Climate policy under uncertainty:
a case for solar geoengineering

Juan B. Moreno-Cruz · David W. Keith

Abstract Solar Radiation Management (SRM) has two characteristics that make it useful for managing climate risk: it is quick and it is cheap. SRM cannot, however, perfectly offset CO_2-driven climate change, and its use introduces novel climate and environmental risks. We introduce SRM in a simple economic model of climate change that is designed to explore the interaction between uncertainty in the climate's response to CO_2 and the risks of SRM in the face of carbon-cycle inertia. The fact that SRM can be implemented quickly, reducing the effects of inertia, makes it a valuable tool to manage climate risks even if it is relatively ineffective at compensating for CO_2-driven climate change or if its costs are large compared to traditional abatement strategies. Uncertainty about SRM is high, and decision makers must decide whether or not to commit to research that might reduce this uncertainty. We find that even modest reductions in uncertainty about the side-effects of SRM can reduce the overall costs of climate change in the order of 10%.

1 Introduction

It appears to be technically feasible to engineer an increase in albedo, a planetary brightening, as a means to offset the warming caused by carbon dioxide (CO_2) and other greenhouse gases through Solar Radiation Management (SRM) (Keith and Dowlatabadi 1992; Keith 2000; Crutzen 2006; Shepherd et al. 2009). However, the cooling produced by SRM does not exactly compensate for the warming caused by

J. B. Moreno-Cruz (✉)
School of Economics, Georgia Institute of Technology,
221 Bobby Dodd Way, Atlanta, GA 30332, USA
e-mail: juan.moreno-cruz@econ.gatech.edu

D. W. Keith
Kennedy School of Government and School of Engineering and Applied Sciences,
Harvard University, 29 Oxford Street, Cambridge, MA 02138, USA

CO_2-driven climate change; and any particular method of SRM will no doubt entail other risks and side-effects (e.g. Bala et al. 2008; Ricke et al. 2010). Nevertheless, SRM may be a useful tool to mange climate risks (Wigley 2006). In this paper we ask how optimal policy is affected by risk regarding the side-effects of SRM, in the face of uncertainty about the magnitude of the damages caused by CO_2-driven climate change.

To answer this question we construct a simple model that captures the following stylized facts about climate change and SRM:

1. *The carbon-climate system has inertia.* There is a lag between the response of the climate system and the anthropogenic carbon emissions that cause climate change. The inertia of the carbon-climate system makes it impossible to quickly reduce climate risk by reducing emissions, as it is expected that 40% of the peak concentration of CO_2 will remain in the atmosphere 1000 years after the peak is reached (Solomon et al. 2009).

2. *Climate change damages are uncertain.* The amount of climate change resulting from a given emissions trajectory is uncertain, as are the resulting economic (or other) damages. Moreover, this uncertainty is irreducible over a timescale of decades during which we will make near-term decisions about emissions abatement (Morgan and Keith 1995; Zickfeld et al. 2010).

3. *SRM is fast.* A reduction in the incoming radiation has relatively instantaneous effects on global temperature (Caldeira and Matthews 2007; Robock et al. 2008). Nature gives an example of how quickly temperature responds to changes in radiative forcing: after Mount Pinatubo's explosion around 20TgS were deposited in the stratosphere, global surface temperatures cooled about 0.5°C over the following year (Soden et al. 2002).

4. *SRM is inexpensive.* At this stage, little is known about the technical costs of SRM, but some preliminary studies have suggested that SRM could offset the increase in global average temperature due to CO_2 at a cost 10 to 1000 times lower than achieving the same outcome by cutting emissions (McClellan et al. 2010; Robock et al. 2009; Shepherd et al. 2009).

5. *SRM cannot eliminate carbon-climate risk.* SRM technologies can intervene to restore the surface temperature by reducing the incoming solar radiation. This intervention, however, cannot eliminate all the damages caused by climate change. In particular, the temperature compensation has a different regional distribution, that leaves the poles under compensated while the equator is over compensated (Caldeira and Matthews 2007). Moreover, the accumulation of greenhouse gases has direct implications on the precipitation patterns (Allen and Ingram 2002); and, in the case of CO_2, ocean acidification (Caldeira and Wickett 2003, 2005).

6. *SRM introduces damages.* There is an increase in the risks of destruction of stratospheric ozone due to SRM implementation (Solomon 1996, 1999). Moreover, sulfuric acid deposition may create health and regional problems (Crutzen 2006); although recent literature suggests these effects are small (Kravitz et al. 2009). Also, recent numerical simulations show that SRM will affect precipitation patterns and volumes, possibly causing droughts in large regions of the planet (e.g. Ricke et al. 2010).

Our goal is to explore the trade-offs between the advantages and disadvantages of SRM in a cost minimizing optimal-decision framework that we intend to be as simple as possible while still capturing all of these stylized facts. The advantages of SRM to manage climate risks are twofold. First, it is inexpensive compared to abatement, and second it allows rapid action avoiding some of the inertia of the carbon system. The corresponding disadvantages of SRM are that it imperfectly compensates for CO_2 driven warming and it may introduce new environmental risks.

In our model, the objective of the decision-maker is to minimize the expected total costs of managing climate change. The costs of climate change are the sum of the costs of abatement and SRM activities plus any economic damages. The costs of abatement and SRM are increasing and convex functions of their arguments, while economic damages are the sum of the damages arising from greenhouse gas concentrations—such as temperature changes and ocean acidification—and those arising from the side-effects of SRM. The damages from temperature are a quadratic function of the change in global surface temperature; which, in turn, is proportional to *radiative forcing*. The damages from ocean acidification arise due to the increase of CO_2 concentrations in the oceans; which, in turn, affects marine life and the economic activities associated with it, i.e. fishing and tourism. The damages arising from the side-effects of SRM are assumed to be a quadratic function of the total level of SRM.

