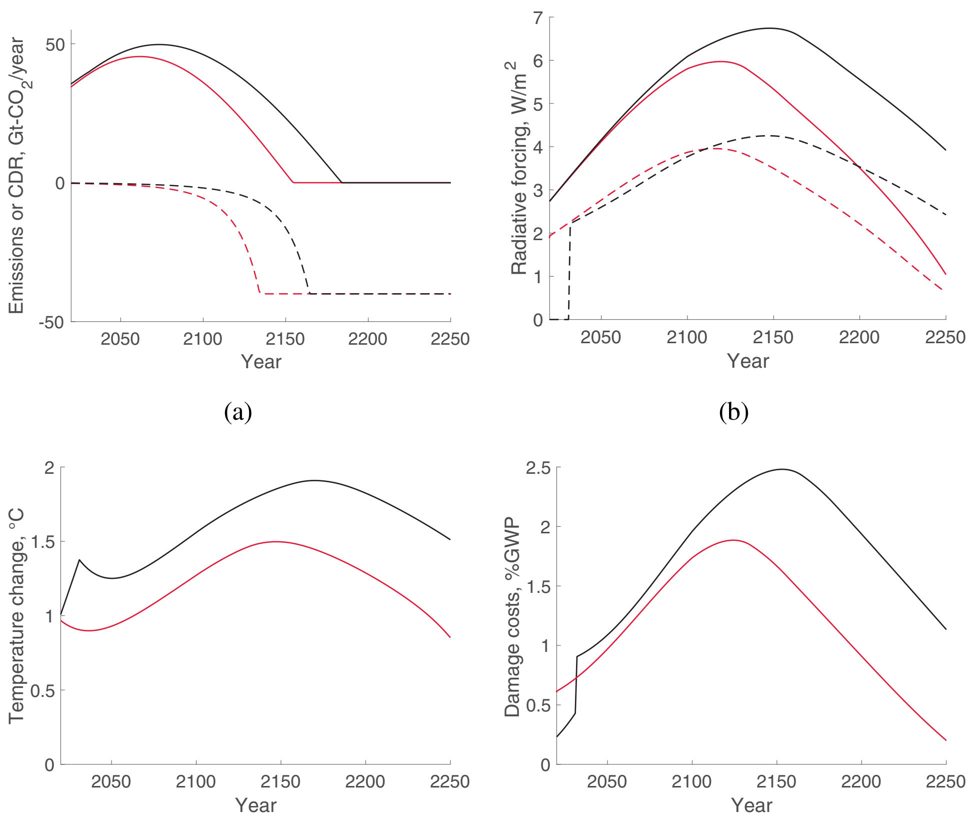

(a) (b)

(c) (d)

Figure B.4. The optimal three-instruments policy portfolio under double pure rate of time preference (in black) and under default rate of 1.5% per year (in red).

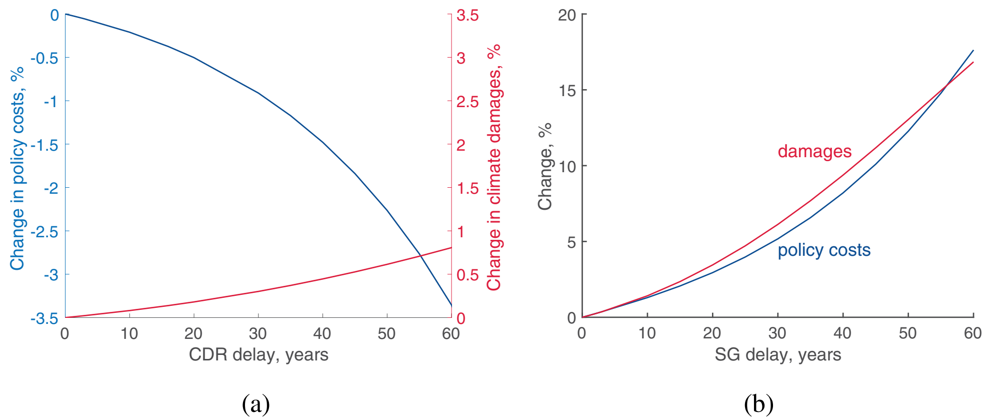

Figure B.5. Percentage change in policy costs and climate damages following a delay in CDR (a) or SG (b) relative to the no-delay case. Change in NPV estimated over the period 2020–2120, using the discount rate of 3% per year.

