



Complementing CO₂ emission reduction by Geoengineering might strongly enhance future welfare

Koen G. Helwegen¹, Claudia E. Wieners^{2,3}, Jason E. Frank¹, and Henk A. Dijkstra^{2,3}

¹Mathematical Institute; Utrecht University, Netherlands

²Institute for Marine and Atmospheric research, Utrecht; Utrecht University, Netherlands

³Centre for Complex Systems Studies, Utrecht; Utrecht University, Netherlands

Correspondence: Claudia E. Wieners (c.e.wieners@uu.nl)

Abstract. Solar Radiation Management (SRM) has been proposed as a means to reduce global warming in spite of high greenhouse gas concentrations and lower the chance of warming-induced tipping points. However, SRM may cause economic damages, and its feasibility is still uncertain. To investigate the trade-off between these gains and damages, we incorporate SRM into a stochastic-dynamic integrated assessment model and perform the first rigorous cost-benefit analysis of sulphate-based SRM under uncertainty, treating warming-induced climate tipping and SRM failure as stochastic elements. We find that SRM has the potential to greatly enhance future welfare and should therefore be taken seriously as a policy option. However, if only SRM and no CO₂ abatement is used, global warming is not stabilised and will exceed 2K. Therefore, even if successful, SRM cannot replace but only complement CO₂ abatement. The optimal policy combines CO₂ abatement and modest SRM and succeeds in keeping global warming below 2K.

10 *Copyright statement.*

1 Introduction

Despite the Paris agreement target to keep global mean temperature change “well below 2K” in order to prevent “dangerous climate change” (UNFCCC, 2015), no decisive reduction of CO₂ emissions has yet taken place (Le Quéré et al., 2018). This has sparked renewed interest in the possibility of cooling the climate system by geoengineering (Crutzen, 2006). Among several suggested approaches, only Solar Radiation Management (SRM), i.e. reflecting part of the incoming solar radiation back into space, has the potential to offset the global mean temperature changes projected by 2100 (Keller et al., 2014).

Of several proposed SRM techniques (Latham et al., 2008; Ahlm et al., 2017; Gabriel et al., 2017; Seneviratne et al., 2018), sulphate aerosol-based SRM is most likely to become feasible in the near future (McClellan et al., 2010; Moriyama et al., 2017). The scheme involves injecting precursor gases such as SO₂ into the stratosphere. This leads to the formation of reflective sulphate aerosols in the lower stratosphere which increase the Earth’s albedo and cause surface cooling. Such cooling – of about 0.4K over several years (Stowe et al., 1990; Thomson et al., 2009) – was observed following the Pinatubo eruption (Stowe et al., 1990; Thomson et al., 2009; Stenchikov et al., 1998; Robock, 2000) which injected 8-10 Mt(S) (megatonnes of sulphur),



mainly as SO_2 , into the stratosphere (Ward, 2009). It is still uncertain whether SRM can completely eliminate future global warming. High aerosol concentrations lead to faster coagulation, which reduces albedo and accelerates deposition (Visioni et al., 2017). One study (Kleinschmitt et al., 2018) suggests that the achievable forcing is limited to $-2\text{W}/\text{m}^2$ while others find that sufficiently strong forcing can be achieved, albeit at very high injection rates (Niemeyer and Timmreck, 2015; Niemeyer and Schmidt, 2017).

The potential benefits of SRM are obvious: a reduction of global warming and associated damages, and a reduced transition likelihood of temperature-related climate tipping points (Cai et al., 2016). However, SRM cannot avert all climate change (Kravitz et al., 2013). In particular, global mean precipitation is expected to decrease (Andrews et al., 2010; MacMartin and Kravitz, 2016), and the spatial precipitation patterns may change. Ocean acidification will continue unless atmospheric CO_2 concentrations are reduced (Tjiputra et al., 2015).

The implementation costs of sulphate SRM are estimated to be $2 - 10 \times 10^9 \$/\text{Mt}$ of injected gas (McClellan et al., 2010; Moriyama et al., 2017), which is modest compared to the world GDP of $80 \times 10^{12} \$$ (data for 2017, World Bank (2017)). However, apart from moral issues (Robock et al., 2009), sulphate SRM may have damaging effects on human health (Effiong and Neitzel, 2016) and the environment (Pitari et al., 2014; Ward, 2009) that are still poorly understood (Irvine et al., 2017). A sudden discontinuation of SRM will cause rapid warming ("termination shock") to levels dictated by greenhouse gas concentrations (Brovkin et al., 2009; Matthews and Caldeira, 2007), which could put more stress on ecosystems and societies than a gradual warming (Trisos et al., 2018).

A cost-benefit analysis of SRM should take into account at least two major uncertainties: the possibility of warming-induced tipping behaviour (whose likelihood is reduced by SRM) and the possibility of SRM failure, either by inefficiency or because damaging side-effects force one to abandon it. In this study we use a stochastic version of the integrated assessment model DICE (Nordhaus, 1992) to compute the optimal policy including CO_2 abatement and SRM. We hereby expand on earlier studies, which often included uncertainty only through parameter sensitivity analysis (Goes et al., 2011; Bahn et al., 2015) or as a simplified two-step decision problem (MorenoCruz and Keith, 2013). Two recent studies (Heutel et al., 2016, 2018) include climate tipping behaviour but employ a simple 4-step look-ahead scheme which is unsuitable for long-term optimisation. We employ dynamic programming (Bellman, 1957) to perform the first rigorous cost-benefit analysis of SRM under uncertainty.

The paper is organised as follows: In Sect. 2.1 we present our model GeoDICE, a stochastic DICE model including Geo-engineering, and in Sect. 2.3, we describe the scenarios employed. The results are presented in Sect. 3.1 (deterministic cases) and Sect. 3.2 - 3.3 (stochastic cases), with a sensitivity analysis in Sect. 3.4. A summary and discussion is presented in Sect. 4.

2 Methods

2.1 GeoDICE: a stochastic DICE model including Geoengineering

We modified DSICE (Cai et al., 2012), a stochastic version of the DICE model (Nordhaus, 1992), to include SRM as an additional policy option (together with CO_2 abatement). To this end, we incorporate the cooling effect, implementation costs,



Symbol	Meaning	Value
α_{CO_2}	Effect CO ₂ on radiative forcing	$5.35W/m^2$
α_{SO_2}	scales sulphate radiative forcing	$65W/m^2$
β_{SO_2}	scales sulphate radiative forcing	$2246Mt(S)/yr$
γ_{SO_2}	scales sulphate radiative forcing	0.23
η	sulphate rad. forcing correction	0.742
b_1	strength temperature response	$0.126 \frac{K}{W/m^2}$
b_2	strength temperature response	$0.0254 \frac{K}{W/m^2}$
τ_{T1}	time scale temperature response	1.89yr
τ_{T2}	time scale temperature response	13.6yr
p_T	precipitation dependence on temp.	$0.0806(mm/day)/K$
p_C	precipitation dependence on CO ₂	$-0.0229(mm/day)/(W/m^2)$
p_S	precipitation dependence on SRM	$-0.0077(mm/day)/(W/m^2)$
ψ_C	econ. damage from CO ₂ conc.	$1.703 \times 10^{-3} K^{-2}$
ψ_T	econ. damage from warming	$0.4(mm/day)^{-2}$
ψ_P	econ. damage from precip. change	$3.31 \times 10^{-8}(ppmv)^{-2}$
ψ_S	econ. damage from SRM	$9.27 \times 10^{-5}(Mt(S))^{-2}$
ψ_{fail}	econ. damage from SRM failure	$0.01/Mt(S)$
Ω	remaining fraction econ. output after tipping	0.9
λ_S	implementation cost SRM	$14 \times 10^9\$/Mt(S)$
κ_{tipp}	tipping probability per year and K warming	$0.00255/yr/K$
κ_{fail}	probability of SRM failure per year	$0.00056/yr$
ρ	pure rate of time preference (constant in time)	$0.015/yr$
δ_K	rate of capital depreciation	$0.065/yr$

Table 1. Model parameters of the GeoDICE model related to the representation of SRM. The carbon model parameters can be found in Table 5 of Joos et al. (2013), and others in DICE/DSICE (Cai et al., 2012).

and environmental damages of SRM into DSICE. A summary of the model parameters and their standard values is given in Table 1.

2.1.1 SRM and Radiative Forcing

To the radiative forcing equation, we add a contribution that depends sublinearly on the sulphur injection rate (Niemeyer and Timmreck, 2015). The total radiative forcing F takes the form

$$F = \alpha_{CO_2} \ln((C_{PI} + C)/C_{PI}) + F_{other} - \eta \alpha_{SO_2} \times \exp[(-\beta_{SO_2}/I_S)] \equiv F_C + F_{other} + F_S. \quad (1)$$



The first term F_C describes the contribution of the increase C in atmospheric CO_2 concentration above the preindustrial value C_{PI} and is the same as in DICE. The second term F_{other} represents the effects of other greenhouse gases (eg. CH_4 , N_2O , halogen compounds) and industrial aerosol. In DICE, this term is prescribed. However, it seems unlikely that a society that makes great efforts towards abating CO_2 emissions does nothing towards combatting other pollutants. Based on data from
5 Ciais et al. (2013) we assume that roughly 30% of the current CH_4 emissions (contributions related to fossil fuel production, e.g. leakage, and biomass burning), 10% of the N_2O emissions (likewise from industry and fossil fuel production) and 100% of the emission of halogen compounds could be abated, whereas 70% of the CH_4 emissions (natural sources and agriculture) and 90% of N_2O emissions (agriculture) cannot be abated. As a rough estimate, about 50% of the radiative forcing stemming from non- CO_2 greenhouse gasses could be abated. For simplicity, we assume that this also holds for (mainly cooling) industrial
10 aerosol. Thus we put

$$F_{other} = F_{other,DICE} \times (1 - \kappa\mu),$$

where $\kappa = 0.5$, μ is the abatement of CO_2 , i.e. the fraction of CO_2 emissions avoided, and $F_{other,DICE}$ is the prescribed contribution used in DICE.