As a simple way to capture climate-carbon inertia we use a two-stage decision framework in which the abatement decisions are made in the first period and SRM decisions are made in the second. In between periods, the decision maker learns the true sensitivity of the climate (Fig. 1). Because temperature depends on cumulative emissions, we assume emissions are irreversible (Matthews et al. 2009) and in that sense, only the level of abatement implemented before learning about the sensitivity of the climate system can help reduce damages caused by temperature changes and ocean acidification. The climate system, however, responds quickly to changes in radiative forcing in the form of SRM. This quickness of response allows SRM to reduce temperature damages after learning about the sensitivity of the climate; hence, eliminating the inertia associated with other forms of climate intervention and abatement. Damages take place in the second period, after SRM decisions are made.

The approach in this paper has proven to be useful for the economic analysis of climate change and we expect it to be equally insightful for the economic analysis of SRM [see Weitzman (2009) for a recent application of a two period model to analyze climate change policy, and Goulder and Mathai (2000) for an example of the use of a cost minimizing framework with increasing and convex costs to analyze climate policy]. Five caveats, however, are important for our analysis:

- The optimal policy assumes a centralized decision maker. In practice, many countries will decide how to implement SRM amongst themselves. The strategic interaction among countries may lead to under or over provision of either SRM or abatement (see Millar-Ball 2012; Moreno-Cruz 2011).
- A centralized decision maker minimizes changes in global mean temperature and other damages at a global scale. By making this assumption, we eliminate all
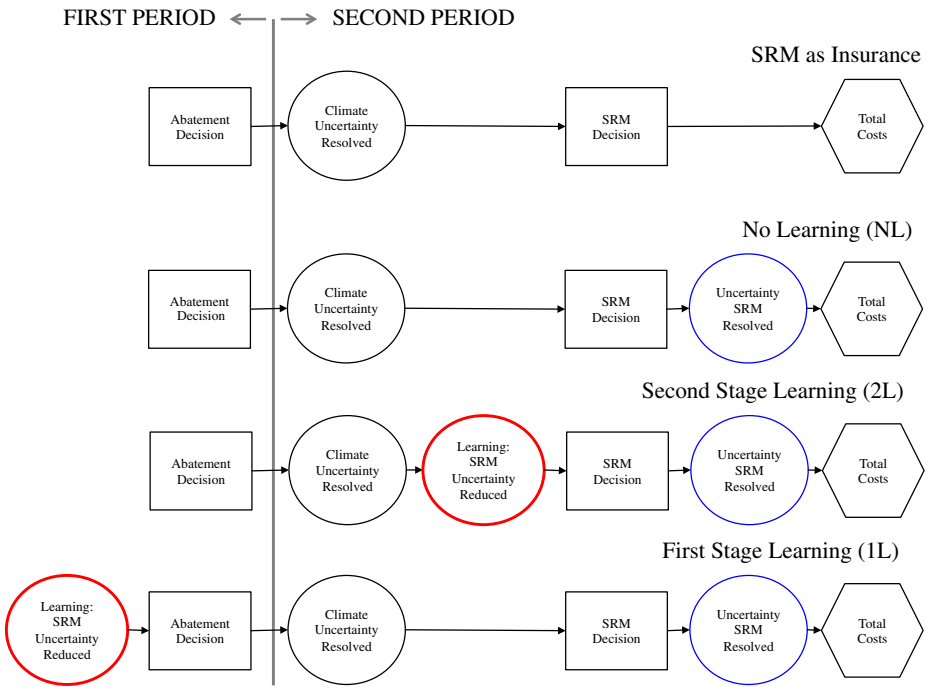
FIRST PERIOD ←——|——→ SECOND PERIOD

SRM as Insurance

Abatement Decision → Climate Uncertainty Resolved → SRM Decision → Total Costs

No Learning (NL)

Abatement Decision → Climate Uncertainty Resolved → SRM Decision → Uncertainty SRM Resolved → Total Costs

Second Stage Learning (2L)

Abatement Decision → Climate Uncertainty Resolved → Learning: SRM Uncertainty Reduced → SRM Decision → Uncertainty SRM Resolved → Total Costs

First Stage Learning (1L)

Learning: SRM Uncertainty Reduced → Abatement Decision → Climate Uncertainty Resolved → SRM Decision → Uncertainty SRM Resolved → Total Costs

Fig. 1 Timing of decisions and information in the two-period model. The top schematic shows the timing when SRM is used as an insurance in Section 3. The bottom three schematics show the scenarios in Section 4. They vary depending on whether decisions are made before or after learning about the SRM damages. Decisions are represented by *rectangles*, while uncertain outcomes are represented by *circles*. SRM uncertainty is represented with *blue circles*. Learning is represented with *red circles*. Payoffs are represented by *hexagons*. The first schematic shows the timing of decisions when there is no learning (NL scenario). The second schematic describes the scenario when learning takes place before SRM decisions are made, but after abatement decisions are made (2L scenario). The third schematic describes the scenario when learning takes place before abatement and SRM decisions are made (1L scenario)

considerations to regional inequalities that may arise from the implementation of SRM [see Moreno-Cruz et al. (2012) and Ricke et al. (2010) for a detailed treatment of the inequalities introduced by SRM]. Nonetheless, understanding the optimal policy is important as it serves as a benchmark with which other policies may be compared.

