

University of Heidelberg

Department of Economics



Discussion Paper Series | No. 540

**The Intergenerational Transfer of Solar Radiation
Management Capabilities and Atmospheric Carbon Stocks**

**Timo Goeschl, Daniel Heyen,
and Juan Moreno-Cruz**

January 2013

The Intergenerational Transfer of Solar Radiation Management Capabilities and Atmospheric Carbon Stocks

Timo Goeschl*
Department of Economics
Heidelberg University

Daniel Heyen
Department of Economics
Heidelberg University

Juan Moreno-Cruz
Department of Economics
Georgia Institute of Technology

October 23, 2012

Abstract

Solar radiation management (SRM) technologies are considered one of the likeliest forms of geoengineering. If developed, a future generation could deploy them to limit the damages caused by the atmospheric carbon stock inherited from the current generation, despite their negative side effects. Should the current generation develop these geoengineering capabilities for a future generation? And how would a decision to develop SRM impact on the current generation's abatement efforts? Natural scientists, ethicists, and other scholars argue that future generations could be more sanguine about the side effects of SRM deployment than the current generation. In this paper, we add economic rigor to this important debate on the intergenerational transfer of technological capabilities and pollution stocks. We identify three conjectures that constitute potentially rational courses of action for current society, including a ban on the development of SRM. However, the same premises that underpin these conjectures also allow for a novel possibility: If the development of SRM capabilities is sufficiently cheap, the current generation may for reasons of intergenerational strategy decide not just to develop SRM technologies, but also to abate more than in the absence of SRM.

Keywords: Geoengineering; Climate Change; Intergenerational Issues; Strategic Behavior.

JEL Codes: D9; O33; Q54; Q55.

*Corresponding author: Bergheimer Str. 20 69115 Heidelberg, Germany. E-Mail goeschl@eco.uni-heidelberg.de

1 Introduction

The economics of climate change have been emphasizing for a long time that concerns for intergenerational equity imply that the current generation needs to reduce its greenhouse gas (GHG) emissions. Such a reduction will benefit future generations by decreasing the damages they are expected to suffer as a result of climate change (e.g. Stern 2006). Consequently, a central concern of the literature on optimal climate policy is to determine the optimal scale of abatement efforts vis-à-vis the business-as-usual scenario and thus to determine the optimal intertemporal trajectory of mitigation (Stern 2006; Nordhaus 2007; Tol 2001).

Recent developments have started to challenge the almost exclusive role of emissions reductions in climate policy. Apart from progress on the economics of adaptation to climate change (Agrawala and Fankhauser 2008), there is growing awareness among economists of the increasing plausibility of so-called “climate engineering” (CE). This term is a shorthand for deliberate large-scale interventions in the Earth’s climate system with the aim of limiting the damages of excessive atmospheric carbon stocks (Keith 2000). A variety of approaches run under the heading of “climate engineering” or “geoengineering”. Most of the attention focuses—due to feasibility, effectiveness, and cost—on solar radiation management (SRM) technologies. These technologies enable an increase in the Earth’s albedo with respect to solar radiation, e.g. by the dispersal of reflective aerosols in the stratosphere. This allows the Earth to tolerate higher atmospheric carbon stocks while keeping global mean surface temperatures within acceptable bounds. Merely a theoretical possibility decades ago (Fleming 2010), it is now considered feasible that an ambitious R&D program could deliver effective SRM technologies within a few decades (Ridgwell et al. 2012).

A setting in which GHG emissions reductions are not the only option for addressing climate change damages, but compete with or are complemented by investment in geoengineering capabilities raises an entirely new set of questions. Some early assessments of climate engineering conclude that a future generation with access to such technologies would be able to handle the damages associated with stock pollutants at a surprisingly low direct variable cost (Klepper and Rickels 2012; Barrett 2008). However, these technologies are not understood to be “magic bullets”. The current consensus is that SRM interventions will involve significant indirect costs through side effects, such as changes in precipitation patterns, and that these side effects will themselves raise issues of how benefits and costs are distributed (Klepper and Rickels 2012; Moreno-Cruz and Keith 2012; Ricke et. al. 2010). As a result, a decision-maker will have to carefully weigh the social benefits of avoided carbon damages and the social costs of SRM-induced side effects when making a decision on the deployment of geoengineering measures. Some papers have started to establish the conditions under which the deployment of SRM may or may not be meaningful (Bickel and Agrawal 2012; Goes et al. 2011; Moreno-Cruz and Keith 2012), to consider the impacts of a geoengineering capability on international negotiations (Barrett 2008; Moreno-Cruz 2010) and how to regulate deployment (Barrett 2008; Victor 2008).

The starting point of this paper is the observation that at the present time, the technological capabilities for a deployable SRM system do not exist. These capabilities would

have to be created at a cost, and the costs of developing these capabilities are considered substantial (Klepper and Rickels 2012). It would fall to the current generation to devote some of its resources to investment into a R&D program that would make these capabilities available, perhaps 30 years from now, to a future generation (Schelling 1996). But is it a rational course of action for the current generation to develop these capabilities for a future generation that can then decide on its use? And how would a decision to develop such SRM capabilities impact on abatement efforts?

Despite agreement on the basic premises of the decision problem, the current literature offers at least five different conjectures regarding the rational course of action. One conjecture, termed “arming the future” argues in favor of SRM R&D since the technology would act as some type of insurance in the event that the sensitivity of the climate system with respect to the carbon stock turns out to be high (Moreno-Cruz and Keith 2012). This rationale of SRM R&D has been advanced by a number of commentators (Gardiner 2010; Schneider 1996 cited in Gardiner 2011, p.9) and enjoys considerable support. Some have even gone so far as to claim that, since SRM is an imperfect substitute for emissions abatement, a positive SRM R&D decision will not affect GHG emissions abatement (Bunzl 2009). This conjecture, which we dub “abatement invariance”, contrasts with a third conjecture, namely that investment in SRM R&D detracts from mitigation and will lead to less abatement effort (Keith 2010, Shepherd 2009). This possibility of a reduction in mitigation effort by the current generation has been characterized as a manifestation of “moral hazard” that a rational course of action would avoid (Shepherd 2009, Hale 2012).

In other conjectures, the concerns lie with the behavior of the future generation. Future decision-makers may engage in climate engineering under circumstances and on a scale that the current generation would not hold to be optimal. At least three mechanisms are thought to explain why future generations may deviate from the judicious course of action. One is behavioral: Given the impression of an impending or immediate “climate emergency”, there is a belief that decision-makers will succumb to emotional factors that favor SRM deployment as a ‘quick fix’ (e.g. Bodansky 1996).¹ Similarly, while sunk cost of its development should on rational grounds be immaterial to the decision on whether to deploy SRM technology, there is ample historical experience that use of a capability is often regarded as necessary to justify large sunk cost (Gardiner 2011). The second mechanism lies in the political economy of science and technology: Researchers and industry funded to carry out research on certain technologies become interest groups in favor of technology deployment, creating a vocal and influential lobby for SRM use (Jamieson 1996).² The final mechanism is the potential for a genuine change of preferences regarding SRM once its availability has become a fact of life for a future generation.³ Irrespective of the mechanism, a bias in favor of SRM interventions,

¹The American Meteorological Society, for instance, believes that “[...]geoengineering technologies, once developed, may enable short-sighted and unwise deployment decisions, with potentially serious unforeseen consequences.” (2009)

²This concern is also present among the general public. Mercer et al (2011) report very strong agreement on the statement that “Research will lead to a technology that will be used no matter what the public thinks.” among participants in a large-scale international survey of public perceptions of climate engineering.

³Already in one of the early reports by the US National Academy of Sciences on climate change, the authors

once those technological capabilities have been created, is the argument underpinning two additional conjectures advanced in the literature. One is that the rational course of action for today’s generation is to rule out research on geoengineering measures in order to prevent the future generation from acting against its own best interests (Keith et al. 2010).⁴ The fifth conjecture is a different interpretation of the “moral hazard” argument by Bunzl (2009) and turns the logic on its head. It postulates that in the face of a pro-CE bias among the future generation, the rational course of action is to offer SRM as a backstop for a “climate emergency” and to abate less. The reason is that in a world in which the future generation uses SRM indiscriminately, the carbon stock damages imposed on the future generation are consistently lower than in the absence of SRM capabilities. As a result, the current generation would be treated more equitably if it allowed itself to emit more (Hale 2012).