The third term F_S describes the influence of sulphate SRM. The sulphate injection leads to a (negative) radiative forcing
15 at the top of the atmosphere which is given by $\alpha_{SO_2} \times \exp[(-\beta_{SO_2}/I_S)^{\gamma_S}]$, as found in an atmospheric chemistry modelling study (Niemeyer and Timmreck, 2015), where I_S is the annual injection rate of sulphur into the stratosphere (measured in Mt(S)/year). To achieve a modest radiative forcing of $-2\text{W}/\text{m}^2$ at the Top Of the Atmosphere (TOA), an annual injection of 10 Megatonnes of sulphur (Mt(S)), equivalent to one Pinatubo eruption, is required, whereas to achieve a forcing of $-8.5\text{W}/\text{m}^2$ (offsetting the greenhouse gas forcing projected for 2100 under the RCP8.5 scenario), 100 Mt(S)/year is needed. However,
20 due to fast adjustment processes, the TOA radiative forcing is not sufficient to predict the impact on global mean surface temperature (Kravitz et al., 2015). For example, it is found (MacMartin and Kravitz, 2016) that to compensate $7.42\text{W}/\text{m}^2$ forcing from quadrupling CO_2 , the solar constant would have to be reduced by $(4.2 \pm 0.6)\%$, which amounts to $10.1\text{W}/\text{m}^2$ TOA (taking into account the Earth's albedo). In other words, the top of atmosphere radiative forcing arising from changes in the solar constant is less efficient than forcing caused by CO_2 by a factor of $\eta = 0.742$. We assume here that sulphate SRM has
25 the same efficiency factor η as solar dimming, since both processes take place above the troposphere, and multiply the sulphate SRM contribution to F by this factor in eq. (1). Note that there is still considerable uncertainty about the forcing efficiency of SRM. For example, Tilmes et al. (2018) find higher efficiencies and an almost linear relationship for injection rates up to 25Mt(S)/yr , while Kleinschmitt et al. (2018) suggests that the maximal radiative forcing achievable with sulphate SRM might be limited to $2\text{W}/\text{m}^2$. This possibility that SRM is much less efficient is qualitatively included in the Realistic Storyline
30 scenario described below, which captures that SRM may never work at all. We impose an upper bound of $I_S \leq 100\text{Mt(S)/yr}$ on the injection rates, i.e. we do not allow them to exceed ≈ 10 Pinatubo eruptions per year. This upper limit is a much higher injection rate than considered in most detailed studies of the environmental and climate effects of SRM.



2.1.2 Carbon cycle and Climate response

We replace the carbon-climate part of DSICE by an emulator of full-fledged climate model simulations (Aengenheyster et al., 2018; MacMartin and Kravitz, 2016). We also include global mean precipitation as a proxy for the residual climate change (changes remaining if SRM is employed to keep global mean temperature constant).

- 5 As in DICE, CO₂ can be emitted by industrial processes and landuse change. The former contribution is calculated within our model, the latter is prescribed externally, using the same values as DICE. We model carbon concentrations based on the Green's function found by Joos et al. (2013).

Current CO₂ concentrations $C(t)$ (above pre-industrial) can be computed from emissions E at all previous times $t' < t$:

$$C(t) = \int_{t' < t} G_C(t - t') E(t') dt', \quad (2)$$

- 10 where $G(t - t')$ is the Green's function determining how a unit emission pulse contributes to the concentrations $t - t'$ years later, and $E(t')$ an emission pulse at time t' . Following Joos et al. (2013), $G_C(s)$ can be represented as a sum of exponentials, $G_C(s) = a_0 + \sum_{n=1}^N a_n e^{-s/\tau_n}$, with $n = 3$, and the temporal evolution of C can be rewritten as

$$C(t) = C_0(t) + \sum_{n=1}^N C_n(t), \quad (3a)$$

$$dC_0/dt = a_0 E, \quad (3b)$$

15 $dC_n/dt = a_n E - \frac{1}{\tau_n} C_n.$ (3c)

- Here a_0 represents the fraction of carbon emissions staying permanently in the atmosphere. The model parameters a_n, τ_n were obtained from a multi-model study (Joos et al. (2013), see their Table 5) and thus represent a best estimate for the behaviour of the carbon cycle, provided that nonlinear effects (e.g. saturation of carbon sinks with increasing CO₂ concentrations) are small. The initial values are $C_0 = 39.01 \text{ppmv}$, $C_1 = 35.84 \text{ppmv}$, $C_2 = 21.74 \text{ppmv}$, $C_3 = 4.14 \text{ppmv}$, since our model does not start
 20 at pre-industrial times (1765) but in 2005.

For the global mean temperature change (relative to pre-industrial), we follow the same approach, fitting the temperature response to a 1-year pulse of radiative forcing obtained by a multi-model study (MacMartin and Kravitz, 2016) onto a sum of exponentials, obtaining

$$T = T_1 + T_2 \quad (4a)$$

25 $dT_n(t)/dt = b_n F - \frac{1}{\tau_{T_n}} T_n$ (4b)

where F is the radiative forcing from eq. (1) and other parameters are in Table 1. For the temperature response to a radiative forcing pulse, there is no permanent response T_0 . The initial values (year 2005) are $T_1 = 0.466 \text{K}$ and $T_2 = 0.436 \text{K}$.

The response of global mean precipitation P to CO₂-induced or SRM-induced radiative forcing is based on MacMartin and Kravitz (2016) and can be split into a temperature-driven increase of $2.5\%/K$ and an instantaneous contribution due to CO₂



and SRM (Andrews et al., 2010). More specifically, the response G_P^f of the global mean precipitation P to a one-year-long $1W/m^2$ pulse of radiative forcing from agent f (f stands for CO_2 or a change in the solar constant) in year 0, obtained by MacMartin and Kravitz (2016), can be expressed as

$$G_P^f(t) = b_f \delta_{t,0} + a G_T(t), \quad (5)$$

5 where G_T is the temperature response to a $1W/m^2$ forcing and $\delta_{t,0} = 1$ if $t = 0$ and 0 else. This means that in the year of the forcing pulse, the fast response b_f plays a role, whereas in later years, the precipitation response is determined by the temperature response. As before, we use the result for reduction in the solar constant as a proxy for sulphate SRM. By lack of data, the fast response to other forcing agents constituting F_{other} is ignored. With these results, the change in global mean precipitation w.r.t pre-industrial can be written as

$$10 \quad P(t) = p_T T(t) + p_C F_C(t) + p_S F_S(t). \quad (6)$$

Here, $F_C > 0$ (CO_2 warms the globe), $F_S < 0$ (SRM cools), $p_T > 0$ and $p_C < p_S < 0$. The reason for the last relation is that increased CO_2 concentrations warm the atmosphere more quickly than the sea surface, which temporarily leads to a more stable stratification, suppressing convective rainfall. For SRM, the instantaneous response is weaker. Therefore, if SRM were employed such as to cancel the global mean temperature change ($F_S = -F_C - F_{other}$, hence $T = 0$), the slow responses

15 stemming from temperature change would cancel and the fast response to CO_2 would dominate, reducing P .

We use P as a proxy for residual climate change, i.e. for all effects which remain even if global mean temperature changes are cancelled by SRM.

2.1.3 The Damage function and SRM costs

As in DICE, the gross domestic product (GDP) \bar{Y} is diminished by climate-related damage and by expenditures for climate

20 policy (CO_2 abatement and SRM implementation). Including these losses, we retain for the net output:

$$Y = \Omega \frac{1}{1+D} \Lambda \bar{Y} - \lambda_S I_S. \quad (7)$$

Here, Ω describes damage due to tipping points (see sect. 2.1.4). If tipping has occurred, then $\Omega = 0.9$ (reducing the economic output), else $\Omega = 1$ (output not reduced). $D \geq 0$ describes non-tipping damage (discussed below). Λ is a factor describing the abatement costs ($\Lambda < 1$ in case of abatement, and $\Lambda = 1$ in case of no abatement) taken over from DICE. Since CO_2 emission

25 is proportional to \bar{Y} , so are abatement costs (the more economic output, the more CO_2 emissions and hence the higher the costs of eliminating a fraction μ of these emissions). $\lambda_S I_S$ is the implementation cost of SRM, which we assume to be linear in the injection rate I_S and independent of \bar{Y} . Two studies (McClellan et al., 2010; Moriyama et al., 2017) suggest that the costs for lifting gasses to 20km height are of the order of $2 - 10 \times 10^9 \$/Mt$ of injected gas. Taking an intermediate value of $7 \times 10^9 \$/Mt$, and assuming that the gas used is SO_2 (which has twice the molecular weight of elementary S), this amounts to

30 $14 \times 10^9 \$/Mt(S)$.