- Use of a static two-period model limits its application in two important aspects. First, it cannot be used to analyze time dependent optimal policies where SRM is introduced incrementally. Instead we concentrate on SRM as a tool to deal with low-probability high-consequence impacts that are colloquially referred to as "climate emergencies". Second, the model cannot address damages due to rapid temperature changes associated to the sudden interruption of an SRM program.

- By considering damages only in terms of reduction in economic output we are neglecting aspects of the problem that do not easily fit this framework such as non-monetary environmental values. This assumption, of course, neglects the

ethical issues associated with the direct manipulation of the climate implied by SRM. We believe these ethical concerns are crucial for the analysis of SRM as are the issues of uncertainty and inertia that we treat here.

The rest of the paper proceeds as follows. In Section 2 we introduce and calibrate the model. In Section 3 we introduce uncertainty on the climate system and analyze the role of SRM in dealing with high-impact, low-probability outcomes. In Section 4, we deal with the uncertainty attached to the damages from SRM and analyze the value of reducing this uncertainty. We draw conclusions in Section 5.

2 A general description of the model

2.1 Temperature, abatement and SRM

When the concentration of greenhouse gases increases in the atmosphere it alters the balance between incoming solar radiation and outgoing terrestrial radiation, resulting in an increase in the mean global temperature of Earth. *Radiative forcing* describes how the radiation balance is altered by human activity. Radiative forcing, R, is a function of the concentration of CO_2 in the atmosphere, S, relative to the preindustrial level, S_0:

$$R = \beta \, \ln \left(S/S_0 \right) \tag{1}$$

where, according to the IPCC (2007), $\beta = 5.35$ watts-per-meter-squared [Wm^{-2}]. Abatement, which we denote by A, refers to measures that reduce the concentration level of CO_2 in the atmosphere. In particular, assume that $S = S_{BAU} - A$, where S_{BAU} is the business as usual concentration of CO_2 in the atmosphere measured in parts per million [ppm].

Changes in mean global temperature, ΔT—measured in °C—are defined as a linear function of radiative forcing, R:

$$\Delta T = \lambda R \tag{2}$$

where λ is the climate sensitivity parameter with units °C m^2/W.

When SRM is introduced in the model, the relation between CO_2 concentrations and temperature is altered. We measure SRM, G, in terms of its radiative forcing potential and, since temperature change is a linear function of radiative forcing, Eq. 2 can be written as:

$$\Delta T(A, G) = \lambda \left(\beta \, \mathrm{Ln} \left(\frac{S_{BAU} - A}{S_0} \right) - G \right) \tag{3}$$

2.2 Economic damages

We represent total climate damages as the sum of impacts from three different sources: temperature, SRM and uncompensated CO_2 damages (e.g. ocean acidification). Following Nordhaus (2008), we assume temperature damages are quadratic. Following Brander et al. (2009), damages from ocean acidification are also quadratic on the concentration of CO_2. We assume that SRM damages are also a quadratic function

of the total level of SRM.[1] To be able to compare the different sources of impacts, we express damages in terms of reductions in economic output. Thus, total damages are given by:

$$D(A, G) = \eta_S(S_{BAU} - A)^2 + \eta_T\lambda^2 (\Delta T(A, G))^2 + \eta_G G^2 \qquad (4)$$

where $\eta_S(S_{BAU} - A)^2$ are the damages caused by ocean acidification and other uncompensated damages from CO_2, $\eta_T\lambda^2 (\Delta T(A, G))^2$ are damages caused by temperature changes, and $\eta_G G^2$ are the damages caused by the side-effects of SRM. In Eq. 4, when A equals S_{BAU} and G equals zero, damages are zero. However, when A is less that S_{BAU}, damages are always positive, showing the inability of SRM to perfectly compensate for greenhouse gas driven climate change (see bottom panel in Fig. 2). That is, although technically SRM can reduce temperature changes to zero, it may do so at the expense of other economic damages.

2.3 Implementation costs

We assume that abatement costs are increasing and convex. In particular, following Nordhaus (2008), we have:

$$\Lambda(A) = K_A A^{\alpha} \qquad (5)$$

where K_A has units [\$/ppm] and $\alpha = 2.8$.

Following Keith and Dowlatabadi (1992) we assume that SRM costs are linear and given by

$$\Gamma(G) = K_G G \qquad (6)$$

where K_G has units [\$/(Wm^{-2})].

Total social costs are the sum of the implementation costs, given by Eqs. 5 and 6, and the economic damages given by Eq. 4. The optimal policy consist of the level of abatement and the level of SRM that minimize total social costs.

2.4 Calibration

We use the year 2100 as our planning horizon, a common target in the analysis of climate change policy.[2] To calibrate our model, we use the projected costs and damages in 2100 reported by the DICE-2007 model (Dynamic Integrated Model of Climate and the Economy) (Nordhaus 2008). We complete the information needed for our calibration using data from the IPCC (2007) and publications related to the costs of SRM. The information given below is, unless otherwise noted, from Nordhaus (2008). The assumptions and calibrated values are summarized in Table 1.

We calibrate costs and damages as percentages of global GDP, when we report dollar values we assume global GDP to be around \$50 trillion per year (World Bank, World Development Indicators). Although not relevant for our study, incorporating

---

[1]There is not evidence of how steep the damages from SRM are. By choosing quadratic damages we are assuming they have the same weight as other climate related damages.