Against this background of competing conjectures, this paper makes two contributions. The first is to extend the economic literature on the intergenerational aspects of climate policy in the direction of technology transfers. To our knowledge, the transfer of a mitigating technology, such as SRM, to a future generation that bears the damage costs of a stock pollutant has not been explicitly considered before. The second contribution is to demonstrate how simple economic analysis can provide a useful starting point for discriminating between the different conjectures on the ‘right’ course of action regarding a portfolio of both mitigation and SRM R&D activities. We provide these contributions by developing the simplest two-generations model that captures the four joint premises that appear to underpin all conjectures. These common elements are (i) that the current generation cares about the future generation sufficiently to be concerned about the stock damages of atmospheric carbon, (ii) that there may be a pro-SRM bias in the future, (iii) that both abatement and R&D involve a cost today and (iv) that both determine, in an environment characterized by uncertainty about the damages associated with atmospheric carbon, the future generation’s carbon stock and technological capabilities. Using this model, we study the subgame perfect behavior of the current generation in order to determine which of the conjectures above can arise, and if so, under which conditions.

Our results demonstrate that the possibility of transferring mitigating technologies into the future both re-emphasizes important economic insights on climate policy and raises important new issues. For example, even in the absence of a pro-SRM bias, the presence of an SRM option offsets current abatement. Far from constituting an instance of “moral hazard” (Bunzl 2009), this is simply a result of the partial substitutability between abatement and SRM that a current generation will rationally want to exploit. At the same time, the presence of a pro-SRM bias will constitute an important source of potentially powerful strategic distortions between generations that support some of the conjectures, but not all. A failure to

raise the possibility that interest in CO₂ may generate or reinforce a lasting interest in national or international means of climate and weather modification; once generated, that interest may flourish independent of whatever is done about CO₂ (p. 470).

⁴Ruling out CE research could take the shape of the recent explicit ban on climate engineering R&D that was declared at the Conference of the Parties of the U.N. Convention on Biological Diversity in Nagoya in 2010.

carefully define those comparisons that are meaningful may be to blame for some of the confusion. As the analysis makes clear, rising R&D costs lead *ceteris paribus* always to ruling out SRM research for rational reasons. The same holds for an increasing pro-SRM bias. However, providing no SRM R&D will always be combined with higher abatement levels than if SRM was made available. At the same time, we find no support for the conjecture that abatement weakens due to a distortion between generations if R&D is undertaken. Quite the opposite: If SRM is made available, abatement will never fall below the non-distortion benchmark but increase the higher the distortion between generations. That abatement will increase even if SRM is provided is a new finding that has not been discussed before. The intuition is that an altruistic current generation can and will want to partially offset a pro-SRM bias among the future generation by providing more abatement today, thus reducing the incentives to deploy SRM in the future.

We proceed as follows. In the following section, we develop the simple two-generations model that captures the salient components of the SRM R&D debate in the most parsimonious and tractable fashion. In section 3, we provide a major step in the debate by defining the meaningful comparison point by the way of a benchmark. Section 4 derives the equilibria of the intergenerational game and establishes the main propositions. In section 5, we conclude.

2 The Model

2.1 The setting

Here, we develop a setting that pares an intergenerational decision problem involving both a stock pollutant and the development of an imperfect backstop technology down to its very essentials. Figure 1 provides a graphical representation of the setting, which features four periods. While the building blocks on stock pollution, abatement, and damages are taken from the mainstream literature, the novel aspects are the inclusion of the climate engineering option and the associated intergenerational issues so as to capture the common narrative elements described in the introduction.

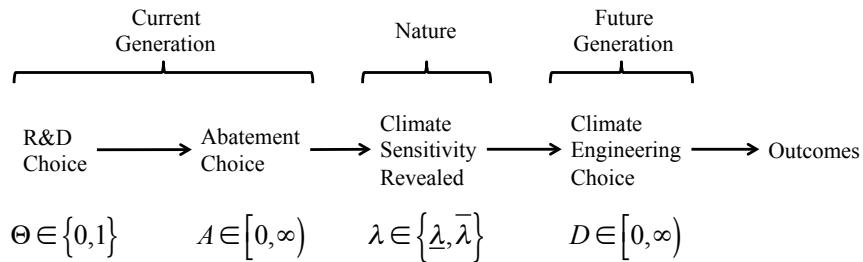


Figure 1: Timing of the game.

The basic set-up consists of two non-overlapping generations, termed “current” and “future”. The current generation decides in period 1 on whether to invest in R&D for future SRM capabilities ($\Theta = 1$) or not ($\Theta = 0$) and chooses pollution abatement level A in period

2.⁵ R&D of SRM involves technological as well as regulatory and institutional costs that will allow the future generation to use the climate engineering technology. These costs could be significant (Klepper and Rickels 2012). In period 3, nature resolves the key scientific uncertainty about climate change (Roe and Baker 2007), namely the climate sensitivity parameter λ that in turn determines the marginal damages associated with the future global carbon stock. In period 4, the future generation decides whether to deploy the SRM technology in order to counteract the damages T from the unabated pollution stock. If it deploys SRM ($D > 0$), the future generation reduces damages from the carbon stock, but suffers environmental damages G associated with the use of SRM, e.g. in the form of disruptions of the hydrological cycle. If no SRM is used ($D = 0$), future society faces the full temperature damages T .

As other tractable economic models of climate change, we specify the damages associated with the global carbon stock as caused by increased temperatures above historical averages. We employ the temperature damage function T by Moreno-Cruz and Keith (2012) of the form

$$T = \lambda^2 (R_0 - A)^2 . \quad (1)$$

The damage function consists of two arguments: The first argument is the squared carbon sensitivity of the climate to a doubling of CO₂, λ . The second argument is the square of the net deviation in the carbon stock from historical levels ($R_0 - A$), which consists of the business-as-usual increase in the carbon stock R_0 minus the abatement effort A .⁶ From the vantage point of the current generation, i.e. in periods 1 and 2, the carbon sensitivity is a random variable with two possible realizations: a carbon-sensitive value of $\bar{\lambda} > 0$ with probability p and a carbon-insensitive value $\underline{\lambda}$ ($0 < \underline{\lambda} < \bar{\lambda}$) with probability $(1 - p)$. As expression (1) makes clear, abatement is productive as it reduces the expected value of pollution damages associated with climate change. Along with all moves $\{\Theta, A\}$, the pollution damage function (1) is common knowledge.

Abatement costs are usually assumed to be convex in abatement efforts. As a simple approximation, we model the abatement cost function of the quadratic type

$$X = \alpha A^2 , \quad (2)$$

with increasing marginal abatement cost $2\alpha A$. The R&D process is modeled in a deterministic fashion, following Goeschl and Perino (2007): R&D for a functioning SRM technology requires payment of a fixed amount K in period 2 and delivers the climate engineering technology by period 4. The cost of not providing the SRM technology ($\Theta = 0$) is zero.

The second option to counteract temperature damages, which can be combined with abatement, is the deployment of SRM. Like the abatement level A , we measure the amount D of deployed SRM in terms of compensated carbon stock (cf. Moreno-Cruz and Keith 2012) such that temperature damages read $T = \lambda^2 (R_0 - A - D)^2$. In that sense, abatement and

⁵The sequentiality of the decisions on Θ and A is purely for ease of presentation. A simultaneous choice of Θ and A leads to identical results.

⁶Temperature damages are assumed to be quadratic in temperature increases. The latter are assumed to be linear in the pollutant stock change

SRM are equivalent regarding temperature damage compensation. They differ, however, in terms of timing and costs.

