Journal of Economic Behavior & Organization xxx (2016) xxx–xxx

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

# Climate tipping points and solar geoengineering☆

Garth Heutel[a], Juan Moreno-Cruz[b,*], Soheil Shayegh[c]

[a] Department of Economics, Georgia State University, Atlanta, GA, USA
[b] School of Economics, Georgia Institute of Technology, Atlanta, GA, USA
[c] Department of Global Ecology, Carnegie Institution of Washington, Stanford, CA, USA

ARTICLE INFO

*Article history:*
Received 15 September 2015
Received in revised form 2 July 2016
Accepted 5 July 2016
Available online xxx

*JEL classification:*
C6
D6
D8
D9
Q54
Q58

*Keywords:*
Solar geoengineering
Climate tipping points
Climate policy

ABSTRACT

We study optimal climate policy in the presence of climate tipping points and solar geoengineering. Solar geoengineering reduces temperatures without reducing greenhouse gas emissions. Climate tipping points are irreversible and uncertain events that can alter the dynamics of the climate system. We analyze three different rules related to the availability of solar geoengineering, and we model three distinct types of tipping points. Before reaching the tipping point, the introduction of solar geoengineering reduces the amount of mitigation, lowers temperatures and increases carbon concentrations. The capacity of solar geoengineering to deal with climate damages depends on the type of tipping point. Solar geoengineering is most effective at dealing with tipping points that affect the responsiveness of temperature to carbon, and it is least effective at dealing with tipping points that cause direct economic losses.



## 1. Introduction

The accumulation of greenhouse gases in the atmosphere is associated with an increase in Earth's surface temperature, affecting economic performance and ecosystems as a whole. As temperature rises, the probability of crossing a climate tipping point (CTP) increases. CTPs are large, rare, difficult to predict, and irreversible disturbances of the carbon-climate system. The most common examples of such events are the collapse of the thermohaline circulation or the disintegration of the West Antarctic Ice Sheet. Solar geoengineering (SGE), and more specifically solar radiation management, has been proposed as a way of limiting the probability of reaching a climate tipping point. By reducing the amount of radiation reaching Earth's surface, temperatures can be kept at a level below which tipping points can occur even without reducing greenhouse gas concentrations. In this paper we analyze optimal climate policy in the presence of CTPs when both emissions reductions (mitigation) and SGE are available, using both a theoretical model and numerical simulations.

As a first stage in our analysis, we build a parsimonious analytical model of climate change economics with CTPs and SGE. We model a CTP as an irreversible event that changes the dynamics of the climate-carbon system, resulting in an output

☆ We thank Derek Lemoine and seminar participants at Tennessee, Stanford, Harvard, and the Southern Economic Association annual meeting for comments. This research was partly funded by the an SRG-A from the Ivan Allen College of Liberal Arts.

* Corresponding author.
*E-mail addresses:* gheutel@gsu.edu (G. Heutel), morenocruz@gatech.edu (J. Moreno-Cruz), sshayegh@carnegiescience.edu (S. Shayegh).

http://dx.doi.org/10.1016/j.jebo.2016.07.002
0167-2681/

loss relative to the state of the world before the CTP is reached. The planner's problem is solved using stochastic dynamic programming techniques that allow us to accommodate the post-CTP transition in the system. The probability of reaching the tipping point is a function of atmospheric temperature. An important characteristic of our model is that while mitigation efforts affect temperatures only in future periods, SGE affects temperatures in the *same* period it is implemented. We explore three different SGE rules currently discussed in the governance literature. The first rule is a *Ban*, in which society chooses not to engage in SGE under any circumstances. In the second rule, SGE is freely used in combination with mitigation. We call this the *Unconstrained* rule. Third, we consider a rule where SGE is allowed only when temperatures surpass the climate tipping point. This is called the *Reparation* rule, since SGE can be thought of as only a "last-resort" policy in the event that the tipping point is reached. Using the analytical model, we identify different roles for mitigation and SGE. While both instruments help reduce damages before and after reaching the CTP, SGE can reduce the risk of crossing the temperature threshold more quickly than can mitigation.

We then incorporate SGE into a quantitative integrated assessment model, the DICE (Dynamic Integrated Climate-Economy) model, following Heutel et al. (2015), to simulate a richer set of alternative scenarios allowing for both mitigation and SGE in the presence of CTPs. In the quantitative model, we relax many assumptions of the analytical model and confirm the results presented in the theory. The simulation model allows us to consider three distinct types of tipping points: a Climate Feedback CTP in which the climate sensitivity (the responsiveness of temperature to the carbon stock) is changed after the CTP, a Carbon Sink CTP in which the carbon dynamics are changed after the CTP, and an Economic Loss CTP in which there is a direct gross output loss from the CTP. We quantify the effects of alternative SGE rules under the three CTP specifications on several outcome variables, including temperature, carbon stock, and the optimal carbon tax. The *Ban* rule yields a carbon tax that is twice as high as the tax under the other two rules. Under the *Unconstrained* rule, the risks associated with the tipping point are largely avoided. The *Reparation* rule reduces damages and carbon taxes only after the threshold is crossed but affects the trajectory of the pre-CTP policy, since it makes triggering a CTP less costly. Under all rules, and contrary to what has been expressed previously in the geoengineering literature, a substantial amount of mitigation is optimal to deal with the risks of climate change.