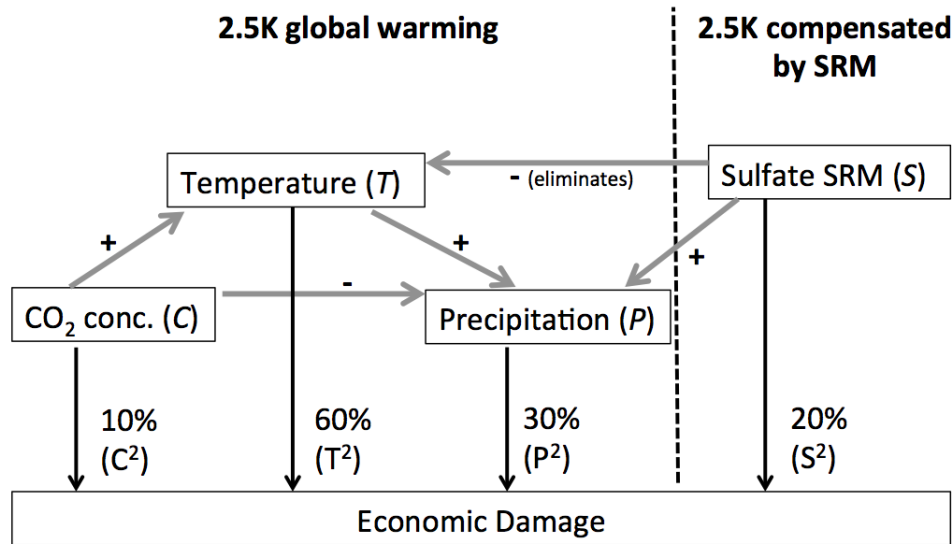


Figure 1. Graphical representation of the damage function. The black thin arrows represent contributions to the damage function, while grey arrows depict how the climate variables influence each other (+ and - stand for increasing and decreasing effects, respectively). The percentages for C , T , and P are based on the contributions of these variables for the standard case of 2.5K warming in equilibrium and in absence of SRM. In this standard case, our damage equals that of the DSICE model (Cai et al., 2012). The sulphur injections that would be needed to offset 2.5K warming cause a damage of 20% of the standard damage function.

While the original DICE model assumes that climate-induced damage D scales with the square of temperature change T ($D(T) = \psi_0 T^2$ with constant ψ_0), we keep the quadratic structure but split the damage function into three climate-related contributions and one contribution representing the damage inflicted by sulphate SRM (see fig. 1):

$$D(T, P, C, I_S) = \psi_T T^2 + \psi_P P^2 + \psi_C C^2 + \psi_S I_S^2 \quad (8)$$

- 5 where T , P , C are the changes (w.r.t. pre-industrial) of global mean temperature, global mean precipitation and atmospheric CO₂ concentration, respectively, and I_S the sulphur injection rate in $Mt(S)/yr$. Note that both positive and negative precipitation changes P are considered damaging, because both require humanity to adapt. The coefficients ψ (values: see Table 1) are chosen such that for the standard case of 2.5K warming in equilibrium without SRM, our total damage equals that of the original DICE model, and the contributions by T , P , and C are 60%, 30%, and 10%, respectively. The standard case was
- 10 determined by running the climate module of GeoDICE to equilibrium with constant greenhouse gas concentrations, such as to obtain $T = 2.5K$. Following Aengenheyster et al. (2018), and approximately in agreement with RCP scenarios, it was hereby assumed that other forcing agents (other greenhouse gases and aerosols constituting F_{other}) contribute 14% of the total radiative forcing. These other forcing agents are not assumed to cause direct damage. The damages associated with the annual sulphur injections needed to offset a warming of 2.5K are assumed to equal 20% of the standard case damage.



Previous studies (Heutel et al., 2016, 2018) likewise split the damage function, but without including residual climate change (P). They assumed increasing CO_2 concentrations to cause 20% of the damage, evenly split over oceanic and atmospheric CO_2 . We reduced this to 10% because atmospheric CO_2 is not known to cause substantial direct damage and may even be beneficial to plant growth (Ciais et al., 2013). The splitting between T and P is somewhat arbitrary, but is based on the rough assumption that, although precipitation changes can have substantial impact, much of the damage is either determined by temperature (especially sea level rise, a major contributor) or at least strongly influenced by it (e.g. hurricanes). The damage related to SRM depends on the injection rate, not on the percentage of compensated greenhouse gas forcing as was (somewhat unrealistically) assumed in earlier studies (Heutel et al., 2016, 2018). The choice of ψ_S is again somewhat arbitrary, as virtually no data on the economic damage of SRM is available. However, our main conclusions are unaffected by the exact choice of the parameters ψ (see Sect. 3.4).

2.1.4 Tipping points and SRM failure

Climate change may not only lead to smooth and predictable damages, but also induce low-probability, high-impact, irreversible events such as a collapse of ice sheets (Cai et al., 2015, 2016). The chance of such tipping behaviour is thought to increase with temperature. We take tipping into account in a stylised way, assuming that there is one tipping event that, once activated, reduces GDP by 10% for all subsequent time steps (i.e. $\Omega = 0.9$ in eq. (7)). The likelihood of tipping obeys

$$L_{\text{tipp}} = \begin{cases} 0 & T < 2K \\ (T - 2K) \times \kappa_{\text{tipp}} & T > 2K \end{cases}, \quad (9)$$

i.e., it is zero if the global mean temperature change $T < 2K$, but increases linearly with warming above $2K$. The constant κ_{tipp} is chosen such that in a scenario where the policy maker uses only abatement and remains unaware of possible tipping behaviour, the probability of tipping within 400 years is 50%. This order of magnitude of the likelihood and damage of tipping is consistent with earlier studies (Cai et al., 2016).

We also take into account the possibility that SRM has to be abolished (e.g. due to new scientific evidence about unforeseen dangers) by assuming that each year, there is a probability κ_{fail} that SRM may not be applied anymore in the future. The cumulative probability of SRM failure over 400 years is 20%. Unless stated otherwise, we include no economic damage related to SRM failure, because humanity is optimistically assumed to realise such dangers and abandon SRM in time.

2.2 Optimisation and Performance Measures

As in DICE, we assume that all decisions are made by a single policy maker who aims to optimise the welfare of the (homogeneous) world population. As in DICE, welfare depends entirely on consumption. The economic output is spent on investment $I = rY$ and global consumption $Lc = Y - I$ where r is the saving rate, and L and c are the world population and per capita consumption, respectively. We assume a fixed saving rate of $r = 22\%$. The utility u (which can be thought of as the current "happiness" of the world population) depends on c : $u = L(c^{1-\gamma} - 1)/(1 - \gamma)$ with $\gamma = 2$ and the quantity to be maximised is



the welfare (time-integrated, discounted utility):

$$W = \sum_t u(t)e^{-\rho t} \quad (10)$$

where t is (discrete) time and ρ the rate of pure time preference. The greater ρ , the less does the far future count towards W . The morally correct value of ρ has been fiercely discussed (Stern et al., 2007; Lilley, 2012; Ackerman, 2007). Here, we will not
 5 join the ethical debate on the ‘correct’ value, but use the standard value of 1.5% and perform a sensitivity study with $\rho = 0.5$ (see Sect. 3.4).

The decision variables are the amount of CO₂ abatement μ (the fraction of CO₂ avoided) and the sulphur injection rate I_S . The model is integrated in yearly time steps, but the decision variables μ and I_S are changed only once a decade to save computational effort. The "policy" (sequence of values for μ and I_S) is optimised such as to maximise the expected
 10 welfare (expectation value of W) over a time horizon till 2400, though the far future is heavily discounted. The optimisation is performed using Dynamic Programming (see appendix). Once an optimal policy is found, it is evaluated by running an ensemble of 5000 members with this policy and Monte-Carlo realisations of the stochastic elements (climate tipping and SRM failure). The best policy is the one yielding the highest expectation value of the welfare W . For easier comparison, we define a performance measure based on the improvement of W under policy π with respect to a no-action policy ($\mu = 0, I_S = 0$):

$$15 \quad \zeta(\pi) = 100\% \times \frac{W_\pi - W_0}{W_{AD} - W_0} \quad (11)$$

where W_π , W_0 , and W_{AD} are the expectation values over the Monte-Carlo ensemble of the welfare associated with policy π , the no-action case, and the optimal policy for the deterministic Abatement-only case (the policy that would be optimal if the decision maker may only use abatement, and no climate tipping occurs), respectively. By construction, the relative performance is 100% for Abatement-only in the deterministic case.

20 To evaluate the performance of policies, one may not only be interested in the ensemble mean, but also the range of possible values. In particular, a risk-averse policy maker might also be interested in the worst (or a relatively bad) case. Hence we present two additional performance measures, based on the 10th and 90th percentile of the Welfare W_π . Similar to eq. (11) we define $\zeta_X(\pi)$, the X-percentile relative performance of a policy π , as

$$25 \quad \zeta_X(\pi) = 100\% \times \frac{W_{\pi,X} - W_0}{W_{AD} - W_0} \quad (12)$$

where $W_{\pi,X}$ is the Xth percentile of the welfare (discounted cumulated utility) for policy π . Note that W_{AD} and W_0 are still the mean (i.e. not percentiles) welfare associated with the optimal policy for the deterministic Abatement-only case, and the no action case, respectively.

2.3 Scenarios

In Sect. 3.1 - 3.2 we first consider three stylised policy scenarios. The first is Abatement-only, in which the decision maker
 30 is allowed to use CO₂ abatement but no SRM. The second is SRM-only, in which the decision maker uses only SRM, until an SRM failure occurs, after which only abatement may be used. This scenario represents a society which does not reduce



Policy	ζ	peak SRM	Ab. 50%	Ab.90%	SCC
Abatement-only	100%	/	2114	2212	35
SRM-only	186%	*	/	/	21
Abatement+SRM	238%	35.1	2134	2243	20

* SRM does not peak, but keeps increasing until the upper limit of $100Mt(S)/yr$.

/ = Not applicable

Table 2. Comparison of policies in the deterministic setting (no tipping, no SRM failure). Abatement-only means that no SRM is used, SRM-only means that no abatement is used (unless SRM fails; see text), and in Abatement+SRM both are used. The performance ζ (see eq. (11)) is a measure of the increase in expected cumulated discounted utility w.r.t. the no-action scenario, and is normalised such as to yield 100% for Abatement-only. The column ‘peak SRM’ contains the highest SRM values (in $Mt(S)/yr$) over all time steps. ‘Ab 50%’ and ‘Ab 99%’ show the year in which the abatement reaches 50% and 99%, respectively. SCC is the social cost of carbon in $(\$/2005)/t(C)$.