The use of SRM involves negligible direct costs (Barrett 2008), but causes collateral damages because of unintended negative impacts. These negative impacts comprise changes in the hydrological cycle and increase in air pollution, according to current assessments (Ricke et al. 2010; Kravitz et al. 2009). The general form of the damages G from SRM is a higher-order polynomial involving the volume of aerosol injected into the stratosphere (D), the net carbon stock increase ($R_0 - A$), and various particle characteristics (Bala et al. 2010; Caldeira and Wood 2008; Ricke et al. 2010). Recent simulations with general circulation models have shown, however, that the changes in temperature and precipitation are disproportionately driven by the linear term attached to the volume of SRM (Ban-Weiss and Caldeira 2010; Moreno-Cruz et al. 2012). A first approximation of the current generation's assessment of SRM damages would therefore specify an essentially linear relationship between damages and the amount D of SRM of the form

$$G = \rho D \quad (3)$$

with ρ denoting the constant marginal value of SRM damages. An immediately obvious economic viability condition for SRM is that $\rho < 2\alpha R_0$. Otherwise, the marginal cost of even the last unit of abatement is lower than the marginal damage of SRM, implying that SRM is never a competitive substitute for abatement.

The final component to make the setting reflect the current literature is a device that captures the possible presence of a bias of the future generation in favor of SRM deployment. The bias implies that unit for unit, the future generation values SRM damages generally less than the current generation for reasons of behavioral mechanisms, political economy, or genuine preference shifts. As a shorthand for this divergence between current and future generations, we introduce a bias parameter $\beta \in [0, 1]$ such that from the future generation's vantage point, damages are $\tilde{G} = (1 - \beta)\rho D$. If there is no bias ($\beta = 0$), then the current and the future generation have an identical relative valuation between temperature damages and SRM damages.

2.2 Objectives and equilibrium concept

Against this background, we now define the objective functions, payoffs, and strategies of both generations. The future generation's objective is to minimize the sum of pollution damages and SRM damages, taking the choices of the current generation on abatement and R&D and the realization of climate sensitivity λ as given. Given R&D on SRM, the future generation can determine its optimal volume of SRM, D^* , that is how much of the inherited carbon stock should be compensated for by means of SRM. The optimal volume of SRM is a function of the abatement level A , the realized climate sensitivity λ and the bias parameter β . The future generation's problem is therefore

$$\min_{D \in [0, R_0 - A]} \left\{ \lambda^2 (R_0 - A - D)^2 + (1 - \beta)\rho D \right\}. \quad (4)$$

It is both intuitively obvious and clear from the first-order conditions that use of SRM requires that $(1 - \beta)\rho < 2\lambda^2(R_0 - A)$, i.e. the marginal damage of the first unit of SRM must be below the temperature damages thus avoided. Otherwise, the corner solution $D^* = 0$ is optimal. For $D^* > 0$, the optimal amount is given by $D^*(A, \lambda, \beta) = R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2}$.

From the vantage point of the current generation, the combination of the two possible realizations of climate sensitivity λ and the conditions on corner and interior solutions give rise to three possible deployment profiles. One profile, D_{uncond} , implies use of SRM by the future generation irrespective of whether the climate turns out to be sensitive or insensitive to carbon. Another, D_{never} , features no use of SRM, even if the climate turns out to be sensitive to carbon. The third, D_{cond} , conditions the use of SRM on the climate sensitivity, using SRM for a sensitive outcome and not using SRM for an insensitive one. The profiles are defined by

$$D_{\text{uncond}} = \begin{cases} R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \bar{\lambda} \\ R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \underline{\lambda} \end{cases}, \quad D_{\text{cond}} = \begin{cases} R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \bar{\lambda} \\ 0 & \text{if } \lambda = \underline{\lambda} \end{cases}$$

$$D_{\text{never}} = \begin{cases} 0 & \text{if } \lambda = \bar{\lambda} \\ 0 & \text{if } \lambda = \underline{\lambda} \end{cases}. \quad (5)$$

Due to a decrease of perceived damages caused by SRM, an increase in the bias β increases the SRM amount deployed. For $\beta = 1$, which refers to future generation attributing no damages to the use of SRM, temperature damages will be fully compensated, $T = 0$.

A quick inspection of the deployment profiles shows that the current generation, given its belief about the future generation's β , determines which profile will be chosen in the future through its choice of abatement A . More specifically, there is a lower and an upper threshold on abatement, with $\underline{A}_{\text{crit}} = R_0 - \frac{(1-\beta)\rho}{2\lambda^2}$ and $\bar{A}_{\text{crit}} = R_0 - \frac{(1-\beta)\rho}{2\lambda^2}$ respectively, such that

$$D^* = \begin{cases} D_{\text{uncond}} & \text{if } A \leq \underline{A}_{\text{crit}} \leq \bar{A}_{\text{crit}} \\ D_{\text{cond}} & \text{if } \underline{A}_{\text{crit}} \leq A \leq \bar{A}_{\text{crit}} \\ D_{\text{never}} & \text{if } \underline{A}_{\text{crit}} \leq \bar{A}_{\text{crit}} \leq A \end{cases}. \quad (6)$$

Through sufficiently high abatement, therefore, the current generation can ensure that SRM will never be used. Conversely, by abating little, available SRM would always be used at some positive level, even if the climate turns out to be relatively insensitive to carbon stocks. Finally, abatement levels between the two critical thresholds give rise to SRM deployment that is conditional on the realization of carbon sensitivity of the climate. Here, SRM would only be used in the eventuality of a sensitive climate.⁷

The optimal choice of costly SRM R&D Θ in period 1 and abatement A in period 2 constitute the essence of the current generation's problem. In its choice, the current generation takes into account both the costs to itself (in the form of R&D and abatement costs) and the costs borne by the future generation (in the form of damages). The cost minimization

⁷Note that the deployment profiles D_{uncond} and D_{cond} become indistinguishable at the abatement level $\underline{A}_{\text{crit}}$. The same holds for D_{cond} and D_{never} at \bar{A}_{crit} .

objective requires that current and future costs need to be made comparable through an appropriate discount factor δ . The objective function is then given by

$$\min_{\Theta, A} C(\Theta, A | \beta) = \min_{\Theta, A} \left\{ \Theta K + \alpha A^2 + \delta(1-p) (\underline{\lambda}^2 (R_0 - A - D(A, \underline{\lambda}, \beta))^2 + \rho D(A, \underline{\lambda}, \beta)) + \delta p (\bar{\lambda}^2 (R_0 - A - D(A, \bar{\lambda}, \beta))^2 + \rho D(A, \bar{\lambda}, \beta)) \right\}. \quad (7)$$

For ease of exposition, we set $\delta = 1$ (no discounting). Also note that the objective function of the current generation makes explicit reference to β since the bias determines profile and amount of SRM deployment. We thus analyze functions of the form $C(\Theta, A | \beta)$ or, equivalently, the form $C(\Theta, A | D)$ where D denotes a certain deployment profile.

The optimal abatement levels differ depending on whether there is R&D on SRM or not. If no SRM R&D is carried out, $\Theta = 0$, the future generation cannot deploy the technology. This simplifies the objective function since $D = 0$ for all A , λ and β . The associated optimal abatement level is

$$A_{\text{NoR\&D}} = \frac{p\bar{\lambda}^2 + (1-p)\underline{\lambda}^2}{\alpha + p\bar{\lambda}^2 + (1-p)\underline{\lambda}^2} R_0 > 0, \quad (8)$$

which is a fraction of the business-as-usual increase in the carbon stock R_0 . In line with intuition, abatement in the absence of SRM increases with a higher probability of a carbon-sensitive climate ($\frac{d}{dp} A_{\text{NoR\&D}} > 0$) and for higher levels of sensitivity ($\frac{d}{d\lambda} A_{\text{NoR\&D}} > 0$, $\frac{d}{d\bar{\lambda}} A_{\text{NoR\&D}} > 0$). The higher marginal abatement costs (higher α), the lower the abatement level: While costless abatement ($\alpha = 0$) leads to full compensation, $A_{\text{NoR\&D}} = R_0$, abatement goes down to zero if marginal abatement costs tend to infinity.