Our approach closely resembles that of Lemoine and Traeger (2014). That paper uses a recursive version of DICE to consider CTPs where policymakers learn about the position of the tipping point and where the costs associated with crossing the tipping point are a function of the state of the economy at the time it is crossed. Like our paper, Lemoine and Traeger (2014) model a Climate Feedback CTP and a Carbon Sink CTP. To that, we add the Economic Loss CTP, as in Cai et al. (2013) and Lontzek et al. (2015). Furthermore, we add SGE to the model. To the best of our knowledge, our paper is the first to incorporate SGE in a model with CTPs.[1]

The use of SGE as part of the portfolio of options has been suggested in the literature under diverse scenarios. The use of SGE as an insurance against catastrophic climate change has been proposed early in the literature (Keith, 2000; Victor, 2008; Keith et al., 2010; Moreno-Cruz and Keith, 2013). The idea of SGE as a complement to mitigation is proposed in the literature as a way to achieve any given temperature level at lower costs for society (Wigley, 2006; Moreno-Cruz and Keith, 2013; Heutel et al., 2015). Finally, banning SGE has been proposed because of the large uncertainties surrounding the unintended consequences of SGE implementation (Barrett, 2008; Blackstock and Long, 2010), the asymmetry of impacts this intervention may have (Moreno-Cruz, 2015; Moreno-Cruz and Keith, 2013), and the difficulty in regulating implementation (Victor, 2008). In our optimal policy context, we abstract from the political economy of implementing solar geoengineering, a difficult task as suggested by the large literature on the governance of geoengineering (e.g. Barrett et al., 2014).

A unique contribution of this paper in terms of methods is to model stochastic parameter values, rather than merely performing sensitivity analyses. Other studies have traditionally considered only sensitivity analyses but failed to develop a solution for the stochastic model. Among the papers that have actually modified DICE to include stochastic parameters are Baker and Solak (2011), Kolstad (1996) and Lemoine and Traeger (2014). However, none of these papers have included SGE as a policy option.

Other papers have added SGE to integrated assessment models and examined the policy implications. Bickel and Lane (2009) and Goes et al. (2011) make several modifications to the DICE model, including allowing SGE and refining the climate dynamics. Their specification imposes an exogenous intermittency in SGE which makes it less effective. They present summaries of policies with an optimal mix of mitigation and SGE (subject to the intermittency). In contrast to Goes et al. (2011), Bickel and Agrawal (2013) find that under some scenarios a substitution of SGE for mitigation can pass a cost–benefit test. Gramstad and Tjøotta (2010) include SGE in DICE and conduct a cost–benefit analysis of SGE under various assumptions about the level undertaken and its costs. Emmerling and Tavoni (2013) use a different integrated assessment model, WITCH, to model SGE and mitigation policy.[2] None of these papers consider the possibility of climate tipping points.

The rest of the paper is organized as follows. In Section 2, we present our analytical model and its predictions. Section 3 describes our numerical model and solution method, and Section 4 reports numerical results.

[1] By contrast, our earlier paper (Heutel et al., 2015) and several others add SGE to an integrated assessment model but without CTPs.
[2] See Heutel et al. (2015) for a more thorough comparison between our approach and previous papers introducing SGE in integrated assessment models.

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002

## 2. Theoretical model

We consider the case of a regulator who solves an infinite-horizon optimization problem with the goal of minimizing the total costs of climate change. In the model, the temperature threshold that triggers CTPs is uncertain, there are different types of tipping points, and SGE and mitigation are imperfect substitutes.