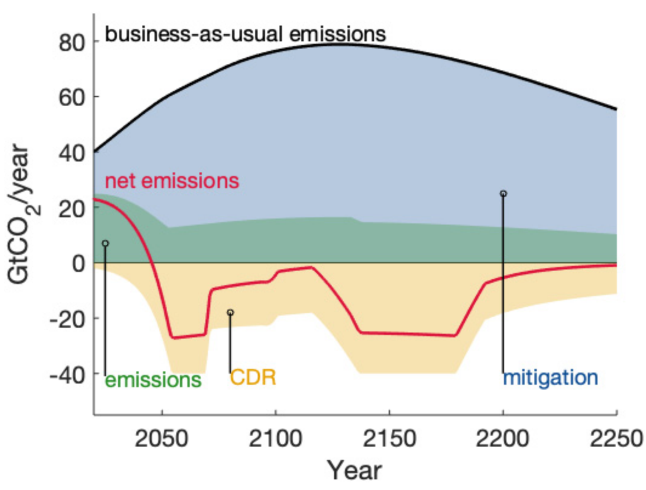

Figure B.6. Results for the 2° target climate policy when policy portfolio is limited to mitigation and CDR: mitigation (blue area), emissions (green area), and CDR (yellow area). Red line highlights net emissions.

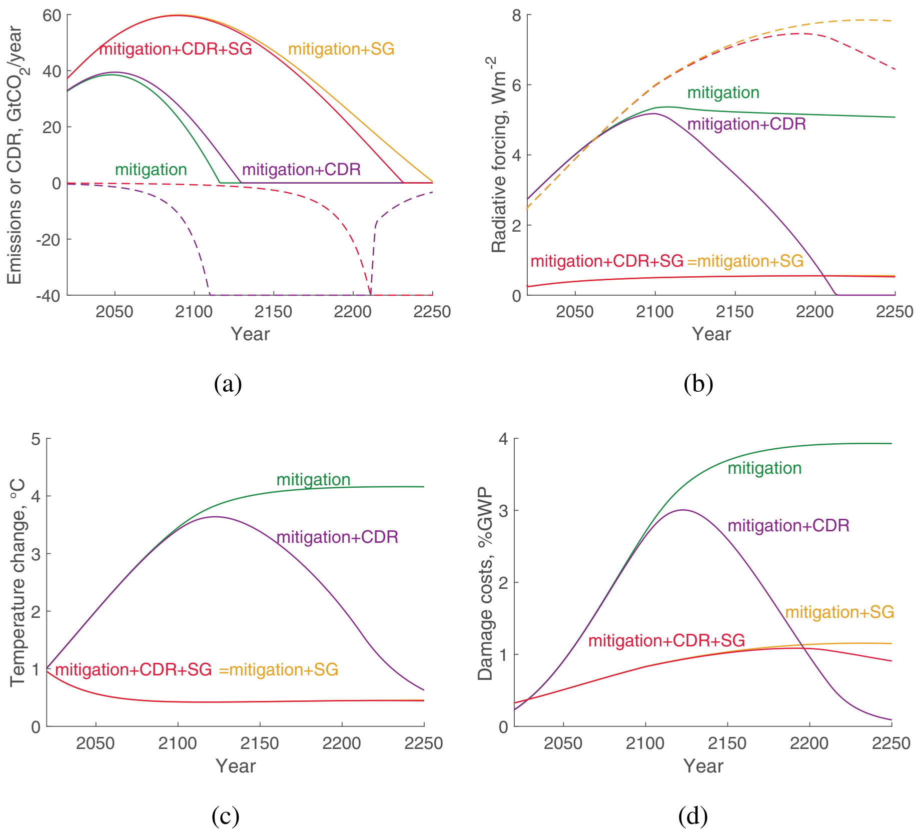

Figure B.7. Optimal policy and its implications in the case when SG efficacy is 10°, illustrated for alternative policy portfolio compositions (a)–(d).