[2]As suggested by the reviewers, we analyzed the results for different target years, 50 years from now and 150 years. All the qualitative results are the same.
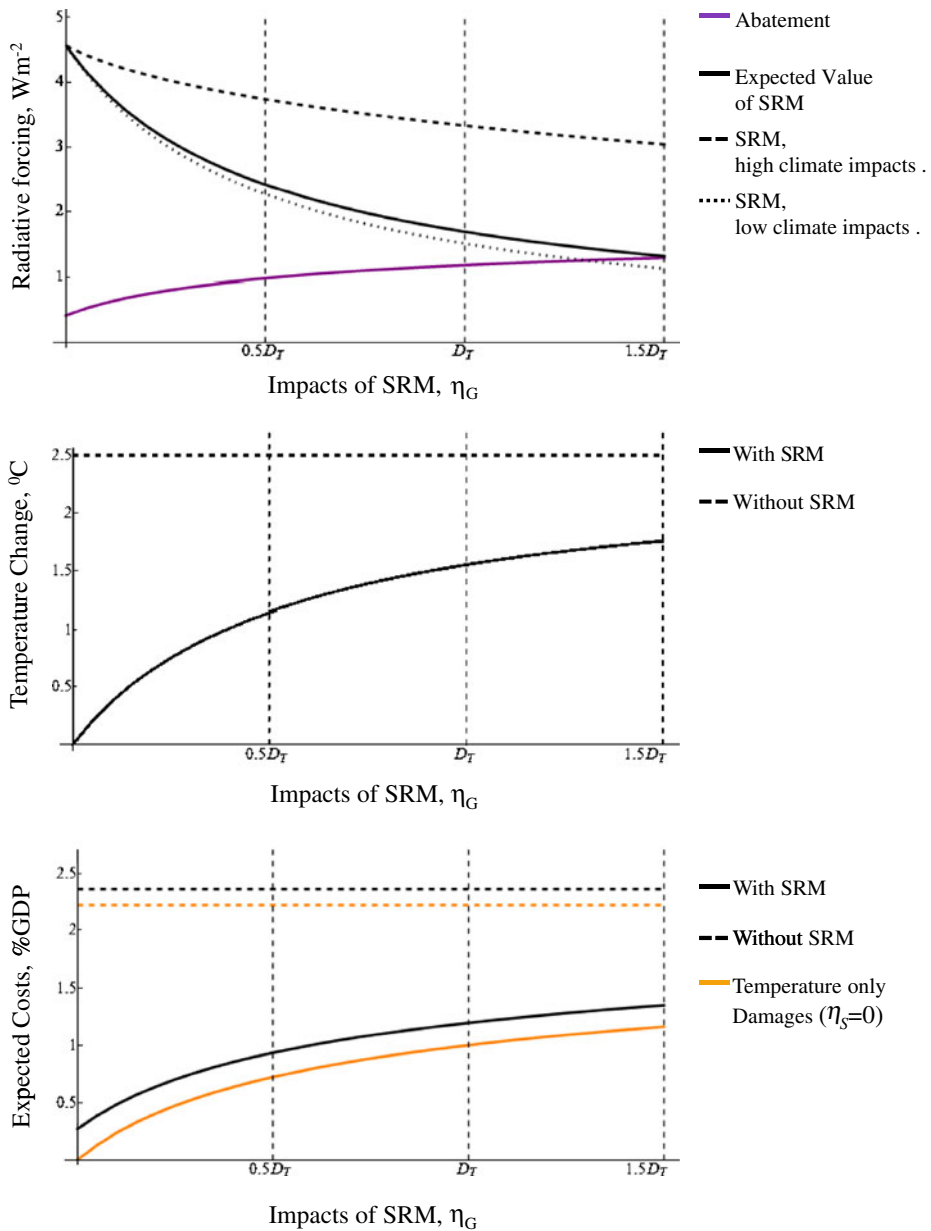
Fig. 2 Optimal climate policy. The horizontal axis is the impacts of SRM expressed as a fraction of the business-as-usual climate damages. For example, when $\eta_G = 0.5D_T$, the impacts of SRM are equivalent to 50% of the damages from CO_2-driven climate change. The vertical axis is in units of radiative forcing (Wm^{-2}). The *top panel* shows the optimal policy measured in terms of radiative forcing potential (Wm^{-2}). The *middle panel* presents the temperature change measured in °C. The *solid line* shows the results with SRM, and the *dashed line* shows the results without SRM. The *bottom panel* shows the expected costs of implementing the optimal policy as a fraction of global GDP. The *orange lines* show the expected total costs with only temperature damages. The difference between the *solid black line* and the *solid orange line* is the fraction of costs that cannot be compensated using SRM

Table 1 Calibration of model

Assumptions	
Climate parameters	Definitions
$S_{BAU} = 685$ ppm	Business as usual concentration of CO_2
$S_0 = 270$ ppm.	Preindustrial concentration of CO_2
$R_{BAU} = 4.98$ W/m^2	Business as usual radiative forcing
$T_{BAU} = 4.28°$C	Business as usual temperature change