Optimal policy depends on the state of the world and the dynamics of the carbon-climate system. We use the following set of first order difference equations to represent the dynamics of the carbon-climate system:

$$S_{t+1} = e_t^{BAU} - m_t + (1-\delta_t)S_t \tag{1}$$

$$T_{t+1} = \lambda_t \left( \ln\left(\frac{S_{t+1}}{S_0}\right) - \theta_t g_{t+1} \right) + (1-\gamma_t)T_t \tag{2}$$

$$S_0 > 0 \ \text{ and } \ T_0 > 0 \ \text{ given.}$$

Eq. (1) captures carbon-cycle dynamics. $S_t$ is the stock of carbon in the atmosphere, $e_t^{BAU}$ is business-as-usual emissions of greenhouse gases, $m_t$ is mitigation, and $\delta_t$ is the absorption capacity of the planet. Eq. (2) shows how temperature, $T_t$, responds to changes in radiative forcing at time $t$. The radiative forcing potential of carbon dioxide depends on the carbon stock $S_t$ relative to its pre-industrial level $S_0$. $g_t$ is the amount of SGE implemented at time $t$ expressed in units of radiative forcing, and $\theta_t \in \{0, 1\}$ represents the rule regarding the availability of SGE at time $t$: $\theta_t = 1$ when SGE is available and $\theta_t = 0$ when it is not. $\lambda_t$ represents the climate sensitivity of the system that transforms radiative forcing into temperature levels. Finally, some fraction of the heat stored in the atmosphere escapes and some other fraction is absorbed by the oceans; this effect is captured by the term $\gamma_t T_t$, where $\gamma_t$ is the heat transfer parameter (Naevdal and Oppenheimer, 2007).

Eqs. (1) and (2) represent the inertia of the climate-carbon system in a simple way that highlights the main difference between mitigation and SGE: temperature in period $t+1$ is a function of mitigation in period $t$; SGE in period $t$, on the other hand, affects temperatures in the same period. Therefore, mitigation efforts in period $t$ create benefits in future periods but can do little to reduce the warming we have already committed to for this period, while SGE can alter temperatures more quickly, thus reducing the inertia of the climate-carbon cycle. This is the most important difference between mitigation and SGE in this context; SGE's capacity to deal with climate risks lies in the fact that the carbon-climate system responds to its implementation more quickly than it responds to mitigation.

We model a climate tipping point as an irreversible change in the climate-carbon system that occurs after a given temperature threshold is crossed. The CTP is triggered by reaching a specific temperature; it is not a function of carbon concentrations or any other variables. We define the vector $\nu_t$ to capture the state of the climate system at time $t$. Before the temperature threshold is crossed, $\nu_t = \nu$, and after it is crossed, $\nu_t = \tilde{\nu}$. The parameters that are in $\nu_t$, and how they change after the tipping point is crossed, can vary depending on the type of tipping point.[3] We describe the effects of the tipping points in more detail when describing the numerical model later.

Total costs are the sum of the costs of implementing mitigation, $m_t$, and SGE, $g_t$, plus the damages associated with climate change. The implementation costs are given by $c(m_t, g_t)$, where $c_m > 0$, $c_{mm} > 0$, $c_g > 0$, $c_{gg} > 0$ and $c_{mg} = 0$. Damages are given by $D(T_t, S_t, g_t)$ and are a function of the current state of the world. They are increasing and convex in temperature and atmospheric carbon concentrations, that is $D_T > 0$, $D_{TT} > 0$, $D_S > 0$, $D_{SS} > 0$. Solar geoengineering also create damages, $D_g > 0$ and $D_{gg} > 0$.

The exact location of the temperature threshold leading to a CTP is unknown to the regulator, but the probability of crossing the threshold is known to be an increasing function of the temperature at time $t$. In this specification of CTPs, the probability of crossing the threshold is captured by an endogenous hazard function given by $h(T_{t+1})$. This hazard function captures the idea that as temperature increases, the likelihood of crossing the threshold in the next period also increases.

We solve the regulator's problem via backwards induction. We first analyze the situation after the CTP has been crossed; we call this the *post-CTP* regime. Then we analyze the situation before the CTP has been crossed; we call this the *pre-CTP* regime. After the threshold is crossed, the value function is given by $V(S_t, T_t, \tilde{\nu})$, where $\tilde{\nu}$ captures the state of the dynamics of the climate-carbon system. We obtain the solution to $V(S_t, T_t, \tilde{\nu})$ by solving the following Bellman equation:

$$V(S_t, T_t, \tilde{\nu}) = \min_{m_t, g_t} \{c(m_t, g_t) + D(T_t, S_t, g_t) + \beta_t V(S_{t+1}, T_{T+1}, \tilde{\nu})\} \tag{3}$$

subject to Eqs. (1) and (2). As the last argument of the value function we write $\tilde{\nu}$ rather than $\nu_t$ to indicate that, after the CTP has been crossed, these parameters are fixed at their post-CTP values.

Before crossing the CTP, the Bellman equation of this problem is as follows:

$$V(S_t, T_t, \nu) = \min_{m_t, g_t} \{c(m_t, g_t) + D(T_t, S_t, g_t) + \beta[(1 - h(T_{t+1}))V(S_{t+1}, T_{t+1}, \nu) + h(T_{t+1})V(S_{t+1}, T_{t+1}, \tilde{\nu})]\} \tag{4}$$

where $V(S_t, T_t, \nu)$ is the value function in period $t$ given the pre-CTP state of the world. With probability $1 - h(T_{t+1})$ the system remains unchanged (so $\nu_{t+1} = \nu$), and with probability $h(T_{t+1})$ the CTP is crossed (so $\nu_{t+1} = \tilde{\nu}$).