CO₂ emissions but relies entirely on SRM (until it fails). The third is Abatement+SRM, wherein the decision maker can use both Abatement and SRM, unless SRM fails, after which only abatement is used. These three scenarios are first discussed in a deterministic setting (Sect. 3.1), i.e. in absence of climate tipping and SRM failure, before addressing them in the full model with uncertainty (Sect. 3.2).

5 While the previous stylised scenarios serve to isolate specific effects, we also present a more Realistic Storyline (see Sect. 3.3) where we assume that SRM will become possible only in 2055, and only at 30% probability. In the first decade where SRM is allowed, there is a 20% probability of SRM failure; in the second decade, 10%; in the third decade 5% and after that 1% per decade, i.e. after some decades of testing, failure becomes less likely.

3 Results

10 3.1 The deterministic case

As a reference, we first consider the deterministic case, i.e. without SRM failure and tipping points, in the three stylised scenarios (see Fig. 2 and Table 2). Allowing SRM in addition to abatement delays abatement by 2-3 decades, but does not replace it (Fig. 2a). CO₂ concentrations in Abatement+SRM (Fig. 2c) peak slightly later than in Abatement-only and reach higher values (875ppmv instead of 741ppmv). SRM helps to reduce global warming considerably: The global mean temperature change T peaks at 1.6K for Abatement+SRM, but at 3.1K for Abatement-only (Fig. 2d).

15 In SRM-only, CO₂ concentrations reach 2000ppmv in 2260 and continue to increase (Fig. 2c). The temperature increase T continues to rise, reaching 5.4K in 2400 (Fig. 2d), although it is lower than in the no-action case (neither SRM nor abatement). Due to the sub-linear increase of the radiative forcing with SRM, very high injection rates would be needed to stabilise T with SRM-only, so that the damage related to sulphate injection outweighs the climate damages. Compared to SRM-only, considerably less SRM is needed in Abatement+SRM, namely $\approx 35Mt(S)/year$ (Fig. 2b; 3 – 4 Pinatubo eruptions per year), yet T remains much lower.

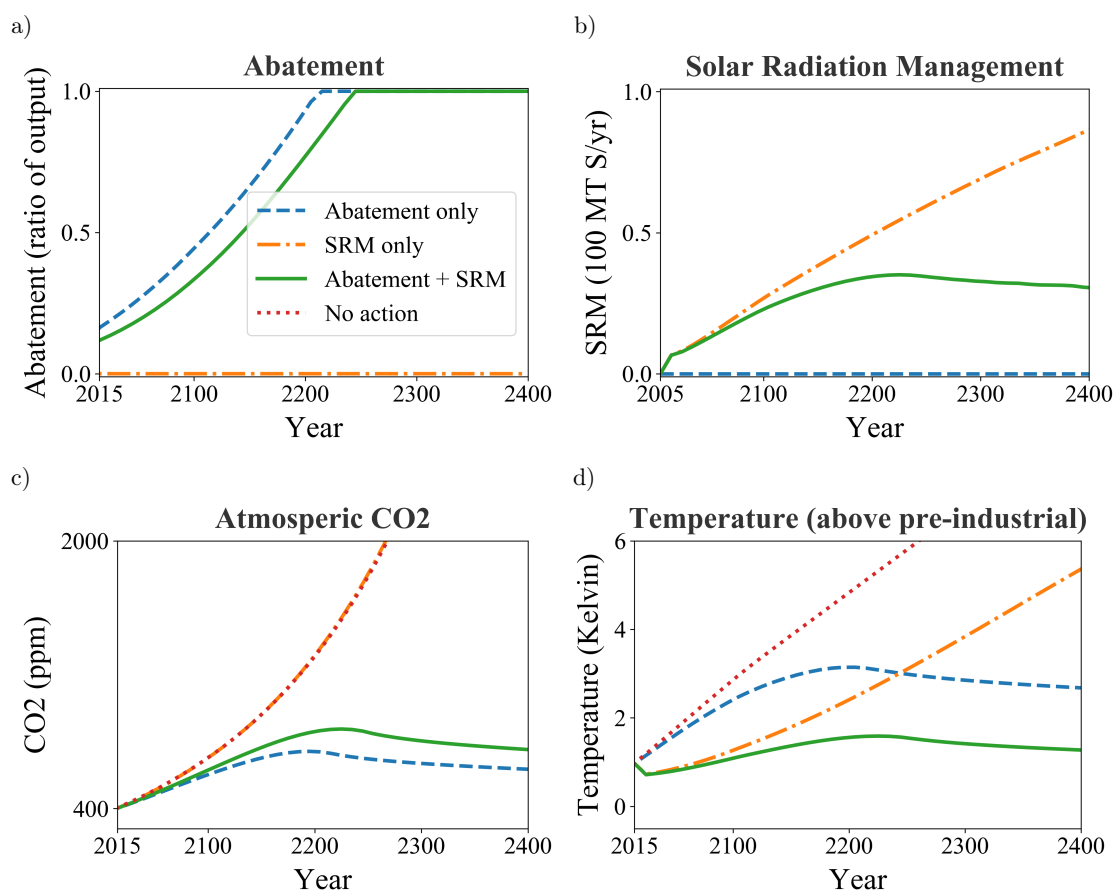


Figure 2. Optimal policy and climate model results for the three stylised scenarios in the deterministic setting. a) Abatement (fraction of CO₂ emissions avoided), b) SRM (Mt(S)/yr), c) atmospheric CO₂ concentration in ppmv, d) global mean temperature above pre-industrial (K). The blue dashed line represents the Abatement-only scenario, yellow dash-dotted: SRM-only, solid green: Abatement + SRM, red dotted (only plot c and d): no climate action (i.e. neither abatement nor SRM).



Policy	ζ	ζ_{10}	ζ_{90}	SRM fail	Tipping	peak SRM	Ab. 50%	Ab.90%	SCC
No action	0%			/	96.2%	/	/	/	45
Abatement-only (det. policy ^{**})	100%			/	49.5%	/	2114	2212	42
Abatement-only	105%	77%	121%	/	37.8%	/	2095	2215	41
SRM-only	181%	179%	185%	19.8%	60.96%	*	/	/	23
Abatement + SRM	219%	220%	223%	20.2%	6.2%	35.0	2139	2242	20
Realistic Storyline	125%	78%	190%	79.9%	30.1%	31.4	2106	2234	37

* SRM does not peak, but keeps increasing until the upper limit of $100Mt(S)/yr$.

** Tipping can occur, but the policy maker ignores this and chooses the policy which would be optimal in the deterministic case.

/ = Not applicable

Table 3. Comparison of policies in the stochastic setting, i.e. including climate tipping and SRM failure. No action means that neither abatement nor SRM are used; other scenarios are explained in Sect. 2.3. The performance measures ζ , ζ_{10} and ζ_{90} are given in eq. (11) and eq. (12). The columns ‘SRM fail’ and ‘Tipping’ show the probability that SRM failure or climate tipping occurs before 2415. The column ‘peak SRM’ contains the highest SRM value (in $Mt(S)/yr$) over all time steps and over all ensemble members. This corresponds to members in which no SRM failure or climate tipping occurred, at least before the time of the SRM peak. ‘Ab 50%’ and ‘Ab 99%’ show the year in which the abatement reaches 50% and 99%, respectively. SCC is the social cost of carbon in $(\$(2005)/t(C))$.

The relative performance $\zeta(\pi)$ for SRM-only and Abatement+SRM, becomes 186% and 238%, respectively (see Table 2). (By construction, $\zeta = 100\%$ for Abatement-only in the deterministic setting.) The reason for the better performance of SRM-only compared to Abatement-only is that SRM-only yields lower temperatures and a higher utility in the first two centuries, which contribute most to the cumulative utility due to discounting. In addition, postponing damage is beneficial as it allows
 5 more time for accumulating capital.

3.2 The effect of uncertainties

Next, we include the stochastic elements, temperature-induced tipping and SRM failure, and determine again the optimal policies for each scenario, prior to evaluating the optimal policy by means of a Monte-Carlo ensemble (see Sect. 2.2). In Fig. 3b - 1, we plot the policy (abatement and SRM) carbon concentration and temperature for the three stylised scenarios. The
 10 plots depict some sample paths of individuals Monte-Carlo runs (thin blue lines), the range of possible outcomes (shading) and the ensemble mean (red line). For comparison, the results from the deterministic case (compare Fig. 2) are also plotted (blue dashed lines).

In the Abatement-only scenario, the danger of tipping initially leads to higher abatement (Fig. 3d) than in the deterministic case, although the temperature is not kept below the $2K$ threshold (see Fig. 3j). If tipping occurs, the abatement decreases
 15 again, as there is no further tipping point to be avoided. The relative performance is 105% (see Table 3, row 3), i.e. it slightly improves when the decision maker takes tipping into account (compare Table 3, row 2). Recall that the reference scenario for W_{AD} uses the policy that would be optimal in absence of tipping, i.e. the policy maker ignores climate tipping.