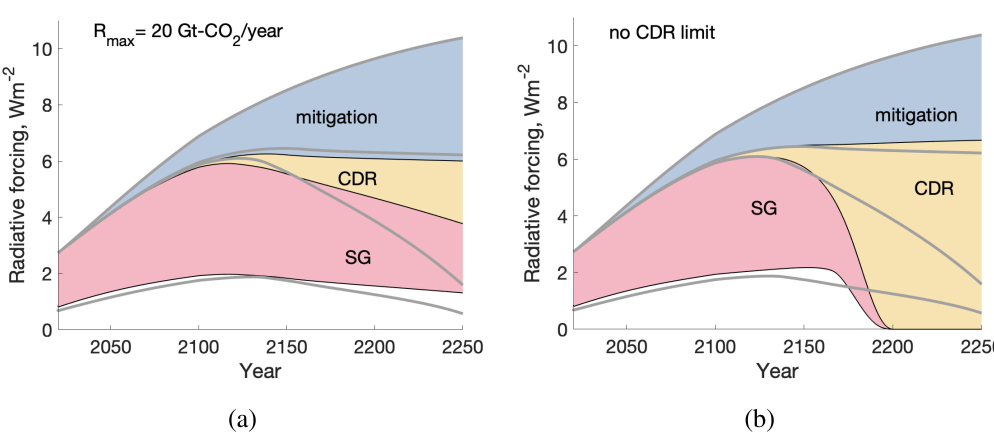

Figure B.8. Optimal climate policy portfolio with three instruments. Contributions of policy instruments to a reduction in BAU RF for the cases when the limit on annual CDR is set at 20 Gt-CO_2/year (a) and under no limit on annual CDR (b). Gray lines indicate reference results for three-instrument portfolio from Fig. 3(a).

References

Bahn, O, M Chesney, J Gheyssens, R Knutti and AC Pana (2015). Is there room for geoengineering in the optimal climate policy mix? *Environmental Science & Policy*, 48, 67–76. https://doi.org/10.1016/j.envsci.2014.12.014.

Bickel, JE and S Agrawal (2013). Reexamining the economics of aerosol geoengineering. *Climatic Change*, 119, 993–1006. https://doi.org/10.1007/s10584-012-0619-x.

Bickel, JE and L Lane (2009). *An Analysis of Climate Engineering as a Response to Climate Change*, p. 60. Frederiksberg: Copenhagen Consensus Center.

Burke, M, SM Hsiang and E Miguel (2015). Global non-linear effect of temperature on economic production. *Nature*, 527, 235–239. https://doi.org/10.1038/nature15725.

Davis, SJ, NS Lewis, M Shaner, S Aggarwal, D Arent, IL Azevedo, SM Benson, T Bradley, J Brouwer, Y-M Chiang, CTM Clack, A Cohen, S Doig, J Edmonds, P Fennell, CB Field, B Hannegan, B-M Hodge, MI Hoffert, E Ingersoll, P Jaramillo, KS Lackner, KJ Mach, M Mastrandrea, J Ogden, PF Peterson, DL Sanchez, D Sperling, J Stagner, JE Trancik, C-J Yang and K Caldeira (2018). Net-zero emissions energy systems. *Science*, 360, eaas9793. https://doi.org/10.1126/science.aas9793.

Dell, M, BF Jones and BA Olken (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature*, 52, 740–798. https://doi.org/10.1257/jel.52.3.740.

den Elzen, MGJ and PL Lucas (2005). The FAIR model: A tool to analyse environmental and costs implications of regimes of future commitments. *Environmental Modeling & Assessment*, 10, 115–134. https://doi.org/10.1007/s10666-005-4647-z.

Eastham, SD, DK Weisenstein, DW Keith and SRH Barrett (2018). Quantifying the impact of sulfate geoengineering on mortality from air quality and UV-B exposure. *Atmospheric Environment*, 187, 424–434. https://doi.org/10.1016/j.atmosenv.2018.05.047.