[3] For example, a CTP that changes the carbon absorptive capacity of the planet would reduce $\delta_t$; a CTP that creates a climate feedback loop would increase $\lambda_t$.

The first order condition with respect to mitigation is[4]:

$$c_m(m_t, g_t) + \beta \left( V_S(S_{t+1}, T_{t+1}, \nu)\frac{\partial S_{t+1}}{\partial m_t} + \underbrace{h(T_{t+1})[V_S(S_{t+1}, T_{t+1}, \tilde{\nu}) - V_S(S_{t+1}, T_{t+1}, \nu)]\frac{\partial S_{t+1}}{\partial m_t}}_{DWI^S} + \underbrace{h_T(T_{t+1})\frac{\partial T_{t+1}}{\partial m_t}[V(S_{t+1}, T_{t+1}, \tilde{\nu}) - V(S_{t+1}, T_{t+1}, \nu)]}_{MHE^S} \right) = 0 \tag{5}$$

where $\frac{\partial S_{t+1}}{\partial m_t} < 0$ and $\frac{\partial T_{t+1}}{\partial m_t} < 0$. The first order condition with respect to SGE is:

$$c_g(m_t, g_t) + \underbrace{D_g(T_t, S_t, g_t) - \lambda_t\theta_t D_T(T_t, S_t, g_t)}_{IDR} + \beta \left( V_T(S_{t+1}, T_{t+1}, \nu)\frac{\partial T_{t+1}}{\partial g_t} + \underbrace{h(T_{t+1})[V_T(S_{t+1}, T_{t+1}, \tilde{\nu}) - V_T(S_{t+1}, T_{t+1}, \nu)]\frac{\partial T_{t+1}}{\partial g_t}}_{DWI^T} \right.$$

$$\left. \underbrace{+h_T(T_{t+1})\frac{\partial T_{t+1}}{\partial g_t}[V(S_{t+1}, T_{t+1}, \tilde{\nu}) - V(S_{t+1}, T_{t+1}, \nu)]}_{MHE^T} \right) = 0 \tag{6}$$

where $\frac{\partial T_{t+1}}{\partial g_t} = -\lambda_t\theta_t(1-\gamma_t) < 0$. The differences between mitigation and SGE can be seen by comparing these two equations. In Eq. (5), the marginal cost of mitigation, $c_m(m_t, g_t)$, equals the expected reduction in marginal climate damages from one extra unit of carbon in the atmosphere, comprised of three effects. The term $V_S$ is the reduction in future climate costs achieved by reducing the stock of carbon in the atmosphere by one unit. The term $h[\tilde{V}_S - V_S]$ is called the "differential welfare impact", $DWI^S$, and captures the difference in the marginal climate costs associated with changes in the carbon stock incurred if the system crosses a CTP.[5] The term $h_T[\tilde{V} - V]$, is the "marginal hazard effect" of mitigation, $MHE^S$, and captures the marginal reduction in the hazard associated with an increase in mitigation (Lemoine and Traeger, 2014). Both $DWI^S$ and $MHE^S$ appear in the equation because of the presence of a CTP and both work in the same direction, increasing the optimal amount of mitigation.

Eq. (6) states that the marginal cost of SGE equals the expected reduction in marginal damages from a small increase in temperature. The first difference between (6) and (5) is that SGE reduces marginal damages in the same period it is implemented, creating an added benefit that is not discounted. We call this new term the "instantaneous damages reduction" effect, $IDR$. This effect is in turn composed of two terms. $D_g > 0$ represents the marginal damages from SGE, while the negative term that includes $D_T$ captures the benefit of reduced warming from SGE. We assume throughout the analytical model that the benefits of SGE outweigh its damages. Otherwise the optimal policy does not involve SGE.[6] The next term, $h[\tilde{V}_T - V_T]$, which we call $DWI^T$, is the differential welfare impact associated with a change in temperature. The third term, $h_T[\tilde{V} - V]$, is the "marginal hazard effect", $MHE^T$, and captures the marginal reduction in the hazard associated with an increase in SGE. All of these terms, $IDR$, $DWI^T$, and $MHE^T$, work in the same direction, increasing the optimal amount of SGE due to the possibility of a CTP.

There is another subtle difference between the two policies that is not apparent from Eqs. (5) and (6): the MHEs from mitigation and from SGE are quite different from each other. Mitigation, through reductions in carbon concentrations, alters temperatures in the next period and all future periods. SGE, on the other hand, affects future and current temperatures. While both instruments have similar effects on future welfare, SGE does not rely on the inertia of the carbon-climate system and has a direct, instantaneous impact. But its impact is limited in the long run since SGE does not reduce $CO_2$ concentrations. As a result, continuous use of SGE is required to maintain a particular temperature level until concentrations are brought down through mitigation. This important difference between mitigation and SGE will become apparent in our numerical simulations.