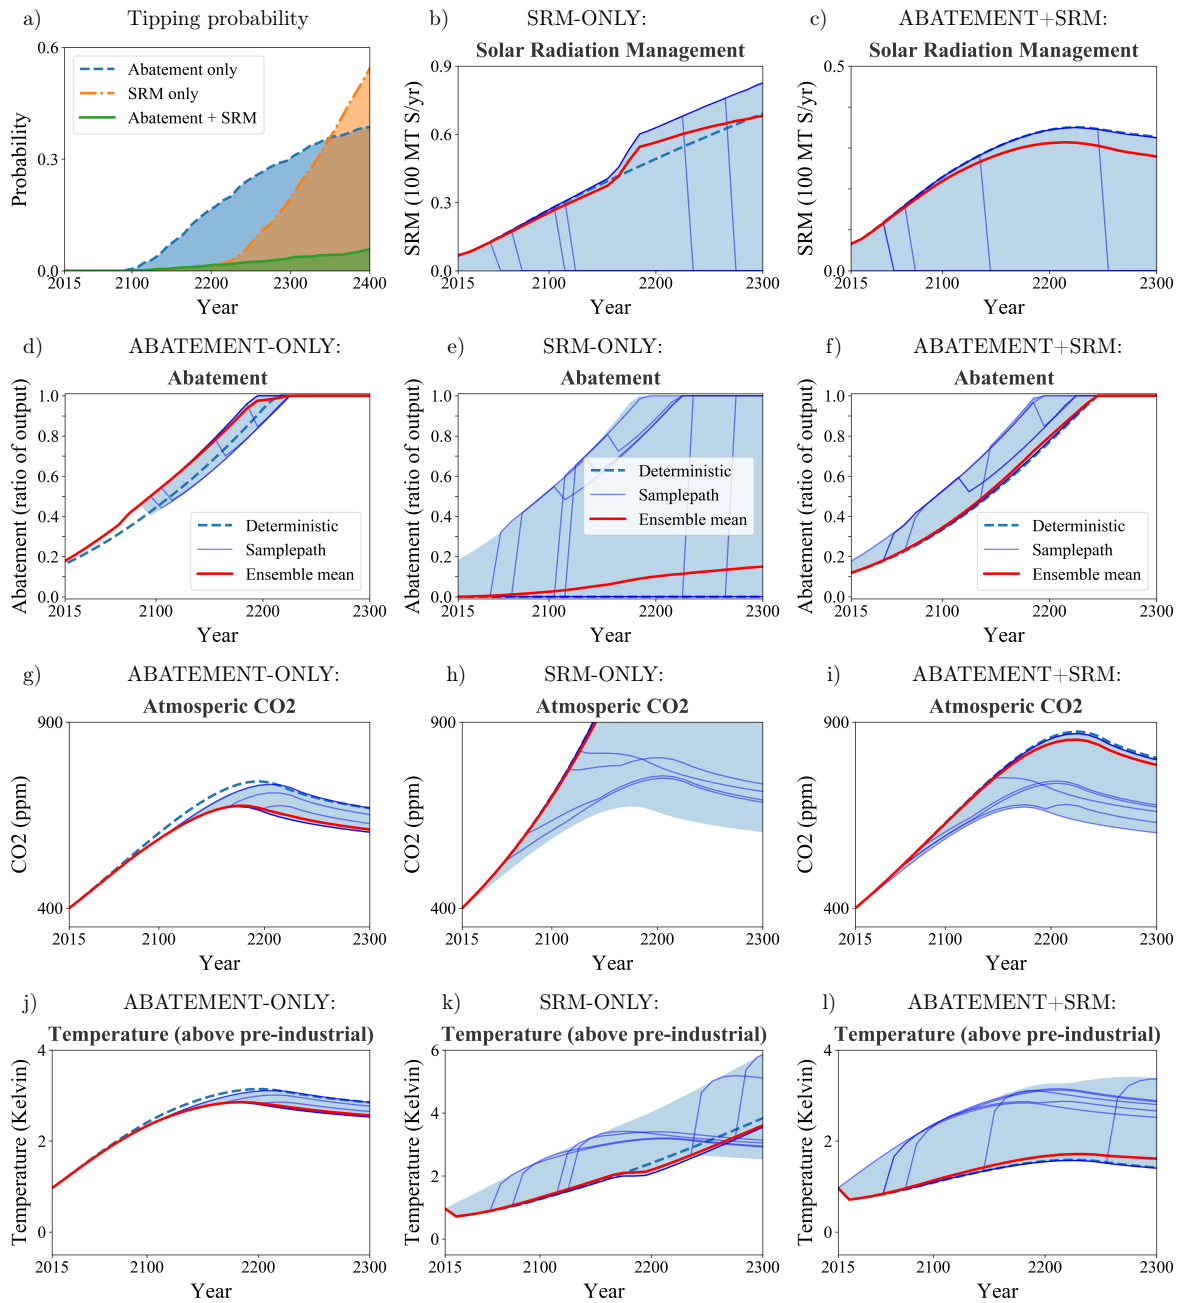


Figure 3. Tipping risk and policy in the stochastic setting (i.e. with tipping point and SRM failure). a) cumulative probability of tipping for Abatement-only (blue dashed line), SRM-only (yellow dash-dotted) and Abatement+SRM (green solid). b-l): policy and climate response for the same scenarios, namely Abatement-only (left column, plots d),g),j)), SRM-only (middle column, plots b),e),h),k)), and Abatement+SRM (right column, plots c),f),i),l)). Variables shown are: SRM deployed (first row, plots b), c)); note the different y-axis scale), abatement fraction (second row, plots d),e),f)), atmospheric CO₂ content in ppmv (third row, plots g),h),i)), and global mean temperature change (last row, plots j),k),l)); note the different y-axis scale). In plots b-l) only the first 285 years of the simulation are shown. The thin blue lines represent a sample of individual ensemble members, the thick red line the ensemble mean and the blue shaded area indicates the range of possible values in the whole ensemble. The dashed blue line depicts the results from the deterministic case (Fig. 2) for reference.



In the Abatement+SRM case, the optimal policy closely resembles the deterministic one if no SRM failure occurs (Fig. 3c,f). Without SRM failure, T stays below $2K$ (Fig. 3l), and hence no tipping occurs. In case of an SRM failure, the temperature suddenly increases and abatement suddenly increases, as the decision maker now tries to limit the warming (and tipping risk) with only abatement. Note that such a sudden increase in abatement may not be feasible in reality. If climate tipping occurs, abatement is reduced again. Compared to Abatement-only, the abatement is delayed by 3-4 decades.

In the SRM-only scenario, the policy again resembles the deterministic case provided no SRM failure occurs and T is below $2K$ (Fig. 3b). When $T = 2K$ is reached, SRM increases sharply to reduce the tipping risk. As before, abatement strongly increases after SRM failure, but is reduced slightly if tipping occurs (Fig. 3e). SRM-only has a performance of 181%, much higher than Abatement-only. However, the chance of climate tipping by the year 2415 is considerably higher for SRM-only (61.0% vs 37.8% for Abatement-only, see Table 3). As in the deterministic setting, the reason is that initially SRM can control the global warming more effectively than abatement, while abatement is a long-term measure. Hence damage is postponed to the far future which is heavily discounted. The cumulative probability of tipping is lower for SRM-only than for Abatement-only until 2350, when the situation reverses (Fig. 3a).

Compared to the deterministic cases, including uncertainty slightly reduces the difference in relative performance between Abatement-only and the scenarios using SRM (compare Table 3 vs Table 2). There are two competing effects: The danger of tipping might favour using SRM, which reduces the tipping probability in the near future, while the possibility of SRM failure reduces the performance of SRM-based scenarios.

In Abatement-only, there is a high spread between the relative performance measures ζ , ζ_{10} and ζ_{90} , compared to SRM-only and Abatement+SRM. This is due to the fact that in most ($> 90\%$) of the ensemble members, SRM keeps global warming below $2K$ at least until ≈ 2200 . Hence SRM postpones climate tipping into the far future (except in the few ensemble members with early SRM failure), while for Abatement-only, tipping can occur as early as 2080. Early tipping greatly reduces the performance because it reduces the GDP for a long period of time and because it is less heavily discounted. For Abatement+SRM, only 6.2% (i.e. $< 10\%$) of the ensemble members show climate tipping, but they strongly affect the mean performance. This explains why, for this scenario, $\zeta < \zeta_{10}$.

Although DICE is too limited to give reliable absolute values of the Social Cost of Carbon (SCC) (van den Bergh and Botzen, 2015), comparing scenarios gives qualitative insight into how SRM affects the SCC (Table 2 and Table 3). For Abatement-only, the SCC in 2015 is $35\$/t(C)$ (in 2005\$) in the deterministic case and $41\$/t(C)$ when including tipping points. For Abatement+SRM, the SCC is $20\$/t(C)$ (both deterministic and stochastic): SRM lowers the SCC by partially compensating the damage caused by CO_2 emissions. For SRM-only, the SCC is only slightly higher, namely $21\$/t(C)$ (deterministic) and $23\$/t(C)$ (stochastic), because SRM suppresses most climate damage in the near future, which is discounted least.