Emmerling, J and M Tavoni (2018a). Climate engineering and abatement: A 'flat' relationship under uncertainty. *Environmental and Resource Economics*, 69, 395–415. https://doi.org/10.1007/s10640-016-0104-5.

Emmerling, J and M Tavoni (2018b). Exploration of the interactions between mitigation and solar radiation management in cooperative and non-cooperative international governance settings. *Global Environmental Change*, 53, 244–251. https://doi.org/10.1016/j.gloenvcha.2018.10.006.

Gillett, NP, MF Wehner, SFB Tett and AJ Weaver (2004). Testing the linearity of the response to combined greenhouse gas and sulfate aerosol forcing. *Geophysical Research Letters*, 31, L14201. https://doi.org/10.1029/2004GL020111.

Giorgi, F (2006). Climate change hot-spots. *Geophysical Research Letters*, 33, L08707. https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2006GL025734 (accessed on 16 June 2020).

Goes, M, N Tuana and K Keller (2011). The economics (or lack thereof) of aerosol geoengineering. *Climatic Change*, 109, 719–744. https://doi.org/10.1007/s10584-010-9961-z.

Goldman Sachs (2019). Carbonomics The Future of Energy in the Age of Climate Change, p. 44, https://www.goldmansachs.com/insights/pages/carbonomics.html.

Gramstad, K and S Tjøtta (2010). Climate engineering: Cost benefit and beyond (MPRA Paper), University Library of Munich, Germany.

Harding, AR, K Ricke, D Heyen, DG MacMartin and J Moreno-Cruz (2020). Climate econometric models indicate solar geoengineering would reduce inter-country income inequality. *Nature Communications*, 11, 227. https://doi.org/10.1038/s41467-019-13957-x.

Heutel, G, J Moreno-Cruz and K Ricke (2016). Climate engineering economics. *Annual Review of Resource Economics*, 8, 99–118. https://doi.org/10.1146/annurev-resource-100815-095440.

Heutel, G, J Moreno-Cruz and S Shayegh (2018). Solar geoengineering, uncertainty, and the price of carbon. *Journal of Environmental Economics and Management*, 87, 24–41. https://doi.org/10.1016/j.jeem.2017.11.002.

House, KZ, CH House, DP Schrag, MJ Aziz (2007). Electrochemical acceleration of chemical weathering as an energetically feasible approach to mitigating anthropogenic climate change. *Environ. Sci. Technol.* 41, 8464–8470, https://doi.org/10.1021/es0701816.

Irvine, P, K Emanuel, J He, LW Horowitz, G Vecchi and D Keith (2019). Halving warming with idealized solar geoengineering moderates key climate hazards. *Nature Climate Change*, 9, 295–299. https://doi.org/10.1038/s41558-019-0398-8.

Irvine, PJ and DW Keith (2020). Halving warming with stratospheric aerosol geoengineering moderates policy-relevant climate hazards. *Environmental Research Letters*, 15, 044011. https://doi.org/10.1088/1748-9326/ab76de.

Kravitz, B, A Robock, L Oman, G Stenchikov and AB Marquardt (2009). Sulfuric acid deposition from stratospheric geoengineering with sulfate aerosols. *Journal of Geophysical Research: Atmospheres*, 114, D14109. https://doi.org/10.1029/2009JD011918.

Kravitz, B, D MacMartin, A Robock, P Rasch, K Ricke, J Cole, C Curry, P Irvine, D Ji, D Keith, JE Kristjánsson, J Moore, H Muri, B Singh, S Tilmes, S Watanabe, S Yang and J-H Yoon (2014). A multi-model assessment of regional climate disparities caused by solar geoengineering. *Environmental Research Letters*, 9, 074013. https://doi.org/10.1088/1748-9326/9/7/074013.

MacMartin, DG, B Kravitz, DW Keith and A Jarvis (2014). Dynamics of the coupled human-climate system resulting from closed-loop control of solar geoengineering. *Clim. Dyn.* 43, 243–258.

McClellan, J, DW Keith and J Apt (2012). Cost analysis of stratospheric albedo modification delivery systems. *Environmental Research Letters*, 7, 034019. https://doi.org/10.1088/1748-9326/7/3/034019.