If SRM R&D is carried out, $\Theta = 1$, the choice of the optimal A is more intricate because a change in A influences the future generation's discrete decision which of the three deployment profiles to choose. The objective function is therefore only piecewise differentiable and the optimal abatement level in each segment is not necessarily an interior solution. At the same time, note that the bias parameter β influences the choice of A only through its impact on the deployment profile; within one profile β has no impact on the abatement choice as $\frac{d}{d\beta} \left(\frac{d}{dA} D(A, \lambda, \beta) \right) = 0$. This is an interesting property that results from the dominance of the linear component in SRM damages and that will be exploited in section 4.

From an abstract point of view, the model set-up and objectives that capture the narratives on whether the current generation should develop SRM capabilities for the future generation define a sequential game with incomplete information. Its basic structure, in particular the technology transfer decision, is a variant of the trust game by Kreps (1990). However, the intergenerational decision problem here features two important differences: One is the availability of the second instrument in the form of abatement, the other the presence of exogenous uncertainty in the form of the random variable λ . The proper solution concept for determining the equilibrium played by the current and future generation is that of sub-game perfection (SP). The current generation, looking forward, employs backward induction to solve problem (7): By determining the optimal play of future generation ($D^* | A, \lambda, \beta$) in period 4 contingent on current generation's choices in periods 1 and 2 and nature's move in

period 3, it identifies its own optimal play $\{\Theta^*, A^*\}$ in periods 1 and 2. However, it has to do so not knowing λ (see Figure 1). The equilibrium concept of SP will admit some combinations of abatement and R&D choices but not others, and thus impose elementary consistency checks on current society’s rational course of action.

3 The Benchmark

Using the framework presented above, the purpose of the following two sections is to explore the consistency of different conjectures regarding the current generation’s rational course of action.

We proceed by constructing and establishing, as the first building block, a suitable benchmark case. Such a case is one in which (i) the bias of the future generation with respect to damages from SRM is set to zero by assumption and (ii) where equilibrium play features the current generation engaging in SRM R&D and the future generation adopting a conditional deployment profile. This choice of a benchmark has three benefits: The first is that in the comparisons in section 4 with situations involving a bias, the strategic distortions introduced by the bias will be clearly identifiable. The second benefit is that this benchmark case demonstrates the consistency of the “arming the future” conjecture for a given parameter set and thus establishes one of the five conjectures as a candidate for the rational course of action. Thirdly, the benchmark case also demonstrates that even in the absence of the bias, the rational course of action involves certain subtleties that the discussion on developing SRM technologies has so far failed to identify.

In the spirit of keeping notation clutter to a minimum, we assign values to some parameters. We restrict the level of abatement to the unit interval by setting $R_0 = 1$. Furthermore, marginal temperature and SRM damages will be regarded in terms of abatement costs by assuming $\alpha = 1$ such that abatement costs are simply A^2 . Finally, we set $\underline{\lambda} = 1$ such that $\bar{\lambda} > 1$ captures temperature damages for high carbon sensitivity, relative to a fixed baseline. These restrictions preserve the relevant degrees of freedom of the model and simplify the analysis substantially.

The benchmark case assumes the absence of a bias of the future generation with respect to SRM damages, $\beta = 0$.⁸ To fix notation, equilibrium play in the benchmark case is characterized by current generations SRM R&D decision $\hat{\Theta}$ and abatement level \hat{A} as well as the future generation’s deployment profile \hat{D} . For the conjecture of “arming the future” to survive the consistency check then requires that equilibrium play can give rise to a decision to conduct R&D on SRM, $\hat{\Theta} = 1$, and an abatement level \hat{A} that induces the future generation to choose the conditional deployment profile of the SRM technology: $\hat{D}(\hat{A}, \beta = 0) = D_{\text{cond}}$. Under these conditions, the current generation will make SRM capabilities available to the future generation alongside an optimal abatement effort to be determined, and the future generation will use SRM technologies only in case that the climate system turns out to be

⁸In terms of interpretation, $\beta = 0$ could mean that the future generation has the same preferences as current generation or that current generation is just ignorant of an existing asymmetry.

carbon sensitive. The future generation, in other words, will be “armed” against inadvertent outcomes in the climate system.

We begin the construction of the benchmark case with the convenient assumption that the abatement level associated with equilibrium play in the “arming the future” conjecture ought to be an interior solution. The interior solution to the current generation’s cost-minimization problem (7) assuming a conditional SRM deployment profile is given by

$$A_{\text{cond}} = \frac{2(1-p) + p\rho}{2(2-p)} . \quad (9)$$

What parameter restrictions follow from designating A_{cond} as our solution \hat{A} ? These parameter restrictions can be summarized as restrictions on the marginal damage from SRM, ρ , and on the R&D cost of SRM, K . Note first that designating A_{cond} as \hat{A} requires that the abatement level A_{cond} chosen by the current generation actually gives rise to a conditional deployment profile by the future generation. Formally,

$$\underline{A}_{\text{crit}}(\beta = 0) < A_{\text{cond}} < \bar{A}_{\text{crit}}(\beta = 0) . \quad (10)$$

The level of abatement therefore has to lie in the middle segment defined in equation (6). Designating A_{cond} as \hat{A} therefore translates into a condition on the parameter ρ : The marginal damages of SRM ρ have to fulfill that $1 < \rho < \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)}$. Together with the economic viability condition for SRM $\rho < 2\alpha R_0 = 2$ (cf. section 2.1), this requires that the marginal damages of SRM fulfill

$$\rho \in \left(1, \min \left\{ 2, \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)} \right\} \right) \quad (11)$$

in order for A_{cond} to be feasible as \hat{A} . The intuition is straightforward: If $\rho < 1$, SRM damages are too small to give rise to conditional deployment of SRM and an unconditional deployment profile would be optimal instead. On the other hand, if $\rho > \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)}$ damages caused by SRM are too large and the optimal deployment profile would be to never use the technology. In this last case the amount of abatement would be larger than \bar{A}_{crit} .⁹

The restrictions on ρ ensure that the deployment profile chosen will indeed be conditional on the realization of λ . Since these restriction cannot by themselves guarantee that there is not some other abatement level that involves lower cost than A_{cond} , restrictions on K are also required. To do so, we compare the total costs of the conditional profile, and associated abatement level A_{cond} , with the costs of the profile without SRM deployment and the unconditional deployment profile (see section 2.2), with their respective optimal abatement levels. First, the total costs of the conditional profile have to be lower than those of the unconditional deployment profile. That is,

$$C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}}) < C(\Theta = 1, \underline{A}_{\text{crit}} | D_{\text{uncond}}) , \quad (12)$$

⁹It should not come as a surprise that the upper bound increases with $\bar{\lambda}$; that is, if the climate system is more climate sensitive, SRM damages can be larger and still allow for a conditional use profile. Less obvious is the fact that the upper bound decreases with p . An increase in p makes the upper limit more stringent because A_{cond} increases with p , but \bar{A}_{crit} does not (whether the future generation deploys SRM does not depend on ex-ante probabilities). As p increases, the current generation has an incentive to increase A_{cond} and minimize now higher expected future costs which in turn reduces the need for SRM deployment.

where $C(\Theta = 1, \underline{A}_{\text{crit}} | D_{\text{uncond}})$ are the minimum costs under the indiscriminate use of SRM.¹⁰ This is equivalent with $\frac{(\rho-1)^2}{2-p} > 0$, which is always satisfied. Restrictions on K therefore need to come from the second comparison between the total costs of the conditional profile and those of the no SRM deployment profile, taking into account that conditional deployment requires investing in SRM R&D. For conditional deployment to involve lower cost requires that

$$C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}}) < C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}}) . \quad (13)$$

This requirement translates into a limit on how costly SRM R&D can be:

$$K < \bar{K} := \frac{p(\rho(2 + p(\bar{\lambda}^2 - 1)) - 2\bar{\lambda}^2)^2}{4(2 - p)\bar{\lambda}^2(2 + p(\bar{\lambda}^2 - 1))} . \quad (14)$$

Taking the above restrictions on ρ and K together, the parameter requirements for the benchmark case are thus fully characterized. More importantly still, there is no evidence of logical contradictions that would rule out equilibrium play according to the “arming the future” conjecture if there is no bias among the future generation: There is nothing inherently contradictory about the current generation providing both abatement and SRM R&D to the future generation such that this technology can be used as a backstop in the event of a carbon-sensitive climate.