3.3 Realistic Storyline

The previous scenarios were very stylised, in order to isolate the impact of SRM and stochastic elements. However, the actual situation is more complex: Presently SRM is not available and we do not know whether it ever will be; yet we might want to decide now whether to pursue (research and development of) SRM. To address this question, we consider a "Realistic Storyline"

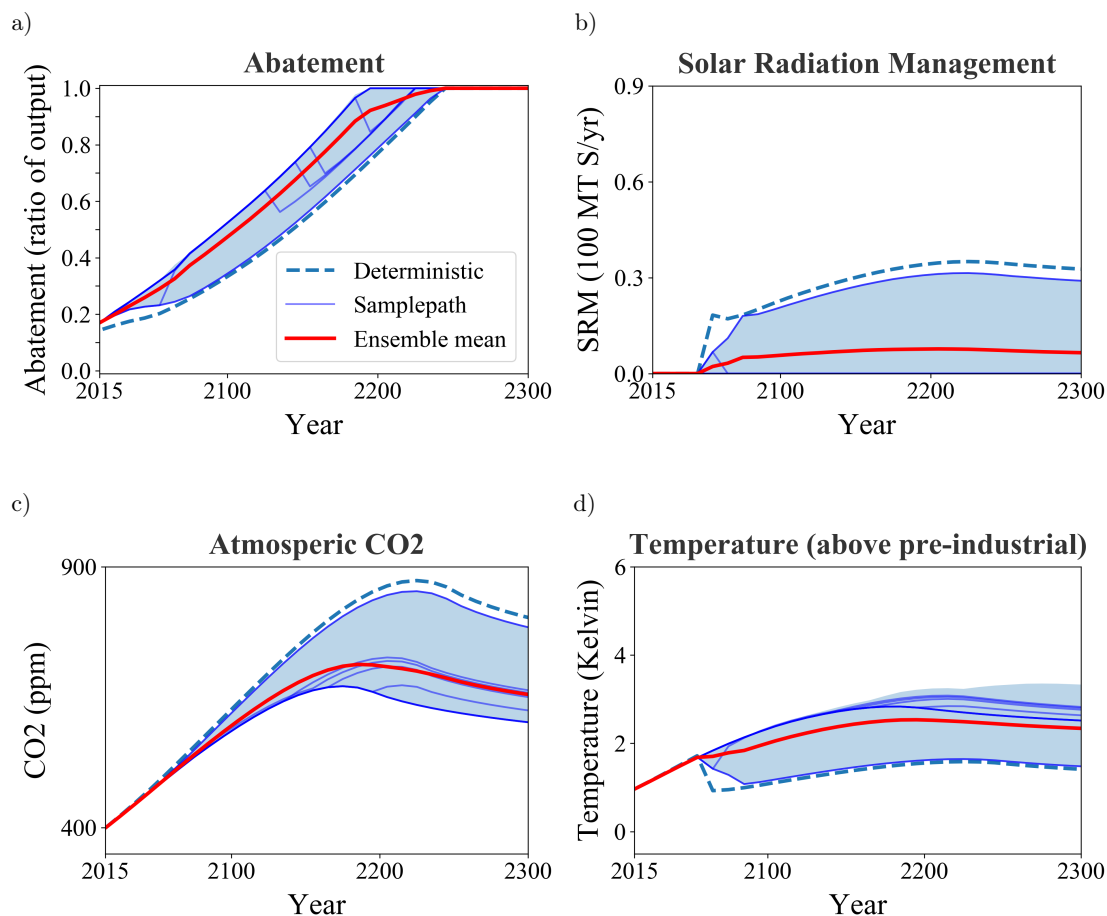


Figure 4. Optimal policy and climate development for the Realistic Storyline scenario. a) Abatement fraction, b) SRM in 100 Mt(S)/yr, c) atmospheric carbon concentration in ppmv, d) global mean temperature change w.r.t pre-industrial. The thin blue lines represent a sample of individual ensemble members, the thick red line the ensemble mean, and the blue shaded area indicates the range of possible values in the whole ensemble. The dashed blue line depicts the results from a deterministic reference case in which SRM becomes available in 2055 certainly and neither SRM failure nor climate tipping occur.



scenario, in which we assume that SRM will become possible only in 2055, and only at 30% probability (in the decades before 2055, there is a certain probability each time step that SRM is declared infeasible, e.g. because scientists identify unacceptable environmental risks). We also assume that after 2055, the probability of SRM failure decreases in time, i.e. with ongoing testing (see Sect. 2.3). As a coarse representation of damages associated with the termination shock, SRM failure is accompanied by a one-time reduction of the GDP by a factor $1 - \psi_{fail} I_S$ where I_S is the injection rate in Mt(S)/yr and ψ_{fail} is given in Table 1. Unlike (irreversible) climate tipping, the termination shock is a short-lived phenomenon and is stronger for large SRM.

In those ensemble members where SRM becomes available in 2055, it is used sparingly in the first time step, because the probability of failure is still high and the decision maker wants to limit the termination shock. In later time steps, SRM is used only slightly less than in the Abatement+SRM scenario, peaking at 31.4% rather than 35%. This difference mainly arises because the decision maker wants to reduce the termination shock: If the termination shock damage is omitted from the Realistic Storyline, SRM peaks at 34.7%.

In the first time step (2015), when the decision maker assumes that SRM will become available with 30% probability only, the abatement is $\mu(2015) = 0.17$, only slightly less than in the Abatement-only scenario where $\mu(2015) = 0.18$. For comparison, in a deterministic reference case in which SRM will be available from 2055 certainly, and no SRM failure or tipping occurs, $\mu(2015) = 0.14$ (see Fig. 4). As time progresses until 2055, the ensemble members diverge: If SRM is already banned, abatement increases, but if a time step has passed without a ban, the decision maker becomes more optimistic that SRM will become feasible and abatement becomes less ambitious. In ensemble members where SRM becomes available, 50% abatement is reached 45 years later than in cases where SRM remains impossible. For current policy, however, the most important point is that in 2015 ("now"), the 30% chance of SRM becoming available does not lead to significant reduction in optimal abatement.

On the other hand, the performance ζ of this scenario is 125% (Table 3), significantly higher than for abatement-only. The lowest 10-th percentile performance, ζ_{10} is very similar to the Abatement-only scenario. In the Realistic Storyline, the low-performance members are those in which SRM never becomes available, and the policy (i.e. trajectory of abatement) in these runs is very similar to Abatement-only. However, ζ_{90} is much higher for the Realistic Storyline than for Abatement-only. This measure is dominated by those members in which SRM becomes available. The total climate tipping risk for the Realistic Storyline is 30.1%, compared to 37.8% in the Abatement-only scenario. The SCC for the Realistic Storyline is $37\$/t(C)$, 12% lower than for Abatement-only.

These comparisons between the Realistic Storyline and Abatement-only indicate that the former performs better. In other words, a policy maker in 2015 should not dismiss SRM prematurely, but keep the option open (by encouraging research and development). If we are lucky and SRM works well, it can greatly enhance future welfare, whereas if it never becomes feasible, we are not worse off than with abatement-only. (Note, however, that we did not include the possibility of a large-scale SRM test with huge unexpected damage, but assumed careful, well-designed research.) So research in SRM should be continued, but not at the cost of strong efforts on abatement.



Scenario	Abate 50%	peak SRM	SCC
Abatement-only, standard	2095	/	41
Ab.+SRM, standard	2139	35.0	20
Abatement-only, low rate of pure time preference ($\rho = 0.5\%$)	2068	/	70
Ab.+SRM, low rate of pure time preference ($\rho = 0.5\%$)	2116	31.1	30
Ab.+SRM, less temp. damage, more precip.damage ($\psi_T \rightarrow \psi_T/2, \psi_P \rightarrow \psi_P \times 2$)	2143	32.6	17
Ab.+SRM, double damage from tipping ($\Omega = 0.8$)	2136	34.8	20
Ab.+SRM, double climate tipping probability ($\kappa_{tipp} \rightarrow \kappa_{tipp} \times 2$)	2137	34.9	20
Ab.+SRM, quadrupled SRM failure probability ($\kappa_{fail} \rightarrow \kappa_{fail} \times 4$)	2121	34.3	23
Ab.+SRM, double damage from SRM ($\psi_S \rightarrow \psi_S \times 2$)	2133	26.8	22

Table 4. Policy metrics of the sensitivity runs. ‘Abate 50%’ is the year in which Abatement reaches 50% ($\mu = 0.5$). ‘peak SRM’ (in $Mt(S)/yr$) is the highest SRM value of the ensemble (over all times and all members) and corresponds to those ensemble members without early SRM failure or climate tipping. ‘SCC’ is the social cost of carbon in $\$(2005)/t(C)$. All simulations were performed in the stochastic settings and are either Abatement-only or Abatement+SRM (abbreviated here as Ab.+SRM). The first two cases, labelled ‘standard’, are repeated from Table 3 for convenience. The sensitivity runs correspond to those discussed in Sect. 3.4.

3.4 Sensitivity Analysis

The results of our model substantially depend on the rate of time preference ρ (see Sect. 2.2). In Abatement-only, reducing ρ from the standard value of 1.5% to a lower value of 0.5% will lead to stronger abatement: 50% abatement is reached 27 years earlier (see Table 4). This is expected, as a lower rate of time preference means that the decision maker gives more weight to the welfare of future generations and is more willing to sacrifice present consumption to reduce climate change. The SCC rises from $\$41/t(C)$ to $\$70/t(C)$. Interestingly, abatement also increases in the Abatement+SRM scenario (50% abatement is reached 23 years earlier) when reducing ρ to 0.5%, while the peak SRM (definition: see Table 3) decreases by about 11%. In other words, a decision maker who cares strongly about the future will choose to reduce CO₂ emissions rather than forcing future generations to rely on SRM, which causes damages and might fail. The SCC rises from $\$20/t(C)$ to $\$30/t(C)$.