Calibration	
Economic parameters	Notes
$K_A = 6.8 \times 10^{-7}$ [%GDP/ppm]	Reducing temperature change in 2100 by 1°C from BAU costs 1% of global GDP and reducing temperatures by 1°C relative to BAU requires a CO_2 reduction of 134 ppm.
$K_G = 5.7 \times 10^{-4}$ [%GDP/(Wm^{-2})]	We assume the costs of SRM are 1% of the costs of abatement (McClellan et al. 2010; Shepherd et al. 2009) and performed sensitivity analysis for values between 10% and 100% of the costs of abatement.
$\eta_T = 0.32$ [%GDP/°C]	If no action to deal with climate change is taken, around 3% of global GDP will be lost in 2100, or $D_T \equiv 0.03 GDP = \frac{1}{2}\eta_T T_{BAU}^2$
$\eta_S = 1.3 \times 10^{-6}$ [%GDP/ppm]	Ocean acidification damages add 10% to the total impacts from climate change (Brander et al. 2009), or $0.1 D_T = \frac{1}{2}\eta_S (S_{BAU} - S_0)^2$

discounting is simple. For example if we assume a discount rate of 1%, the yearly GDP value would be equivalent to $33 trillion. If we assume a discount rate of 7%, yearly GDP would be $7 trillion. Economic growth is equally easy to introduce. Introducing economic growth at a rate of 2.5% will yield a yearly GDP value of $200 trillion. Considering a discount rate close to the rate of economic growth would leave the yearly value of GDP at around $50 trillion.

There is insufficient information to allow us to quantify the risks of SRM, η_G, with any confidence, so we treat them parametrically. In Section 3 we analyze optimal policy as a function of η_G and in Section 4 we introduce uncertainty and learning on η_G.

3 Climate sensitivity uncertainty: SRM as insurance

In this section we analyze the role of SRM in dealing with the uncertainty surrounding the climate's response to changes in the atmospheric concentration of CO_2. Specifically, we made the climate sensitivity parameter, λ, random. We define the random variable $\widetilde{\lambda}$, to introduce the uncertainty of the response of the climate system. $\widetilde{\lambda}$ follows a binomial distribution of the form:

$$\widetilde{\lambda} = \begin{cases} \lambda_H = 2.3 \text{ with probability } p = 0.1 \\ \lambda_L = 0.7 \text{ with probability } 1 - p = 0.9. \end{cases} \tag{7}$$

Notice that the mean of this distribution is 0.86, which is consistent with recent estimates (IPCC 2007).[3] We choose this distribution of $\widetilde{\lambda}$ to capture the idea of low probability-high impact events that are characteristic of fat-tail distributions commonly associated to climate sensitivity (Roe and Baker 2007; Weitzman 2009). This is of course a simple approximation that allows us to introduce risk in the climate system without increasing the complexity of the model. The qualitative results of our paper would remain the same if we introduce a continuous distribution with fat-tails.

As we mentioned in the introduction, to capture climate-carbon inertia, decisions about abatement and SRM are made sequentially. Abatement decisions are made in the first period and SRM decisions are made in the second period. In between periods, the true climate sensitivity is revealed. Here SRM decisions are made under perfect information, but we will relax this assumption in Section 4 (see Fig. 1).

We introduce the imperfection of SRM parametrically; that is, the optimal level of abatement and the optimal level of SRM are a function of the magnitude of the side effects of SRM, η_G. We allow damages from SRM to be higher than those induced by CO_2-driven climate change, so $\eta_G \in [0, 1.5D_T]$ where $D_T = \$11.4 \times 10^{12}/(\text{Wm}^{-2})$. That is, when $\eta_G = D_T$, reducing temperature changes to zero using only SRM may create damages just as large as if temperature were equal to its business as usual level. By setting the upper limit at $\eta_G = 1.5D_T$ we try to highlight the role of SRM as an insurance. This limit, however, can be too high. The most commonly discussed direct impact of SRM (using stratospheric aerosols) is ozone loss. Estimates of the economic losses due to ozone depletion are in the order of US$1.1 trillion between 1987 and 2060. That is equivalent to 0.03% of global GDP, or $(1/100)D_T$ (Environment Canada 1997; Sunstein 2007).

The top panel in Fig. 2 shows the optimal policy. As expected, SRM is a decreasing function of η_G while abatement is increasing in η_G. Thus, abatement and SRM are *technical substitutes*: if SRM is costly, then it is optimal to implement more abatement. Also, the optimal level of SRM is always higher in the high-sensitivity outcome ($\lambda = \lambda_H$) compared to the low-sensitivity outcome ($\lambda = \lambda_L$). This is the result of the assumption that SRM can be chosen after learning about the climate sensitivity of the system. Moreover, in the case of an unlucky outcome, SRM is used more than abatement, even if damages from SRM are higher than D_T.

The middle panel in Fig. 2 shows temperature with and without SRM. We can see that temperature change increases when the damages from SRM increase. Temperature increases because there is a reduction in the level of SRM that is less than compensated by the increase in abatement levels; which results from the fact that abatement costs are increasing and convex.

The bottom panel in Fig. 2 shows the total costs of managing climate change as a function of the marginal damages from SRM, η_G. As expected, total costs are higher when damages from SRM become larger. If SRM was harmless, that is $\eta_G = 0$, the savings relative to the case of no SRM would be around 2% GDP or $1 trillion per

---

[3]The probability distribution described in Eq. 7, albeit quite simple, captures the main characteristics of the climate distribution described by the IPCC. According to the IPCC (2007) "climate sensitivity is likely to be in a range of 2–4.5°C with a best estimate of about 3°C," while "values substantially higher than 4.5°C cannot be excluded." Using the simplified expression for radiative forcing obtained from the IPCC Third Assessment Report, we find that with probability 0.9 climate sensitivity will be 2.6°C and with probability 0.1 climate sensitivity will be 8.5°C. On average, climate sensitivity is 3.2°C.

year, which is equivalent to a reduction in the expected costs of climate change close to 85%. If on the other hand $\eta_G = D_T$, the cost reduction due to the introduction of SRM is around 1.1%GDP or $550 billion per year, which is equivalent to a reduction in the expected costs of climate change close to 50%. To illustrate the role that the uncompensated damages from CO_2 play in the model, we set $\eta_S = 0$ (orange lines in lower panel of Fig. 2). The difference between the black and orange lines are due to costs such as ocean acidification that cannot be compensated by SRM even if there are no damages from SRM, ($\eta_G = 0$).