Before exploring the impact of a bias regarding SRM use among the future generation, it is useful to point out two features of the benchmark equilibrium. One is the abatement level: The abatement level $\hat{A} = A_{\text{cond}}$ is smaller than the abatement level which would be optimal without SRM, $A_{\text{cond}} < A_{\text{NoR\&D}}$. From an economic perspective, this is not surprising: Since abatement is costly, the ability to counteract the adverse effects of a pollutant implies that a higher pollution stock can be tolerated. Contrary to the characterization as “moral hazard” (see introduction), this reduction in abatement is therefore an optimal response. Whether the “moral hazard” argument has traction in a setting in which there is a bias among generations is a key question of the following section: it is conceivable that strategic considerations may lead to a suboptimal reduction on abatement, that is a level of abatement that is lower than A_{cond} . The second instructive feature of the benchmark are its comparative statics. The benchmark abatement level A_{cond} responds in a predictable fashion to an increasing likelihood of a “climate emergency” as well as higher marginal damages of SRM. In both cases, we see higher levels of abatement since $\frac{d}{dp}A_{\text{cond}} > 0$ and $\frac{d}{d\rho}A_{\text{cond}} > 0$. The threshold level of R&D costs also responds predictably to stronger climate damages (the willingness to equip future generation with the technology increases) and higher marginal damages of SRM (opposite effect) since \bar{K} reveal $\frac{d}{d\lambda}\bar{K} > 0$ and $\frac{d}{d\rho}\bar{K} < 0$ where the inequalities follow from (14). Its response to an increase in the probability of a “climate emergency” is less intuitive, however: the sign of $\frac{d}{dp}\bar{K}$ is ambiguous. Increases in p at small levels of p tend to increase the threshold \bar{K} while the opposite is true at high levels of p . The intuition behind this result is that if the “climate emergency” is a low-probability event, the backstop characteristic of

¹⁰The function $C(\Theta, A | D_{\text{uncond}})$ is a convex quadratic function in A with first derivative $-2(\rho - 1)$ at $A = \underline{A}_{\text{crit}}$. This is negative due to the benchmark condition (4). This shows that $\underline{A}_{\text{crit}}$ is the minimizer of $C(\Theta = 1, A | D_{\text{uncond}})$ in $[0, \underline{A}_{\text{crit}}]$

SRM dominates. An increase in the probability in this region strengthens the incentives to provide SRM irrespective of R&D costs. It would not make sense here to sacrifice too much in abatement cost but rather provide the option to counteract climate damages if these turn out to be high. If, however, a certain threshold for p is reached the “climate emergency” is more prevalent and, in order to minimize expected damages, the current generation draws larger levels of abatement. But with higher abatement levels, the role for SRM is reduced and with it, the incentives to pay for R&D. This intuition is reflected in $A_{\text{NoR\&D}}$ being a concave function in p with large increases for small p while A_{cond} is convex in p , thus featuring higher increases for higher levels of the p . The surprising result here is that there is no monotonic relationship between the severity of the responsive climate and the willingness to develop the technology. While the incentives to provide the future generation with SRM increase in the marginal damages $\bar{\lambda}$, that is not necessarily the case in terms of increases in the likelihood p of the “climate emergency”.

4 Strategic distortions

Having established the benchmark equilibrium, we can now study how the rational course of action is impacted when the current generation anticipates a bias of the future generation when assessing SRM damages from geoengineering technologies relative to temperature damages from the carbon stock. In particular, we are interested in the effects on the current generation’s R&D decision and abatement choice $\{\hat{\Theta}, \hat{A}\}$. In the interest of space, we immediately proceed to a full characterization of the intergenerational carbon stock and technology transfer game in β - K -space before providing the analysis that underpins this characterization. We then relate this characterization to the question of which of the five conjectures on current society’s rational course of action survive the consistency test. We find that only some conjectures pass this test. However, a sixth conjecture that has so far not been discussed in the literature emerges as a candidate for how current society could rationally approach this intergenerational decision.

4.1 The full characterization

Figure 2 depicts the subgame perfect equilibria of the intergenerational carbon stock and technology transfer game in β - K -space. The x -axis is defined by the bias parameter β , which denotes the degree to which the future generation discounts SRM damages vis-à-vis temperature damages relative to the current generation. The y -axis is defined by the cost parameter K , which denotes how much the current generation has to sacrifice to provide the future generation with SRM technologies. The upper bound of the x -axis, $\beta = 1$, refers to a level of bias at which the future generation attaches no damages to SRM deployment and therefore compensates temperature damages completely through geoengineering. The upper bound of the y -axis, \bar{K} , derives from expression (14) of the benchmark case and is the level of R&D costs at which the current generation is indifferent between providing the technology

and not.

The benchmark case can be seen in Figure 2 as the line for which $\beta = 0$, the y -axis. The area to the right of the benchmark consists of four distinct zones. The north-east of the area is taken up by zone I in which no R&D is provided for the future generation, i.e. $\Theta = 0$. This zone is associated with a fixed abatement level $A_{\text{NoR\&D}}$ irrespective of the bias β because the future generation's relative valuation of SRM damages is immaterial when no SRM technology is provided. On the y -axis, the boundary of this zone, K^{Ban} , naturally starts at \bar{K} for $\beta = 0$ (see above) and decreases as β increases. As a result, zone I takes up an increasing amount of parameter space for higher levels of bias β . The other three zones, namely zones II, III, and IV, are all associated with the provision of R&D by the current generation, i.e. $\Theta = 1$, but differ with respect to the abatement levels chosen by the current and the deployment profiles chosen by the future generation. The boundaries between zones II and III and zones III and IV are each defined by a critical threshold condition on β and do not depend on K .




Figure 2: Equilibria in the intergenerational decision problem. The axes are the bias of the future generation β and the costs of SRM R&D K .

4.2 Equilibria under SRM R&D ($\Theta = 1$)

With these basic features of Figure 2 established, we now proceed to explain how its geometry depends on the bias parameter β and the cost of R&D, K . To do so, assume for the moment that R&D is always carried out, i.e. $\Theta = 1$.

Small bias (zone II) Then, starting from the benchmark case ($\beta = 0$), we can analyze what happens to abatement and deployment as β increases. Recall that in the benchmark case, equilibrium play consists of abatement of amount $\hat{A} = A_{\text{cond}}$ by the current generation and of a conditional deployment profile D_{cond} by the future generation. Also recall from section 2.2 that as long as the conditional deployment profile is the future generation’s best response, the current generation will find it optimal to stick to the benchmark level of abatement \hat{A} even at higher levels of β .

Expression (6) in section 2 defined the conditions under which sticking to the conditional deployment profile is no longer a best response for the future generation. These conditions were set by two thresholds regarding the current generation’s abatement, $\underline{A}_{\text{crit}}$ and \bar{A}_{crit} . Observe now that both thresholds increase in β . As a result, the abatement level \hat{A} that gave rise to conditional deployment at low levels of β can now give rise to unconditional deployment as a best response. This is the case if β exceeds some critical value β_{crit} . The critical value is the level of β at which $\underline{A}_{\text{crit}}(\beta)$ exceeds $\hat{A} = A_{\text{cond}}$ and is given by

$$\beta_{\text{crit}} := \frac{2(\rho - 1)}{(2 - p)\rho} . \quad (15)$$

Formally, the best-response deployment profile follows

$$D^*(\hat{A}, \beta) = \begin{cases} D_{\text{cond}} & 0 \leq \beta \leq \beta_{\text{crit}} \\ D_{\text{uncond}} & \beta_{\text{crit}} < \beta \leq 1 . \end{cases} \quad (16)$$

A straightforward implication of (16) and the assumption of an interior equilibrium for the benchmark case, which is associated with $\rho > 1$ in equation (11), is the existence of a non-empty interval of $\beta \in [0, \beta_{\text{crit}}]$ for which the benchmark solution will be the equilibrium play despite the presence of a bias. For the geometry of Figure 2, β_{crit} defines the boundary between zones II and III and implies that zone II always exists and that equilibrium play within zone II follows the benchmark case.