### 2.1. Comparing SGE rules

The regulator chooses the optimal levels of mitigation and SGE subject to one of three rules regarding SGE availability. These three rules, that we assume are exogenous to the regulator, encompass different options presented in the solar geoengineering debate:

4 To simplify notation we write $X_Y(Y) \equiv \partial X(Y)/\partial Y$.

5 For the sake of clarity in the text we define $\tilde{V}_S$ as the post-CTP regime value function derivative with respect to $S$: $V_S(S_{t+1}, T_{t+1}, \tilde{\nu})$. Likewise, we define $V_S$, $\tilde{V}$, and $V$.

6 Our assumption produces an interior equilibrium but is a strong assumption. Risks associated with the use of SGE remain largely unknown and uncertain. We abstain here from those uncertainties to concentrate on the risk associated with CTP. This issue is analyzed extensively in Heutel et al. (2015) and Moreno-Cruz and Keith (2013).

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002

(a) **Ban**: SGE is never allowed; $\theta_t = 0$ for all $t$.
(b) **Unconstrained**: SGE is always allowed; $\theta_t = 1$ for all $t$.
(c) **Reparation**: SGE is allowed only after the CTP has been reached at time $t = \bar{t}$; $\theta = 0$ for $t < \bar{t}$ and $\theta = 1$ for $t > \bar{t}$.

We are interested in understanding how the presence of SGE affects the optimal amount of mitigation policies, carbon concentrations, and temperature levels in a world with CTPs. We expect optimal policy to satisfy the following hypotheses.

Before the CTP is crossed, that is, in the pre-CTP regime, we expect:

(i) For mitigation:

$$m^{ban} > m^{reparation} > m^{unconstrained}$$

(ii) For SGE:

$$0 = g^{ban} = g^{reparation} < g^{unconstrained}$$

(iii) For temperature levels:

$$T^{ban} > T^{reparation} > T^{unconstrained}$$

(iv) For atmospheric carbon concentrations:

$$S^{ban} < S^{reparation} < S^{unconstrained}$$

The intuition regarding these hypotheses is as follows. Consider first the *Ban* rule. In this case, SGE is zero at all times, and Eq. (6) does not apply. The $DWI^S$ in Eq. (5) implies that the benefits of mitigation occur in the future, and mitigation reduces damages before and after the CTP is crossed. Mitigation also reduces the propensity to cross the CTP, $MHE^S$, in the immediate future. This implies that the presence of CTPs increases the optimal amount of mitigation, relative to a climate system without CTPs.

Next, consider the *Unconstrained* rule, where SGE can be freely used at any period. This is the case described by Eqs. (5) and (6); both mitigation and SGE are used to tackle climate change. That is, by construction, the *Unconstrained* rule represents the optimal policy, and the outcomes under the other two rules must be sub-optimal.[7] The differential welfare impacts and the marginal hazard effects increase both mitigation and SGE, relative to the case of no CTP. The IDR effect is independent of CTPs, since it is not related to future damages. As a result, any increase in SGE due to the risk of a CTP comes at a cost today, and this cost restrains the amount of SGE. It follows from the discussion above that the planner substitutes away from mitigation and toward SGE, increasing atmospheric carbon concentrations relative to the *Ban* rule.

Finally, consider the *Reparation* rule, where SGE can be used only after the CTP has been crossed. Under this rule, both mitigation and SGE account for the *DWI*, but now the $MHE^T$ cannot be addressed with SGE: once the threshold is crossed, changes in the climate system cannot be reversed even if we substantially reduce temperatures with SGE. Relative to the *Unconstrained* rule, the amount of SGE will be lower, and the amount of mitigation will be higher. This results in higher temperatures and lower carbon concentrations before the threshold is crossed, relative to the *Unconstrained* rule. Relative to the *Ban* rule, mitigation is lower, carbon concentrations are higher, and temperature is lower.

To corroborate our intuition and quantify our analysis, we develop and implement a numerical simulation model that allows us to explore the dynamics of the system in a more comprehensive framework.

## 3. Numerical simulation model

The analysis presented in the preceding section relies on a parsimonious analytical model. In this section, we extend our analysis by modifying an integrated assessment model to incorporate CTPs and the possibility of SGE. The dynamic integrated climate-economy (DICE) model has been widely used to study climate change and optimal climate policy.

[7] This of course follows from the assumption that all costs, damages, and risks of SGE are included in our model. Bans or limits on SGE use are generally recommended due to the fear of unforeseen damages excluded from models. As we discussed above, the *Ban* rule could be optimal if the damages from SGE are much larger than the benefit associated to a reduction in temperatures.