We find that it is still optimal to implement high levels of SRM even if the marginal damages from SRM are higher than those of climate change ($\eta_G = 1.5D_T$). This counter-intuitive result arises because SRM can be implemented after the uncertainty about climate sensitivity is resolved. *The signal advantage of SRM is its quick response: even if damages from SRM are substantially high, it is still valuable to have SRM available, as a complement to abatement measures, in case the climate sensitivity is high.* Consider the counterfactual: without SRM it is difficult to bound climate damages in the face of climate sensitivity uncertainty and inertia as argued by Roe and Baker (2007) and Weitzman (2009).

4 Uncertain SRM: assessing the value of learning about the side-effects

In this section we explicitly introduce uncertainty about the damages from SRM. We examine the effect that reducing this uncertainty has on the optimal policy and the total costs of addressing climate change. Uncertainty about the risk and the effectiveness of SRM may be reduced by researching and engaging in the small scale implementation of SRM. We describe the reduction of uncertainty–achieved by research or otherwise—as *learning*.

The implications of learning for the optimal policy depends strongly on when learning occurs in relation to decisions. We treat three scenarios (Fig. 1). The first scenario assumes no learning (NL), or equivalently that learning occurs after abatement and SRM are chosen. In the second we assume that learning occurs before SRM decisions are made, but after abatement is chosen; we refer to this as *second stage learning—2L*. In the third scenario, we assume that learning occurs before abatement and SRM decisions are made; we refer to this as *first stage learning—1L*.

To introduce risk associated with SRM, we treat the damages due to SRM, η_G, as a random variable $\widetilde{\eta_G}$ that follows the distribution:

$$\widetilde{\eta_G} = \begin{cases} \eta_G^H = D_T & \text{with probability } q = 0.5 \\ \eta_G^L = 0 & \text{with probability } 1 - q = 0.5. \end{cases} \tag{8}$$

which has an expected value of $0.5D_T$. When $q = 0.5$, we have no information regarding whether damages from SRM are larger or smaller than those of climate change. In this case, and due to the linearity of the model imposed by our assumption of quadratic damages, the optimal policy is equal to the case of no uncertainty when $\eta_G = 0.5D_T$. This is also true for other probability distributions that preserve the mean of the original distribution. The linearity of the model with respect to the choice of SRM implies that the decision maker is risk neutral. This very important characteristic allows us to concentrate on the value of learning that reduces uncertainty (Baker 2006).
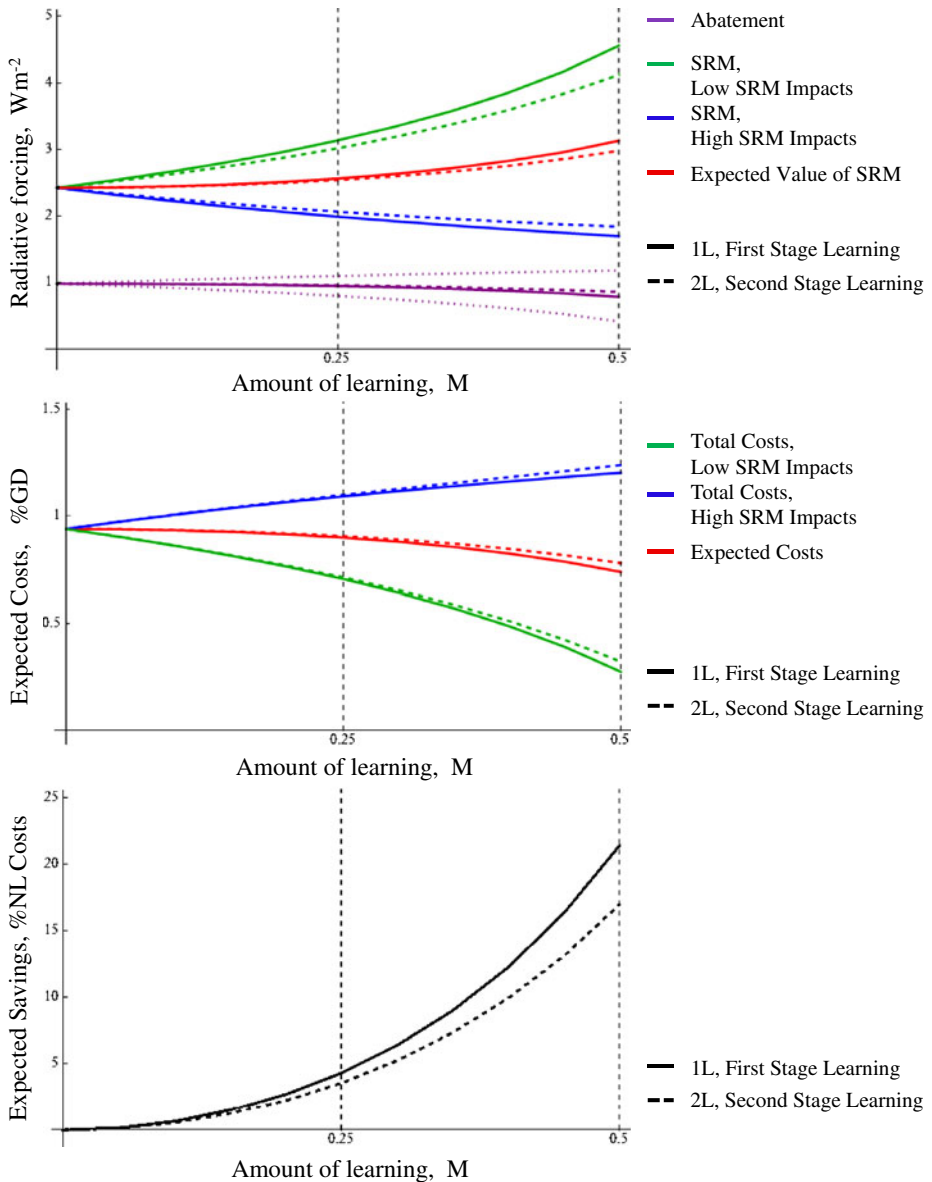
Fig. 3 The effects of learning on the optimal levels of abatement and SRM, and its implications for total costs, as a function of the amount of learning, *M*. The second stage learning (2L) scenario is denoted by *dashed lines*, the first stage learning (1L) scenario is denoted by *solid lines*, and the No Learning scenario corresponds to $M = 0$. The *top panel* shows the effects of learning on the expected level of SRM and abatement. The *blue line* shows the expected level of SRM in the case where learning reveals that the SRM impacts are worse than expected, while the *green line* shows the converse. In *red* is the expected value of SRM when the probability of learning that the damages from SRM are larger or smaller than the damages from climate change is 0.5. The *purple* lines shows the optimal level of abatement, *A*. The *purple dotted lines* show the level of abatement in the 1L scenario. The *middle panel* shows the expected costs, with the same convention as the *top panel*. The *bottom panel* shows the total savings. Total savings are the difference between the total costs of the optimal policy when there is no learning and the corresponding learning scenario