The distribution of the damages between the two major contributors, namely warming and residual climate change, was chosen rather arbitrarily. However, halving ψ_T (warming contribution) and doubling ψ_P (residual contribution) does not qualitatively affect our results. 50% abatement is reached 4 years later in the Abatement+SRM scenario, and SRM peaks at 32.6Mt(S)/year instead of 35.0Mt(S)/year, i.e. the optimal policy still combines a similar abatement with modest SRM. The SCC drops from $\$20/t(C)$ to $\$17/t(C)$. Doubling the damages associated with climate tipping ($\Omega = 0.9 \rightarrow \Omega = 0.8$) only accelerates 50% abatement by 3 years in the Abatement+SRM case and the SCC remains at $\$20/t(C)$. Doubling the likelihood of climate tipping (κ_{tipp}) accelerates 50% abatement by 2 years and likewise does not affect SCC. Increasing the failure probability of SRM (κ_{fail}) by a factor of 4, i.e. such that SRM failure occurs in 80% of the ensemble members rather than 20%, increases the SCC only by 15% in the Abatement+SRM scenario, i.e. from $\$20/t(C)$ to $\$23/t(C)$. The reason is that the



likelihood of SRM failure in the first decades, which are least discounted, is still fairly small. The peak SRM is reduced only by 2%: As long as SRM is available, it is used despite high failure probability. 50% abatement is reached in 2121, rather than 2138, in the ensemble mean. Doubling the damage associated with SRM (i.e. doubling ψ_S) accelerates 50% abatement by 6 years and the SCC rises from $\$20/t(C)$ to $\$22/t(C)$. The peak SRM is reduced by about 23%, to 26.8Mt(S)/yr. Hence even if SRM is twice as damaging as assumed in the standard case, the optimal policy still employs modest SRM as complement to abatement.

To summarise, changes in the damage function and/or likelihood of stochastic events do not qualitatively affect the optimal policy in the Abatement+SRM scenario, which consists of a combination of reasonably high abatement (delayed by a few decades w.r.t Abatement-only in the standard settings) and modest SRM.

10 4 Summary and Discussion

In this paper, we present the first cost-benefit analysis of SRM under uncertainty performed with a rigorous optimisation approach (dynamic programming). From our analysis we draw two conclusions. First, sulphate SRM has the potential to greatly enhance future welfare and should therefore be taken seriously as possible policy option. Second, even if successful, SRM does not replace CO₂ abatement, but complements it.

15 Compared to previous studies (Goes et al., 2011; MorenoCruz and Keith, 2013; Heutel et al., 2018), our results are more optimistic about SRM, which seems partly due to the improved methodology we adopted. For instance, demonstrating that welfare is severely impacted if the decision maker makes wrong assumptions on the SRM-related damages (Bahn et al., 2015) is not a consistent cost-benefit analysis. The analysis by Goes et al. (2011) only considers a full replacement of abatement by SRM, rather than a complementary approach. Compared to Heutel et al. (2018), we find a much stronger reduction in the SCC. 20 However, as discussed previously, their model and optimisation method differ in some crucial points from ours.

Our results should not be interpreted as precise policy recommendations to set, say, exact values of the SCC, as our model is too limited to offer more than a qualitative exploration and comparison of simple scenarios. For example, uncertainty in the climate system is limited to one tipping point, while uncertainty in the climate sensitivity is ignored. Our climate model is based on linear response theory, and although this approach captures many climate feedbacks adequately, it does not capture 25 the possible dependence of the response on the background state, e.g. a saturation of carbon sinks (Aengenheyster et al., 2018).

A controversial component in Integrated Assessment Models such as DICE is the quantification of climate damages (Howard, 2014; van den Bergh and Botzen, 2015), which is highly aggregated and based on very limited data. We introduced additional parameters to the damage function by making a plausible, but rather ad-hoc attribution of climate damages to temperature, global precipitation ('residual climate change') and CO₂ concentrations. Also little is known about the size of ecological, let 30 alone economic, damages associated with SRM. Gaining a better understanding of these damages is essential for conducting a meaningful cost-benefit analysis and ultimately determining a climate policy, hence it should be given a high priority.

The economic component of our model also has limitations: it is assumed that a sudden strong increase in abatement is possible after SRM failure, which might be not feasible in reality. Like DICE, our model has an exogenous innovation and thus



ignores learning-by-doing effects in abatement (and SRM). It also does not capture the fact that the energy transition is partly an investment (e.g. restructuring the energy grid), i.e. abatement costs should reduce once the infrastructure for green energy is set up. These effects may bias our results against abatement.

Despite the large scientific and political uncertainties which need to be overcome, we believe that one cannot afford to dismiss SRM at the current stage, as it has the potential to greatly reduce climate risk and enhance future welfare. However, the scientific uncertainties, especially concerning efficiency and damages of SRM, must be quantified, and the uncertain prospect of SRM becoming available should not tempt us to reduce abatement.

Code availability.

Data availability. The code used (described in the Methods section) is available upon request from the corresponding author.

10 *Code and data availability.*

Sample availability.

Video supplement.

Appendix A: Solving the GeoDICE model

A1 Terminal function

15 Unrealistic behaviour occurs in the last time steps of an optimisation problem, because the decisions made do not influence the future anymore (as the future is not simulated). To avoid this problem, we follow Cai et al. (2016) and run the optimisation over 600 years, while only considering the first 400 as actual simulation and the final 200 years as ‘terminal function’. During termination, tipping can still occur and SRM can be freely chosen, while abatement is set to 1. Due to discounting, the trajectory after 600 years has little relevance for the optimal policy during the first 400 years. Indeed, prolonging the runs to 800 years
20 had a negligible effect on policies during the first 400 years.



A2 Optimisation method

The social planner problem aims at finding the policy that maximises the expected cumulative discounted utility. To solve this problem in the stochastic setting, we apply dynamic programming (Bellman, 1957). This methodology relies on the concept of the value function to obtain the optimal policy via backward reduction. As our state space is continuous and no analytic solution is available, we are forced to adopt some approximation scheme to represent the value function at each time step. Following Cai et al. (2016), we use a Chebyshev approximation, which is well suited for parallelisation. The Chebyshev polynomial is obtained by solving a small optimisation problem at each of a finite number of regularly spaced Chebyshev approximation nodes. We used a fourth-degree Chebyshev polynomial with five approximation nodes per continuous dimension. In combination with the binary state variables for the tipping point and SRM failure, this results in 312500 approximation nodes per time step. This method is developed and discussed extensively in the work by Cai (2009); Cai et al. (2012a, 2016). For a complete overview we refer the reader to these papers and the references therein. Here we outline the methodological choices specific to the present application: the boundaries used for the domain of the Chebyshev polynomial, and adjustments to the value function approximation to accommodate the asymmetry and non-smoothness of the true value function. Additionally, we examine the accuracy of this methodology when applied in the current setting.

15 A2.1 Boundaries

In order to define the Chebyshev approximation nodes, we must first set the boundaries of the region of state space in which we are interested. To do this, we calculate three trajectories in the deterministic model: first, the optimal trajectory (obtained by optimising the whole system in all decision variables with standard deterministic optimisation software); second, a ‘high-emission’ trajectory calculated by setting mitigation and SRM to zero for the whole run; and third, a ‘low-emission’ trajectory, calculated by setting mitigation to one and SRM to zero for the whole run. We subsequently take as domain boundaries for each variable the minimum and maximum over these three trajectories, with an additional margin of minus, respectively plus 30% of these values. For all experiments, we check that all the sample paths in the ensemble remained well within the boundaries of the domain. For approximation nodes close to the boundaries, it will still be possible to select actions that may bring the system outside the boundaries in the next step. Since a Chebyshev polynomial cannot be extrapolated outside its domain, we first project the state onto the region of interest before evaluating the approximate value function.

A2.2 Value Function Smoothing

In the current setting, directly using a Chebyshev polynomial to approximate the value function gives poor results because the value function exhibits an asymmetry and a non-smoothness that a low-degree Chebyshev polynomial cannot capture. The discontinuity is caused by the fact that in states with positive temperatures, SRM is available to reduce them, while in states with negative temperatures this is impossible: therefore, positive temperature deviations are preferred over negative temperature deviations of equal magnitude. This problem is resolved by allowing ‘reverse’ SRM, which generates a radiative forcing of the same magnitude but opposite sign as regular SRM. Allowing such actions changes the value of certain states, thus removing



the asymmetry. This is a purely mathematical construct (we do not assume such reverse SRM is actually possible): the states with modified values are never reached in actual trajectories, and are only considered in the first place because the domain of the Chebyshev approximation must be a hypercube.

The non-smoothness results from the fact that the tipping point can only be crossed after a certain threshold is reached: this generates a discontinuity in the first derivative of the value function. This is resolved by fitting two separate Chebyshev polynomials to the two parts of the value function.

A2.3 Accuracy

We test the accuracy of our optimisation by comparing the resulting policy in a deterministic setting to the policy obtained by regular non-linear optimisation. The difference in action and trajectory is $< 3\%$, while the difference in the SCC is $< 2\%$. For the scenario in which only abatement is allowed, errors are lower ($0.1\text{--}1\%$ for actions and SCC, $0.01\text{--}0.1\%$ for trajectories), which is in line with the accuracy reported by Cai et al (Cai et al., 2016). Good accuracy in the deterministic setting may not generalise to the stochastic setting when the stochasticity itself introduces issues. To guard against this problem we ensure that the value function approximation fits well to the actual value function samples obtained at each time step.

Author contributions. All authors conceived research. C.E.W. and K.G.H. designed model and scenarios and interpreted results. K.G.H. performed simulations. All authors contributed to writing the paper.

Competing interests. The authors declare no conflict of interest.

Disclaimer.

Acknowledgements. Claudia E. Wieners is supported by the Complexity Lab Utrecht (CLUe) of the Centre for Complex Systems Studies at Utrecht University. Henk A. Dijkstra acknowledges support by the Netherlands Earth System Science Centre (NESSC), financially supported by the Ministry of Education, Culture and Science (OCW), Grant no. 024.002.001.