Large bias (zone IV) Starting from the opposite end of the interval of bias, as announced in the beginning of this section, we now examine optimal play for the case of $\beta = 1$ when $\Theta = 1$. Again, from expression (6) and (5) in section 2, it is clear that the future generation will always deploy SRM in a way to fully offset the temperature damages. Deployment will therefore be $D = 1 - A$ irrespective of λ . The current generation’s best response to this is to choose abatement level $A_{\text{uncond}} = \frac{\rho}{2}$ because A_{uncond} minimizes the total cost to the current generation in the face of unconditional deployment in the future, given that R&D is carried out. This abatement level is strictly greater than that which is optimal under conditional deployment, which reflects that it is cheaper for the current generation to counteract the future generation’s “abuse” of SRM technologies through increased abatement, and hence less SRM deployment. Just as with the benchmark case, the question arises over what interval of β this equilibrium play persists, now going in the opposite direction (decreasing β). In the benchmark case, we observed that as long as conditional deployment remained the best response of the future generation, the current generation did not adjust its abatement.

The same logic applies here: As long as unconditional deployment remains the best response of the future generation, the current generation does not deviate from A_{uncond} . How low can β be for unconditional deployment to remain the best response to A_{uncond} ? Broadly analogously to the previous case, the critical value now is the level of β at which $\underline{A}_{\text{crit}}(\beta)$ falls below A_{uncond} and is given by

$$\bar{\beta} := \frac{2(\rho - 1)}{\rho} . \quad (17)$$

Casual inspection makes clear that there is always a non-empty interval of $\beta \in [\bar{\beta}, 1]$ for which a combination of abatement at A_{uncond} by the current generation and unconditional deployment by the future generation will be the equilibrium play when $\Theta = 1$. For the geometry of Figure 2, $\bar{\beta}$ defines the boundary between zones IV and III and implies that zone IV always exists.

Medium bias - abatement as a strategic instrument (zone III) The strategically subtlest case arises for degrees of bias between the two boundaries β_{crit} and $\bar{\beta}$. This is always a non-empty interval (except for the degenerate case that $p = 1$) as is obvious from inspecting the boundary expressions. Zone III therefore exists. In order to understand equilibrium play in this interval, recall that β_{crit} denoted the threshold for values of β above which A_{cond} no longer induced conditional deployment as a best response and that $\bar{\beta}$ denoted the threshold for values of β below which A_{uncond} no longer induced unconditional deployment as a best response. To find the amount for abatement that is both cost-minimizing and subgame-perfect in this interval, it is helpful to start from the threshold $\underline{A}_{\text{crit}}$ which separates those abatement levels that induce unconditional (to the left of $\underline{A}_{\text{crit}}$) from those that induce conditional deployment (to the right of $\underline{A}_{\text{crit}}$). It should be clear that for unconditional deployment, if subgame perfection was not a requirement, the cost-minimizing level of abatement A_{uncond} would lie to the right of $\underline{A}_{\text{crit}}$. Subgame perfection in this β -zone, however, requires an abatement level not above $\underline{A}_{\text{crit}}$ and thus smaller than A_{uncond} . At $\underline{A}_{\text{crit}}$, marginal total costs are still decreasing. Therefore, $\underline{A}_{\text{crit}}$ must be the cost-minimizing subgame perfect choice of abatement if the current generation wants to induce unconditional deployment.

The case of the current generation wanting to induce conditional deployment is a mirror image of the case above: If subgame perfection was not a requirement, the cost-minimizing level of abatement A_{cond} would lie to the left of $\underline{A}_{\text{crit}}$. However, subgame perfection in this zone requires an abatement level not below $\underline{A}_{\text{crit}}$ and thus greater than A_{cond} . At $\underline{A}_{\text{crit}}$, marginal total costs are already increasing. Therefore, $\underline{A}_{\text{crit}}$ must be the cost-minimizing subgame perfect choice of abatement if the current generation wants to induce conditional deployment. Taken together, $\underline{A}_{\text{crit}}$ is the cost-minimizing subgame perfect choice of abatement in zone III. At $\underline{A}_{\text{crit}}$, the conditional and unconditional deployment profiles become indistinguishable because the level of unconditional deployment in the case of carbon-insensitive climate becomes zero. Note that in contrast to the cost-minimizing subgame perfect abatement levels in zones II and IV which are independent of β , equilibrium play in zone III features abatement that increases in β .

We now have a complete picture of equilibrium abatement and SRM deployment in zones

II, III, and IV, i.e. under the assumption that R&D is carried out by the current generation.¹¹ What remains to be shown is under what conditions the provision of R&D indeed constitutes a rational course of action by the current generation. Given the presence of strategic distortions in equilibrium play in zones II, III, and IV depending on the degree of the bias β , not providing SRM technologies in the first place may be the cost-minimizing choice. The key determinant of this decision is the cost of R&D. To this we turn now.

4.3 Equilibrium R&D decision ($\Theta = 0$ vs. $\Theta = 1$)

Recall from above that if no R&D is carried out ($\Theta = 0$), then there is a unique abatement level $A_{\text{NoR\&D}}$ chosen by the current generation. The future generation is forced into never deploying SRM. Therefore, $D = 0$. The expected total costs from this course of action are $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$. For equilibrium play under $\Theta = 1$ to be selected as the rational course of action, it needs to feature lower cost than $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$ despite involving R&D cost K . What is required therefore is to compare total expected costs in each of the zones II, III, and IV with the total expected costs from not carrying out R&D and to determine the condition on K for R&D still to be provided. The condition that results for each zone each constitutes a segment of the boundary K^{Ban} in Figure 2 that separates the no-R&D zone I from the three R&D zones II, III, and IV.

We first compare the boundary between zone II and zone I. The costs we have to compare are $C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}})$ and $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$. We are looking for a condition under which no R&D involves lower cost than providing R&D together with abatement of amount A_{cond} . Simple algebraic operations translate this comparison into a condition on K of the form

$$K > K_{II}^{\text{Ban}}(\beta) = \bar{K} - \frac{p\rho^2}{4\bar{\lambda}^2}\beta^2. \quad (18)$$

The boundary between zone III and zone I involves a comparison of costs $C(\Theta = 1, \underline{A}_{\text{crit}} | D_{\text{cond}})$, which are by the definition of $\underline{A}_{\text{crit}}$ equal to $C(\Theta = 1, \underline{A}_{\text{crit}} | D_{\text{uncond}})$, and $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$. For no R&D to be cost-minimizing requires that

$$K > K_{III}^{\text{Ban}}(\beta) = \bar{K}_{III} - \frac{(p + (2-p)\bar{\lambda}^2)\rho^2}{4\bar{\lambda}^2}(\beta - \beta_{III})^2 \quad (19)$$

with $\beta_{III} = \frac{\rho-1}{\rho} \frac{2\bar{\lambda}^2}{2\bar{\lambda}^2 - p(\bar{\lambda}^2 - 1)}$ and $\bar{K}_{III} = \bar{K} - \frac{p(\rho-1)}{(2-p)(2\bar{\lambda}^2 - p(\bar{\lambda}^2 - 1))}$. It is easy to show that $\beta_{III} \in [0, \beta_{\text{crit}}]$.

The boundary between zone IV and zone I involves a comparison of costs $C(\Theta = 1, A_{\text{uncond}} | D_{\text{uncond}})$ and $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$. For no R&D to be cost-minimizing requires that

$$K > K_{IV}^{\text{Ban}}(\beta) = \bar{K}_{IV} - \frac{p + (1-p)\bar{\lambda}^2}{4\bar{\lambda}^2}\beta^2 \quad (20)$$

¹¹An interested reader might ask how the thresholds that separate zones II, III and IV, β_{crit} and $\bar{\beta}$, depend on model parameters. We have $\frac{d}{dp}\beta_{\text{crit}} > 0$ and $\frac{d}{d\rho}\beta_{\text{crit}} > 0$. Increasing technology damages and likelihood of bad outcomes thus both expand the first region. The intuition is that both changes lead to a higher benchmark level \hat{A} which is then more robust to deviations in SRM damage assessment β . The second threshold, $\bar{\beta}$, depends only on ρ . The derivative is $\frac{d}{d\rho}\bar{\beta} > 0$ and larger than $\frac{d}{d\rho}\beta_{\text{crit}}$ because higher SRM damages reduce its desirability stronger if it is to be always deployed compared to conditional deployment. Thus, increasing damages of the technology make region IV smaller and enlarge both region II and region III.

where $\bar{K}_{IV} = \bar{K} + \frac{(1-p)(\rho-1)^2}{2-p}$.