We assume that learning increases the spread of the original distribution by skewing the probability towards one of the two outcomes. Learning is equally likely to show that the damages from SRM are equal to the damages from climate change, $\eta_G = D_T$, or to show they are zero, $\eta_G = 0$. That is, learning does not change the expected value of η_G. In the case where, with probability 0.5, learning reveals that the impacts are more likely to be worse than expected, the distribution of $\widetilde{\eta_G}$ takes the form:

$$\widetilde{\eta_G} = \begin{cases} \eta_G^H = D_T & \text{with probability } q^H = 0.5 + M \\ \eta_G^L = 0 & \text{with probability } 1 - q^H = 0.5 - M. \end{cases} \quad (9)$$

where $M \in [0, 0.5]$ describes the amount of learning that occurs. On the other hand, if learning reveals that low impacts from SRM are more likely, then the distribution of $\widetilde{\eta_G}$ takes the form:

$$\widetilde{\eta_G} = \begin{cases} \eta_G^H = D_T & \text{with probability } q^L = 0.5 - M \\ \eta_G^L = 0 & \text{with probability } 1 - q^L = 0.5 + M. \end{cases} \quad (10)$$

We present our analysis as a function of M, the amount of learning that occurs. When $M = 0$ no learning has occurred. Whereas when $M = 0.5$, learning has fully eliminated uncertainty.

Figure 3 shows the effects of learning on the optimal policy (top panel), the expected costs of climate change (middle panel), and the net savings or *expected value of information* (bottom panel), as functions of the amount of learning, M. First stage learning (1L) is preferred to second stage learning (2L) for two related reasons. First, it allows better decisions in terms of SRM: SRM is lower when learning reveals high SRM damages and SRM is higher when learning reveals low SRM damages. This tendency is accentuated when learning is larger ($M \to 0.5$). Second, the value of learning is an increasing function of the amount of learning and it is higher under first stage learning (1L).

The top panel in Fig. 3 also shows that the expected level of abatement does not change significantly with early (1L) or late (2L) learning compared to the no learning (NL) scenario. This suggest that, at least for the optimal policy, learning about SRM do not affect the expected value of abatement. Of course, the realized—as opposed to expected—value of abatement does strongly depend on the outcome of learning.

5 Conclusions

We explore a simple model in which a decision maker chooses the level of emissions abatement and SRM that minimizes the costs of climate change in the face of uncertainty about the impacts of both emissions and SRM. We draw two main conclusions. First, imperfect SRM is an effective means to manage the uncertainty in the climate response because it can be implemented quickly after this uncertainty is resolved, providing a tool to manage the inertia in the carbon-climate decision problem. Without SRM, the existence of high-consequence low-probability climate impacts, combined with the irreversibility of emissions, may force very high levels of abatement and hence high costs. In our model, we find that SRM is used in the case of an unlucky (high-impact) outcome even if the damages from SRM exceed the expected damages from climate change. Under the same assumption about

the damages from SRM, SRM is substantially reduced when climate impacts are relatively low.

Second, we find that learning about SRM—that is the value of information associated with reducing the uncertainty about the side-effects of SRM—can reduce the overall costs of climate change in the order of 10%, depending on the amount of learning. Suppose learning about SRM reduced the expected cost of climate change by 5%. We can compare these savings, equivalent to 0.05% of world GDP, with the current spending on SRM research which is less than $10 million per year, or 0.00002% of GDP; though we cannot, of course, conclude that learning will be proportional to spending since we don't know how effective this research will be in reducing uncertainty about SRM. Moreover, this specific numerical result depends on the calibration of the model and on the assumptions about the prior probability distribution over the side-effects of SRM.