References

- Ackerman, F. Debating Climate Economics: The Stern Review vs. Its Critics. Report to Friends of the Earth-UK, 1–25 (2007)
- Aengenheyster, M., Feng, Q., van der Ploeg, F., & Dijkstra, H.A. The point of no return for climate action: effects of climate uncertainty and risk tolerance. *Earth System Dynamics* Vol. 9 Issue 3, p1085-1095. 11p.
- 5 Ahlm, L., Jones, A., Stjern, C. W., Muri, H., Kravitz, B., and Kristjánsson, J. E.: Marine cloud brightening – as effective without clouds, *Atmos. Chem. Phys.*, 17, 13071–13087 (2017).
- Andrews, T., Forster, P. M., Boucher, O., Bellouin, N., & Jones A. Precipitation, radiative forcing and global temperature change, *Geophys. Res. Lett.*, 37, L14701, 701 (2010)
- Bahn, O., Chesney, M., Gheysens, J., Knutti, R. & Pana, A. C. Is there room for geoengineering in the optimal climate policy mix?
10 *Environmental Science and Policy* 48, 67–76 (2015).
- Bellman, R. *Dynamical Programming*. Princeton University Press, Princeton, NJ, USA (1957).
- Brovkin, V., Petoukhov, V., Claussen, M., Bauer, E., Archer, D., & Jaeger, C. Geoengineering climate by stratospheric sulfur injections: Earth system vulnerability to technological failure, *Clim. Change*, 92, 243–259 (2009)
- Cai, Y. *Dynamic Programming and its Application in Economics and Finance* PhD Thesis (Stanford University, 2009).
- 15 Cai, Y., Judd, K. L. & Lontzek, T. S. DSICE: A Dynamic Stochastic Integrated Model of Climate and Economy. SSRN Electronic Journal (2012).
- Cai, Y., Judd, K. L. & Lontzek, T. S. . Continuous-Time Methods for Integrated Assessment Models. Working Paper 18365, National Bureau of Economic Research. URL: <http://www.nber.org/papers/w18365> (2012)
- Cai, Y., Judd, K.L., Lenton, T.M., Lontzek, T.S., & Narita, D. Environmental tipping points significantly affect the cost-benefit assessment
20 of climate policies. *Proc Natl Acad Sci* 201503890 (2015)
- Cai, Y., Lenton, T. M. & Lontzek, T. S. Risk of multiple interacting tipping points should encourage rapid CO₂ emission reduction, *Nature Climate Change* 6, 520–525 (2016)
- Ciais, P., C. Sabine, G. Bala, L. Bopp, V. Brovkin, J. Canadell, A. Chhabra, R. DeFries, J. Galloway, M. Heimann, C. Jones, C. Le Quéré, R.B. Myneni, S. Piao and P. Thornton, 2013: Carbon and Other Biogeochemical Cycles. In: *Climate Change 2013: The Physical Science Basis*. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Crutzen, P. J. Albedo enhancement by stratospheric sulfur injections: A contribution to resolve a policy dilemma? *Clim. Change*, 77, 211–219 (2006)
- 30 Effiong, U. & Neitzel, R. L. Assessing the direct occupational and public health impacts of solar radiation management with stratospheric aerosols, *Environ. Health*, 15(1), 1–9 (2016)
- Gabriel, C. J., Robock, A., Xia L., Zambri, B., and Kravitz, B. The G4Foam experiment: Global climate impacts of regional ocean albedo modification, *Atmos. Chem. Phys.*, 17, 595–613 (2017)
- Goes, M., Tuana, N. & Keller, K. The economics (or lack thereof) of aerosol geoengineering. *Climatic Change* 109, 719–744 (2011).
- 35 Heutel, G., Moreno-Cruz, J. & Shayegh, S. Climate tipping points and solar geoengineering. *Journal of Economic Behavior and Organization* 132, 19–45 (2016).



- Heutel, G., Moreno-Cruz, J. & Shayegh, S. Solar Geoengineering, Uncertainty, and the Price of Carbon. *Journal of Environmental Economics and Management* 87, 24–41 (2018)
- Howard, P. Omitted Damages: What's missing from the Social Cost of Carbon, 2014. Available at https://costofcarbon.org/files/Omitted_Damages_Whats_Missing_From_the_Social_Cost_of_Carbon.pdf
- 5 Irvine, P. J. et al. Towards a comprehensive climate impacts assessment of solar geoengineering. *Earth's Future* 5, 93–106 (2017)
- Joos, F. et al. Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis. *Atmos. Chem. Phys.* 13, 2793–2825 (2013)
- Keller, D.P., Feng, E.Y. & Oschlies, A. Potential climate engineering effectiveness and side effects during a high carbon dioxide-emission scenario. *ncomms* 5, 3304 (2014)
- 10 Kleinschmitt, C., Boucher, O. & Platt, U. Sensitivity of the radiative forcing by stratospheric sulphur geoengineering to the amount and strategy of the SO₂ injection studied with the LMDZ-S3A model. *Atmos. Chem. Phys.*, 18, 2769–2786 (2018)
- Kravitz, B. et al. Climate model response from the Geoengineering ModelIntercomparison Project (GeoMIP), *J. Geophys. Res. Atmos.* 118, 8320–8332 (2013)
- Kravitz, B., MacMartin, D. G., Rasch, P. J., & Jarvis, A. J. A new method of comparing forcing agents in climate models. *J. Climate*, 28, 8203–8218 2015
- 15 Latham, J. et al. Global temperature stabilization via controlled albedo enhancement of low-level maritime clouds. *Phil. Trans. R. Soc. A* 366, 3969–3987 (2008)
- Le Quéré, C. et al. Global Carbon Budget 2017, *Earth Syst. Sci. Data*, 10, 405–448 (2018)
- Lilley, P. The Failings of the Stern Review of the Economics of Climate Change. The Global Warming Policy Foundation (GWPF) Report 20 92012
- MacMartin, D. G., & Kravitz, B. Dynamic climate emulators for solar geoengineering. *Atmospheric Chemistry and Physics*, 16(24), 15789–15799 (2016)
- Matthews, H. D., & Caldeira, K. Transient climate-carbon simulations of planetary geoengineering, *Proc. Natl. Acad. Sci. U. S. A.*, 104, 9949–9954 (2007)
- 25 McClellan, J., Keith, D. & Apt, J. Cost analysis of stratospheric albedo modification delivery systems *Environ. Res. Lett.* 7 034019 (2010)
- Moreno-Cruz, J. B. & Keith, D. W. Climate policy under uncertainty: A case for solar geoengineering. *Climatic Change* 121, 431–444 (2013).
- Moriyama, R., Sugiyama, M., Kurosawa, A. et al. The cost of stratospheric climate engineering revisited. *Mitig Adapt Strateg Glob Change* 22: 1207 (2017)
- Niemeier, U. & Timmreck, C. What is the limit of climate engineering by stratospheric injection of SO₂? *Atmos. Chem. Phys.*, 15, 9129–9141 30 (2015)
- Niemeier, U. & Schmidt, H. Changing transport processes in the stratosphere by radiative heating of sulphate aerosols. *Atmos. Chem. Phys.*, 17, 14871–14886, (2017)
- Nordhaus, W. D. The “DICE” Model: Background and Structure of a Dynamic Integrated Climate–Economy Model of the Economics of Global Warming. Cowles Foundation Discussion Paper No. 1009. Cowles Foundation for Research in Economics: New Haven, CT (1992)
- 35 Pitari, G., et al. Stratospheric ozone response to sulphate geoengineering: Results from the Geoengineering Model Intercomparison Project (GeoMIP), *J. Geophys. Res. Atmos.*, 119, 2629–2653 (2014)
- Robock, A. Volcanic eruptions and climate. *Reviews of Geophysics*, 38, 191–219 (2000)



- Robock, A., Marquardt, A., Kravitz, B., & Stenchikov, G. Benefits, risks, and costs of stratospheric geoengineering. *Geophys Res Lett* 2009, 36:L19703 (2009)
- Seneviratne, S.I. et al. Land radiative management as contributor to regional-scale climate adaptation and mitigation. *Nat. Geosci.* 11: 88–96 (2018)
- 5 L.L. Stowe, R.M. Carey, P.P. Pellegrino, Monitoring the Mt. Pinatubo aerosol layer with NOAA/11 AVRHH DATA, *Geophys. Res. Lett.* 19/2 159 (1992)
- Stenchikov, G. L., et al. Radiative forcing from the 1991 Mount Pinatubo volcanic eruption, *J. Geophys. Res.*, 103, 13837–13857 (1998)
- Stern, V. et al. *The Stern Review*. Government Equalities Office, Home Office (2007)
- Thompson, W. J., Wallace, J. M., Jones, P. D. & Kennedy, J. J. Identifying Signatures of Natural Climate Variability in Time Series of
10 Global-Mean Surface Temperature: Methodology and Insights *J. Clim.* 22, 6120–6141 (2009)
- Tilmes, S. et al. CESM1(WACCM) Stratospheric Aerosol Geoengineering Large Ensemble (GLENS) project. *Bull. Am. Meteorol. Soc.* (2018).
- Tjiputra, J. F., Grini, A. & Lee, H. Impact of idealized future stratospheric aerosol injection on the large-scale ocean and land carbon cycles, *J. Geophys. Res. Biogeosci.* 121, 2–27 (2016).
- 15 Trisos, C. H. et al. Potentially dangerous consequences for biodiversity of solar geoengineering implementation and termination. *Nat. Ecol. Evol.* 2, 475–482 (2018)
- UNFCCC(2015), Adoption of the Paris Agreement, United Nations Framework Convention on Climate Change, United Nations Office, Geneva, Switzerland.
- J.C.J.M. Van den Bergh, W.J.W. Botzen Monetary valuation of the social cost of CO2 emissions: a critical survey *Ecol. Econ.*, 114, pp. 33-46
20 (2015)
- Visioni, D., Pitari, G., & Aquila, V. Sulfate geoengineering: A review of the factors controlling the needed injection of sulfur dioxide. *Atmospheric Chemistry and Physics*, 17(6), 3879–3889 (2017)
- Ward 2009: sulphur dioxide initiates global climate change in four ways *Thin Solid Films* 517, 3188–3203 (2009)
- World Bank (<https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>)