Note that the segments derived above form a continuous boundary K^{Ban} since $K_{II}^{\text{Ban}} = K_{III}^{\text{Ban}}$ and $K_{III}^{\text{Ban}} = K_{IV}^{\text{Ban}}$ at the relevant thresholds β_{crit} and $\bar{\beta}$, respectively. Also note that K^{Ban} varies monotonously in the bias β since $\frac{d}{d\beta} K_i^{\text{Ban}} < 0$ for all segments i . This confirms our earlier intuition that for the same level of R&D costs K , an increase in the bias β renders not carrying out R&D a relative more attractive course of action because the strategic distortions between the generations increase in β .¹²

4.4 Surviving conjectures

In this final substantive section of the paper, we relate the full characterization of the intergenerational carbon stock and technology transfer game back to the conjectures about the rational course of action for current society that appear in the current literature and that we reviewed in the introduction. Based on the preceding analysis, it is now possible to make statements about the degree to which different conjectures can be replicated in a simple model that captures the four common elements of (i) intergenerational altruism, (ii) possible pro-SRM bias, (iii) costly abatement and R&D, and (iv) intergenerational transmission of carbon stocks and technology under uncertainty about the climate system.

We identified five conjectures that share these common elements, yet differ markedly in their conclusions. One conjecture termed “arming the future” postulates that current society will find it rational to provide SRM technologies as a backstop for inadvertent climate outcomes while providing significant abatement. We were able to replicate this conjecture without problems and used it as the benchmark case for understanding the strategic distortions introduced by anticipating a bias among the future generation that implicitly favors the deployment of SRM. By contrast, the “abatement invariance” conjecture, namely that abatement should be unaffected by the decision to conduct R&D into SRM (Bunzl 2009), does not pass the consistency check of our model. The common elements that underpin all conjectures make it almost inevitable that abatement will change as soon as a decision in favor of SRM is taken since both abatement and SRM address the same problem, but feature different cost structures. We also fail to find support for the hypothesis that the anticipation of a large bias among the future generation that favors SRM deployment leads to the current generation slashing abatement. The reason is that if the current generation cares sufficiently for the future generation to take abatement action today, then it will also care sufficiently to take into account the negative side effects of large-scale geoengineering that would result from slashing abatement.

¹²An interested reader might again ask how the boundary K^{Ban} depends on key model parameter. While the comparative statics are algebraically messy, we can build on the continuity and monotonicity of the boundary observed above using a simple trick. This trick involves examining the comparative statics of point $K_{IV}^{\text{Ban}}(\beta = 1)$, i.e. the point at which the boundary coincides with the upper bound on the bias β . By continuity and monotonicity, the comparative statics of this point are qualitatively the same as the comparative statics of the entire boundary. The comparative statics are that $\frac{d}{dp} K_{IV}^{\text{Ban}}(\beta = 1) > 0$, $\frac{d}{d\lambda} K_{IV}^{\text{Ban}}(\beta = 1) > 0$ and $\frac{d}{d\rho} K_{IV}^{\text{Ban}}(\beta = 1) < 0$. A greater likelihood or severity of a “climate emergency” increase the relative size of those zones that involve R&D into SRM while higher SRM damages have the opposite effect.

On the same grounds that challenge the last two conjectures, the model generates a much more upbeat view about the link between SRM R&D and abatement in the presence of a bias. As zone II in Figure 2 illustrates, if the bias is sufficiently weak ($\beta < \beta_{\text{crit}}$), the benchmark abatement level \hat{A} will be maintained in equilibrium. In zone III with $\beta_{\text{crit}} < \beta \leq \bar{\beta}$, the bias is strong enough to render the initial abatement level \hat{A} inconsistent with conditional use by future generation. Anticipating that the technology will be used indiscriminately, current generation reacts by *increasing* abatement to the level $A_{\text{crit}}(\beta)$ at which D_{cond} and D_{uncond} merge. The larger the bias β , the larger the necessary increase in abatement in order to induce the future generation into this specific use of the SRM technology. Finally, in zone IV ($\beta > \bar{\beta}$), a further increase of β has no additional impact on the abatement level which remains at A_{uncond} , irrespective of β .

Interestingly, A_{uncond} can be even higher than the technology denial abatement level $A_{\text{NoR\&D}}$. We have

$$A_{\text{uncond}} > A_{\text{NoR\&D}} \quad \text{if and only if} \quad \rho > \frac{2 + 2p(\bar{\lambda}^2 - 1)}{2 + p(\bar{\lambda}^2 - 1)}. \quad (21)$$

This is, of course, only meaningful if this condition does not preclude technology provision equilibria with $\Theta = 1$. We can actually find equilibria in zone IV that feature “overabatement” relative to $A_{\text{NoR\&D}}$.¹³ Figure 3 summarizes the abatement level over all equilibria.




Figure 3: Abatement in case of $\Theta = 1$ as a function of the bias β . Whether A_{uncond} is larger or smaller than the optimal abatement level under no technology provision, $A_{\text{NoR\&D}}$, depends on model parameter.

In addition to this new equilibrium with SRM R&D provision and abatement increases, Figure 2 captures another consistent conjecture supported by the premises of the model. This

¹³Even though $\rho > \frac{2+2p(\bar{\lambda}^2-1)}{2+p(\bar{\lambda}^2-1)}$ implies that $K_{IV}^{\text{Ban}}(\beta = 1) < 0$ and thus a higher abatement level $A_{\text{uncond}} > A_{\text{NoR\&D}}$ will not realize for $\beta = 1$, there are equilibria with $\beta < 1$. Take, for instance, $p = \frac{1}{10}$, $\bar{\lambda} = 2$ and $\rho = 1.15$. This leads to $A_{\text{uncond}} > A_{\text{NoR\&D}}$; at the same time $K_{IV}^{\text{Ban}}(\bar{\beta}) > 0$, implying that the proposed equilibrium with $\Theta = 1$ and $A = A_{\text{uncond}} > A_{\text{NoR\&D}}$ exists in zone IV.

is that current society will rationally choose not to engage in SRM R&D activities, i.e. a ban on SRM R&D is another course of action that our model confirms in a robust fashion (zone I). This course of action becomes particularly relevant for a large anticipated bias and for high development costs associated with SRM technologies.

We reserve a final comment on the conjecture that investment in SRM R&D will reduce abatement and thus give rise to “moral hazard”. As we point out in section 3, comparing abatement levels under positive and no R&D in the benchmark case shows that abatement in the presence of SRM R&D is smaller. However, rather than constituting a situation where an economic party is imposing a risk burden on some other party without proper compensation, this is an efficient decision that reflects a rational readjustment of its abatement efforts by the current generation.

In sum, our model is able to successfully replicate three out of the five conjectures reviewed in the introduction. “Arming the future”, a R&D ban, and abatement reductions relative to a situation with no SRM R&D all constitute courses of action for current society that a model capturing the common elements among the conjectures can generate as consistent conclusions. Two conjectures, one postulating “abatement invariance” with respect to SRM R&D decisions and one postulating a drastic abatement reduction by current society, cannot be replicated and appear inconsistent with the basic premises in the literature. At a minimum, this means that auxiliary hypotheses and assumptions are necessary in order to demonstrate the consistency of these conjectures. At a maximum, the conclusion is that these conjectures are erroneous. In addition to the replication test of five existing conjectures, our model also shows that the same four common elements that underpin these conjectures give rise to a novel conjecture: It can be a rational course of action of current society to provide more abatement the higher the degree of anticipated bias. In fact, it is possible for this optimal abatement level to even exceed the level that society would rationally provide in the absence of SRM R&D.