In what follows, we provide a brief qualitative description of the DICE model, and in the following sections we describe our modifications to the standard DICE model to incorporate CTPs and SGE. Appendix A of this paper provides more details on the model, including all of the model's equations and parametrization, and an extensive description of our solution algorithm.[8]

As in the 2007 version of DICE, ours is a finite horizon dynamic model with 60 time periods (600 years). It includes a representative agent model of the economy with exogenous technological growth. In each period (a decade), an existing capital stock is used as an input to an aggregate production function. For simplification we assume an exogenous, fixed savings rate: the representative consumer saves a fixed fraction of net output and consumes the rest.[9]

Carbon emissions are a byproduct of economic production. Carbon accumulates in the atmosphere over time and mixes with the carbon stock of the ocean, according to a dynamic transition model. The radiative forcing – the difference between incoming short-wave radiation and outgoing long-wave energy (heat) – is a function of atmospheric carbon. Global temperature is a function of radiative forcing and past temperatures. This climate model is calibrated based on scientific studies.

The economic and climate models are "integrated" together in that increasing global temperatures reduce net economic output. The wedge between gross and net output is an increasing function of temperature, called the damage function. These damages can be avoided by spending on abatement to reduce emissions, and the cost of abatement is calibrated based on engineering and econometric studies.

The model can be used to calculate optimal climate mitigation policy, which maximizes total discounted net consumption by comparing the costs of abatement with the damages from temperature growth. Optimal policy can be expressed by the optimal amount of mitigation in each period as a percentage of emissions abated, $m_t$, or by the optimal carbon price in each period.

### 3.1. Summary of modifications to DICE

Here we briefly summarize our modifications to DICE. These are based on the modifications in Heutel et al. (2015), and more detail is available there, as well as in this paper's appendix. There are six modifications made to DICE to incorporate SGE and CTPs.

#### 3.1.1. SGE intensity

We include a choice variable for the intensity of SGE, $g_t$, analogous to DICE's choice variable for the intensity of mitigation, $m_t$. Thus, in addition to choosing an optimal mitigation path, our model solves for an optimal SGE path. Both $m_t$ and $g_t$ are proportions. $m_t$ is the proportion of emissions that are abated and is between 0 and 1. $g_t$ is the proportion of radiative forcing that is reduced (see below), and it can take values greater than 1.

#### 3.1.2. SGE's effect on radiative forcing

SGE affects the radiative forcing of the Earth's atmosphere, reducing the amount of sunlight entering and thereby reducing temperature. DICE has a dynamic model of temperature based on radiative forcing, and radiative forcing itself is determined by carbon concentrations. SGE reduces radiative forcing directly, therefore quickly reducing temperatures. Setting SGE to $g_t = 1$ corresponds to reducing radiative forcing to its pre-industrial levels. By allowing $g_t > 1$, SGE can reduce forcing even below preindustrial levels, even more quickly reducing temperatures.

#### 3.1.3. SGE implementation cost

Our specification of implementation costs is analogous to DICE's specification of the cost of mitigation. It is a convex (quadratic) function of the intensity of SGE $g_t$. It is calibrated from back-of-the-envelope calculations based on Crutzen (2006), Rasch et al. (2008), and McClellan et al. (2012). Costs are expressed as a fraction of gross output.[10] Implementation costs are small; in our calibration the costs of SGE at intensity $g = 0.1$ is 0.06% of gross output. Instead, the larger risks of SGE come from its potential damages.

#### 3.1.4. SGE damages

SGE may directly cause damages, for instance, by reducing the upper ozone layer (Heckendorn et al., 2009). We model these damages analogously to DICE's specification of damages from climate change, as a proportional loss of potential output. We know of no study that attempts to quantify these damages, and thus this parameterization is inherently uncertain. We attempt to be conservative (i.e., biased against SGE) in our parameterization and assume that full SGE ($g_t = 1$) causes damages

[8] See also Nordhaus (2008) for a summary of the model's assumptions and equations.

[9] In practice, when savings is allowed to be endogenous the savings rate only varies slightly from this fixed value.

[10] The SGE variable is the fraction by which the total radiative forcing is reduced. The magnitude depends on the level of radiative forcing, which itself depends on the level of carbon concentration and the size of the economy. Therefore, the cost of geoengineering is expressed as a fraction of total economic output. While there are alternative modeling strategies (e.g. costs proportional to the amount of geoengineering), we prefer to model SGE costs similarly to the way abatement costs are modeled to minimize the channels that can explain differences in our results. One consequence of this assumption is that climate damages act to make abatement and solar geoengineering cheaper (in terms of consumption).