Meinshausen, M, SCB Raper and TML Wigley (2011). Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 – Part 1: Model description and calibration. *Atmospheric Chemistry and Physics*, 11, 1417–1456. https://doi.org/10.5194/acp-11-1417-2011.

Moreno-Cruz, JB, KL Ricke and DW Keith (2012). A simple model to account for regional inequalities in the effectiveness of solar radiation management. *Climatic Change*, 110, 649–668. https://doi.org/10.1007/s10584-011-0103-z.

Moreno-Cruz, JB and DW Keith (2013). Climate policy under uncertainty: A case for solar geoengineering. *Climatic Change*, 121, 431–444. https://doi.org/10.1007/s10584-012-0487-4.

Moriyama, R, M Sugiyama, A Kurosawa, K Masuda, K Tsuzuki and Y Ishimoto (2017). The cost of stratospheric climate engineering revisited. *Mitigation and Adaptation Strategies for Global Change*, 22, 1207–1228. http://dx.doi.org.ezp-prod1.hul.harvard.edu/10.1007/s11027-016-9723-y.

National Academies of Sciences, Engineering, and Medicine (2018). *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*. Washington, D.C.: National Academies Press. https://doi.org/10.17226/25259.

National Academies of Sciences, Engineering, and Medicine (2021). Reflecting Sunlight: Recommendations for Solar Geoengineering Research and Research Governance. Washington, DC: The National Academies Press. https://doi.org/10.17226/25762.

Nordhaus, WD (2010). Economic aspects of global warming in a post-Copenhagen environment. *Proceedings of the National Academy of Sciences*, 107, 11721–11726. https://doi.org/10.1073/pnas.1005985107.

Nordhaus, W (2018). Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy*, 10, 333–360. https://doi.org/10.1257/pol.20170046.

Pitari, G, V Aquila, B Kravitz, A Robock, S Watanabe, I Cionni, ND Luca, GD Genova, E Mancini and S Tilmes (2014). Stratospheric ozone response to sulfate geoengineering: Results from the Geoengineering Model Intercomparison Project (GeoMIP). *Journal of Geophysical Research: Atmospheres*, 119, 2629–2653. https://doi.org/10.1002/2013JD020566.

Proctor, J, S Hsiang, J Burney, M Burke and W Schlenker (2018). Estimating global agricultural effects of geoengineering using volcanic eruptions. *Nature*, 560, 480–483. https://doi.org/10.1038/s41586-018-0417-3.

Ricke, KL, JB Moreno-Cruz and K Caldeira (2013). Strategic incentives for climate geoengineering coalitions to exclude broad participation. *Environmental Research Letters*, 8, 014021.

Smith, W and G Wagner (2018). Stratospheric aerosol injection tactics and costs in the first 15 years of deployment. *Environmental Research Letters*, 13, 124001. https://doi.org/10.1088/1748-9326/aae98d.

Tilmes, S, JH Richter, B Kravitz, DG MacMartin, MJ Mills, IR Simpson, AS Glanville, JT Fasullo, AS Phillips, J-F Lamarque, J Tribbia, J Edwards, S Mickelson and S Ghosh (2018). CESM1(WACCM) stratospheric aerosol geoengineering large ensemble project. *Bulletin of the American Meteorological Society*, 99, 2361–2371. https://doi.org/10.1175/BAMS-D-17-0267.1.

The U.S. Global Change Research Program (USGCRP) (2021). About USGCRP. https://www.globalchange.gov/about (accessed on 20 May 2021).

van Vuuren, DP, AF Hof, MAE van Sluisveld and K Riahi (2017). Open discussion of negative emissions is urgently needed. *Nature Energy*, 2, 902–904. https://doi.org/10.1038/s41560-017-0055-2.

Vinca, A, M Rottoli, G Marangoni and M Tavoni (2018). The role of carbon capture and storage electricity in attaining 1.5 and 2°C. *International Journal of Greenhouse Gas Control*, 78, 148–159. https://doi.org/10.1016/j.ijggc.2018.07.020.