5 Concluding discussion

If feasible, human interventions into the Earth’s climate system would represent a novel method for future generations to limit the damages of the atmospheric carbon stock that they will inherit from current society. Such interventions, running under the term of “geo-engineering” or “climate engineering”, would create undesirable side effects of their own, but could conceivably be considered a realistic option. Solar radiation management (SRM) technologies are considered one of the likeliest forms of geoengineering to be developed and deployed.

R&D into SRM raises the possibility of passing on to the next generation not just abatement efforts in the form of a carbon stock that is below business-as-usual. In addition, or instead, the current generation could pass on a technology that can partially remedy the damages of excessive atmospheric carbon stocks. Natural scientists, philosophers, ethicists, and other scholars have started to develop several conjectures on how current society should decide on the right combination of SRM R&D and abatement efforts. By contrast, economists

have so far not examined the intergenerational issues implicit in this technological possibility. This is despite the fact that economics has the potential to contribute substantially to the discussion thanks to its powerful conceptual tools for the analysis of intergenerational transfers. It is also despite the fact that the issue of intergenerational technology transfers to address the intergenerational transfer of stock pollution has so far attracted little attention in economics, in contrast to other types of intergenerational transfers.

The results of the present paper are a first step towards addressing the intergenerational issues of SRM R&D in the context of climate change. By developing a simple analytical model that formalizes several common elements in the wider debate about the correct course of action for the current generation in this intergenerational game, we harness the powers of game theoretic analysis for the purpose of understanding more about the problem. The comparison of the diverse conjectures about the rational course of action to develop SRM capabilities and the attempts to replicate their logic in a tractable model adds rigor to the debate and allows distinguishing between consistent conjectures and those that require either auxiliary hypotheses or correction. The same rigor allows us to identify solutions that have so far escaped attention, such as our finding that abatement may actually be higher when SRM capabilities are developed, and challenges loose argumentation, such as the claim regarding the presence of “moral hazard”.

For the economic debate, this paper represents a starting point for considering more systematically than before how to integrate the development of technological capabilities into intergenerational models. This integration is far from complete. It also adds to an emerging literature that examines the potential role of behavioral factors such as hyperbolic discounting, paternalism, and bounded rationality in the context of interactions across generations. The context of geoengineering provides a conducive and potentially consequential setting for more research of this type.

References

- [1] S. Agrawala and S. Fankhauser. *Economic Aspects of Adaptation to Climate Change: Costs, Benefits and Policy Instruments*. OECD Publishing, 2008.
- [2] American Meteorological Society. Geoengineering the climate system: A policy statement of the american meteorological society. *Bulletin of the American Meteorological Society*, 90(9):1369–1370, 2009.
- [3] G. Bala, K. Caldeira, and R. Nemani. Fast versus slow response in climate change: implications for the global hydrological cycle. *Climate dynamics*, 35(2):423–434, 2010.
- [4] G.A. Ban-Weiss and K. Caldeira. Geoengineering as an optimization problem. *Environmental Research Letters*, 5(3):034009, 2010.
- [5] S. Barrett. The incredible economics of geoengineering. *Environmental and Resource Economics*, 39(1):45–54, 2008.

- [6] J.E. Bickel and S. Agrawal. Reexamining the economics of aerosol geoengineering. In review at *Climatic Change*, August 2012.
- [7] D. Bodansky. May we engineer the climate? *Climatic Change*, 33(3):309–321, 1996.
- [8] M. Bunzl. Researching geoengineering: should not or could not? *Environmental Research Letters*, 4(4):045104, 2009.
- [9] K. Caldeira and L. Wood. Global and arctic climate engineering: Numerical model studies. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1882):4039–4056, 2008.
- [10] J.R. Fleming. *Fixing the sky: the checkered history of weather and climate control*. Columbia University Press, 2010.
- [11] S.M. Gardiner. Is ‘arming the future’ with geoengineering really the lesser evil? some doubts about the ethics of intentionally manipulating the climate system. *Climate ethics*, pages 323–336, 2010.
- [12] S.M. Gardiner. Some early ethics of geoengineering the climate: a commentary on the values of the royal society report. *Environmental Values*, 20(2):163–188, 2011.
- [13] M. Goes, N. Tuana, and K. Keller. The economics (or lack thereof) of aerosol geoengineering. *Climatic change*, 109(3):719–744, 2011.
- [14] T. Goeschl and G. Perino. Innovation without magic bullets: Stock pollution and R&D sequences. *Journal of Environmental Economics and Management*, 54(2):146–161, 2007.
- [15] B. Hale. Getting the bad out: Remediation technologies and respect for others. In W.P. Kabasenche, M. O’Rourke, and M.H. Slater, editors, *The Environment. Philosophy, Science, and Ethics*. MIT Press (MA), 2012.
- [16] D. Jamieson. Ethics and intentional climate change. *Climatic Change*, 33(3):323–336, 1996.
- [17] D.W. Keith. Geoengineering the Climate: History and Prospect. *Annual Review of Energy and the Environment*, 25(1):245–284, 2000.
- [18] D.W. Keith. Engineering the planet. In *Climate Change Science and Policy*. Schneider, S. and Mastrandrea, M., 2012.
- [19] D.W. Keith, E. Parson, and M.G. Morgan. Research on global sun block needed now. *Nature*, 463(7280):426–427, 2010.
- [20] G. Klepper and W. Rickels. The real economics of climate engineering. *Economics Research International*, 2012.

- [21] B. Kravitz, A. Robock, L.D. Oman, G.L. Stenchikov, and A. Marquardt. Sulfuric acid deposition from stratospheric geoengineering with sulfate aerosols. *Journal of Geophysical Research*, 114(D14), 2009.
- [22] D.M. Kreps. *Game theory and economic modelling*. Clarendon Press Oxford, 1990.
- [23] A.M. Mercer, D.W. Keith, and J.D. Sharp. Public understanding of solar radiation management. *Environmental Research Letters*, 6(4):044006, 2011.
- [24] J.B. Moreno-Cruz. Mitigation and the geoengineering threat. Available at: <http://works.bepress.com/morenocruz/3>, 2010.
- [25] J.B. Moreno-Cruz and D.W. Keith. Climate policy under uncertainty: a case for solar geoengineering. *Climatic Change*, 2012. Available at <http://www.springerlink.com/content/1824m4unw0472803/>.
- [26] J.B. Moreno-Cruz, K.L. Ricke, and D.W. Keith. A simple model to account for regional inequalities in the effectiveness of solar radiation management. *Climatic change*, 110(3):649–668, 2012.
- [27] W.D. Nordhaus. A review of the “stern review on the economics of climate change”. *Journal of Economic Literature*, 45(3):686–702, 2007.
- [28] K.L. Ricke, M.G. Morgan, and M.R. Allen. Regional climate response to solar-radiation management. *Nature Geoscience*, 3(8):537–541, 2010.
- [29] A. Ridgwell, C. Freeman, and R. Lampitt. Geoengineering: taking control of our planet’s climate? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1974):4163–4165, 2012.
- [30] G.H. Roe and M.B. Baker. Why is climate sensitivity so unpredictable? *Science*, 318(5850):629–632, 2007.
- [31] T.C. Schelling. The economic diplomacy of geoengineering. *Climatic Change*, 33(3):303–307, 1996.
- [32] J.G. Shepherd. *Geoengineering the climate: Science, governance and uncertainty*. Royal Society, 2009.
- [33] N. Stern. *The Stern review report on the economics of climate change*. Cambridge University Press, 2006.
- [34] R.S.J. Tol. Equitable cost-benefit analysis of climate change policies. *Ecological Economics*, 36(1):71–85, 2001.
- [35] US National Academy of Sciences. *Changing Climate: Report of the Carbon Dioxide Assessment Committee*. National Academy Press, 1983.

- [36] D.G. Victor. On the regulation of geoengineering. *Oxford Review of Economic Policy*, 24(2):322–336, 2008.