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002

equal to 3% of gross output. This is on par with the magnitude of climate change damages in DICE from a 6 degrees Celsius temperature increase.

### 3.1.5. Climate change damages directly from carbon

In DICE, climate change damages are a function of global temperature only. Since SGE reduces temperatures but not atmospheric or ocean carbon concentrations, in our model damages from climate change are separated out between damages from temperature, from atmospheric carbon concentrations, and from ocean carbon concentrations. High ocean carbon concentrations result in ocean acidification, which can lead to damages (Brander et al., 2012). High atmospheric carbon concentrations may yield benefits (Pongratz et al., 2012) or damages (Bony et al., 2013). Just like with damages from SGE, these damages are mostly unknown, and therefore this calibration must be rather arbitrary. We keep the total level of climate change damages identical to the calibrated level in DICE. We assume that the majority (80%) of climate change damages come directly from temperature, but a small amount of damages may come from ocean concentrations (10%) and from atmospheric concentrations (10%). As shown in Heutel et al. (2015), this implies that SGE is not a perfect substitute for mitigation.[11]

### 3.1.6. Climate tipping points

The incorporation of climate tipping points into DICE along with SGE is unique to this paper and not found in Heutel et al. (2015). CTPs are modeled as irreversible events (in dynamic programming language, absorbing states), meaning that once we hit a tipping point we enter a new state in terms of climate or economic systems, and there is no possibility of returning to the old state. We consider three types of CTPs: two affecting climate dynamics and one affecting economic output. The first two CTPs are analogous to the two CTPs modeled in Lemoine and Traeger (2014); the third is analogous to the CTP modeled in Cai et al. (2013) and Lontzek et al. (2015).

(i) **Climate feedback**: Crossing this CTP strengthens the temperature feedback loop by increasing the marginal effect of carbon on temperature. Numerically, after this CTP is crossed the climate sensitivity variable increases from 3 °C to 5 °C.[12]
(ii) **Carbon sink**: Crossing this CTP reduces the natural capacity of the planet to absorb carbon. Numerically, after crossing this CTP, carbon sinks are weakened by 50%.[13]
(iii) **Economic loss**: Crossing this CTP causes a permanent loss in gross economic output equal to 1% of net output.

As in Lemoine and Traeger (2014), the probability of reaching a CTP in the next period is a function of the atmospheric temperature in the current period. A CTP is reached once we cross a threshold temperature, but that threshold temperature is uncertain before it is crossed. The CTP threshold temperature takes a uniform distribution. The minimum value is the current temperature (since once the current temperature has been reached, we know the CTP threshold cannot be below it). The maximum value of the threshold temperature distribution is calibrated so that the expected value of the threshold temperature is 2.5 $^{\circ}C$ in 2005. Therefore, in each period, the probability distribution of the threshold temperature in the next period is uniform between $T_t$, the current temperature, and $\overline{T}$, the upper limit temperature:

Tipping points are introduced in the DICE model as a binary variable. The value of this variable is set to zero before crossing the tipping point; once a tipping point is crossed the variable changes to one and stays at one for the rest of the simulation. Depending on the type of CTP, subsequent state variables (including temperature, carbon concentration, and economic output) are calculated.

## 3.2. Solution algorithm

Here we briefly describe our solution algorithm, which is based on the methodology introduced in Shayegh and Thomas (2015). The appendix provides more details and conducts several robustness and validity tests.

The evolution of the climate-economy system under uncertain tipping points is modeled as a Markov decision process. We define $S_t$ as a state variable with multiple dimensions. For this problem, the state variable has eight dimensions: capital, atmospheric temperature, lower ocean temperature, atmospheric carbon concentration, upper ocean carbon concentration, lower ocean carbon concentration, radiative forcing, and a binary state variable capturing whether or not the CTP has been crossed. Given the values of the state variable parameters at each time step, the mitigation action, the SGE action, and the realization of uncertainty (crossing the tipping point), we can calculate the state variable parameters for the next time step. The model is solved assuming a finite time horizon of 60 periods, where each period represents 10 years.

[11] In Heutel et al. (2015), we explore this calibration using sensitivity analysis. The qualitative behavior of the system remains the same so long as the temperature damages dominate the outcomes.

[12] Climate sensitivity measures the steady-state temperature increase due to doubling atmospheric carbon levels. In our analytical model, this CTP amounts to an increase in $\lambda_t$ in Eq. (2). See the appendix for details.

[13] In our analytical model, this is a decrease in $\delta_t$ in Eq. (1). See the appendix for details.