The model is a highly simplified representation of the problem and its applicability is limited by the caveats presented in the introduction to this paper. Also, the model is afflicted by the same constraints attached to any model of climate policy that supposes a single decision maker; namely, no strategic interaction, no asymmetries and therefore, no distributional issues. We have used, however, a calibration of climate damages and abatement that is widely used and is representative of results derived in many complex models. Hence, the limitations of the model likely do not affect its main result; that is, SRM is valuable for managing climate risk, not because of its low cost, but because it can be implemented quickly if we discover that climate impacts are high, a "climate emergency."

References

Allen MR, Ingram WJ (2002) Constraints on future changes in climate and the hydrologic cycle. Nature 419:224–232

Bala G, Duffy PB, Taylor KE (2008) Impacts of geoengineering schemes on the global hydrological cycle. Proc Natl Acad Sci 105:7664–7669

Baker E (2006) Increasing risk and increasing informativeness. Oper Res 54(1):26–36

Brander LM, Rehdanz K, Tol RSJ, van Beukering PJH (2009) The economic impact of ocean acidification on coral reefs. ESRI working paper 282

Caldeira K, Matthews HD (2007) Transient climate-carbon simulations of planetary geoengineering. Proc Natl Acad Sci 104:9949–9954

Caldeira K, Wickett ME (2003) Anthropogenic carbon and ocean pH: the coming centuries may see more ocean acidification than the past 300 million years. Nature 425:365

Caldeira K, Wickett ME (2005) Ocean model predictions of chemistry changes from carbon dioxide emissions to the atmosphere and ocean. J Geophys Res Oceans 110:C9

Core Writing Team, Pachauri RK, Reisinger A (eds) (2007) IPCC 2007: climate change 2007: synthesis report. Tech report, Intergovernmental Panel on Climate Change

Crutzen PJ (2006) Albedo enhancement by stratospheric sulfur injections: a contribution to resolve a policy dilemma? Clim Change 77:211–219

Environment Canada (1997) Global costs and benefits of the Montreal protocol

Goulder LH, Mathai K (2000) Optimal CO_2 abatement in the presence of induced technological change. J Environ Econ Manage 39:1–38

Keith DW (2000) Geoengineering the climate: history and prospect. Annu Rev Energy Environ 25:245–284

Keith DW, Dowlatabadi H (1992) A serious look at geoengineering. Eos Trans Am Geophys Union 73:289, 292–293

Kravitz B, Robock A, Oman L, Stenchikov G, Marquardt AB (2009) Sulfuric acid deposition from stratospheric geoengineering with sulfate aerosols. J Geophys Res 114:D14109

Matthews HD, Gillett NP, Scott PA, Zickfeld K (2009) The proportionality of global warming to cumulative carbon emissions. Nature 459:829–833

McClellan J, Sisco J, Suarez B, Keogh G (2010) Geoengineering cost analysis. Final report, Aurora Flight Sciences Corporation, Cambridge, Massachusetts

Millar-Ball A (2012) The Tuvalu syndrome: can geoengineering solve climate's collective action problem? Clim Change 110:1047–1066

Moreno-Cruz J (2011) Mitigation and the geoengineering threat. Mimeo, Unviersity of Calgary

Moreno-Cruz J, Ricke K, Keith D (2012) A simple model to account for regional inequalities in the effectiveness of solar radiation management. Clim Change 110:649–668

Morgan G, Keith DW (1995) Subjective judgments by climate experts. Environ Sci Technol 29(10)

Nordhaus W (2008) A question of balance: weighing the options on global warming policies. Yale University Press, 234 pp

Ricke K, Morgan G, Allen M (2010) Regional climate response to solar radiation management. Nature Geosci 3:537–541

Robock A, Oman L, Stenchikov G (2008) Regional climate responses to geoengineering with tropical and Arctic SO2 injections. J Geophys Res 113:D16101

Robock A, Marquardt A, Kravitz B, Stenchikov G (2009) Benefits, risks, and costs of stratospheric geoengineering. Geophys Res Lett 36:L19703

Roe GH, Baker MB (2007) Why is climate sensitivity so unpredictable? Science 318:629–632

Shepherd J, Caldeira K, Haigh J, Keith D, Launder B, Mace G, MacKerron G, Pyle J, Rayner S, Redgwell C, Watson A (2009) Geoengineering the climate: science, governance and uncertainty. The Royal Academy

Soden BJ, Wetherald RT, Stenchikov GL, Robock A (2002) Global cooling after the eruption of Mount Pinatubo: a test of climate feedback by water vapor. Science 296:727–730

Solomon S, Portman RW, Garcia RR, Thomason LW, Poole LR, McCormick MP (1996) The role of aerosol variations in anthropogenic ozone depletion at northern midlatitudes. J Geophys Res 101:6713–6727

Solomon S (1999) Stratospheric ozone depletion: a review of concepts and history. Rev Geophys 37:275–316

Solomon S, Plattnera G-K, Knutti R, Friedlingstein P (2009) Irreversible climate change due to carbon dioxide emissions. Proc Natl Acad Sci 106:1704–1709

Sunstein CR (2007) Of Montreal and Kyoto: a tale of two protocols. Harv Environ Law Rev 31:1–65

Weitzman M (2009) On modeling and interpreting the economics of catastrophic climate change. Rev Econ Stat 91:1–19

Wigley TML (2006) A combined mitigation/geoengineering approach to climate stabilization. Science 314:452–454

Zickfeld K, Morgan MG, Frame DJ, Keith DW (2010) Expert judgments about transient climate response to alternative future trajectories of radiative forcing. PNAS 107:12451–12456