The state variables and action variables are continuous (except for the binary CTP indicator), so finding an exact solution for this problem through conventional backward induction methods is infeasible. Therefore, to solve the model, we use the two-step-ahead approximation method introduced in Shayegh and Thomas (2015). The approximation technique was tested and tuned in the deterministic case and then applied to the stochastic model.[14] In this technique, at each time step $t$, a value function $\overline{V}_t$ is defined and used to approximate the future utility from taking a candidate action $a_t$:

$$\hat{V}_t(S_t) = \max_{a_t}\{U_t(S_t, a_t) + \overline{V}_t(S_t)\} \tag{7}$$

where $\hat{V}_t(S_t)$ is the optimal value of state $S_t$ based on the value approximation.

We define $\overline{V}_t$ as a linear combination of utilities of two subsequent states in the future. We construct these two states deterministically by assuming that uncertainty in the model remains at its current level observed in time $t$ for the next two time steps $t+1$ and $t+2$. The first look-ahead state $S_{t+1}$ is constructed given the current state ($S_t$) and action ($a_t$). After that, we apply a predefined action $a_0$ twice to construct the next look-ahead state $S_{t+2}$ and calculate the utilities of these next two states. The utilities of the two states $S_{t+1}$ and $S_{t+2}$ under the assumptions about the value of the uncertain parameter and the predefined action $a_0$, are used to calculate the approximate value function $\overline{V}_t$. The predefined action $a_0$ is the myopic optimizer action that maximizes the immediate utility ($U_t$) at each time step. Since $U_t$ is a smooth and concave function, the existence of this action is guaranteed. The linear coefficients of the approximate value function $\overline{V}_t$ are updated after each iteration. More information and discussion about the details of this algorithm are provided in the appendix.

The advantage of this technique is in using endogenous parameters to calculate the value function approximation by assuming a deterministic trajectory for the two steps into the future at any given time. The deterministic trajectory allows us to calculate the utilities of these two future steps and bring them back to the present time using an artificial and tunable discount rate. The adjusted value is then used as a proxy for the uncertain value of all future states. These values reflect the social utility under the deterministic assumption and are used to construct the value function of the current state. The optimal action (mitigation and SGE) is found by maximizing this value function. The algorithm starts at time $t = 1$ and progresses until the last time step. After calculating all value functions, these values are used to update the coefficients of the approximate value function in previous states. The algorithm then iterates until the error (the difference between approximate values of $\overline{V}_t$ and optimal values of $\hat{V}_{t+1}$) converges to zero.

The algorithm is developed in MATLAB and is available upon request. The full description of the model and approximation algorithm is presented in the appendix.

## 4. Simulation results and discussion

In this section we discuss the simulation model results for the three types of tipping points under the three different rules regarding the availability of SGE. We analyze the optimal climate policy portfolio of mitigation and SGE, and the resulting temperature, carbon concentrations, carbon price, and welfare. We run simulations for two alternative cases that allow us to highlight the role of SGE to deal with CTPs. First, we consider the case where the tipping point is never reached, but the possibility of a tipping point affects the incentives to mitigate and to implement SGE in a pre-CTP regime. Second, we consider the case where the tipping point is crossed in an arbitrary year, 2085, and show how SGE affects both the pre- and post-CTP regime.

### 4.1. Case 1: pre-CTP regime policy

The results presented in this case are conditional on not having crossed the CTP during the first 120 years of the simulation.[15] The objective is to observe how the option of SGE affects the optimal pre-CTP policy.

In Fig. 1 we present the optimal mitigation (row 1), and SGE (row 2). In Fig. 2 we present atmospheric temperature in degrees Celsius above preindustrial level (row 1), and atmospheric carbon concentration in GtC (row 2). In all figures, we present optimal policy and outcomes under the three different CTPs: Climate feedback (column 1), Carbon sink (column 2), and Economic loss (column 3).[16] In each panel, the horizontal axis shows the year of the simulation.

We also simulate a baseline economy that does not include CTPs or the possibility of a CTP; we call these the no-CTP simulations.[17] Under this baseline we consider both the *Ban* rule and the *Unconstrained* rule regarding SGE availability (the *Reparation* rule does not apply in the no-CTP case). The no-CTP baseline allows us to see the effect of CTPs themselves, in addition to the effect of SGE rules. In the no-CTP, Ban SGE scenario (the *orange-dashed* line in Fig. 1), mitigation levels increase over time from around 15% in 2015 to about 54% by 2125. We can see in Fig. 2 that the associated temperature peaks at

[14] To test the accuracy of this solution algorithm, we use it to replicate the results in Lemoine and Traeger (2014). This exercise is described in the appendix. Our approach does not require a reduction in the dimension of the state space, and we are able to solve the problem using the full set of transition equations used in the original DICE model.

[15] Although our simulations are run for 60 periods (600 years), as is standard in the literature in these figures we present only the outcomes for the initial periods, in our case the first 12 periods.

[16] The appendix also presents the carbon price and net welfare outcomes.

[17] This is distinct from the pre-CTP regime policy of Case 1, in which CTPs are possible but never crossed.

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002

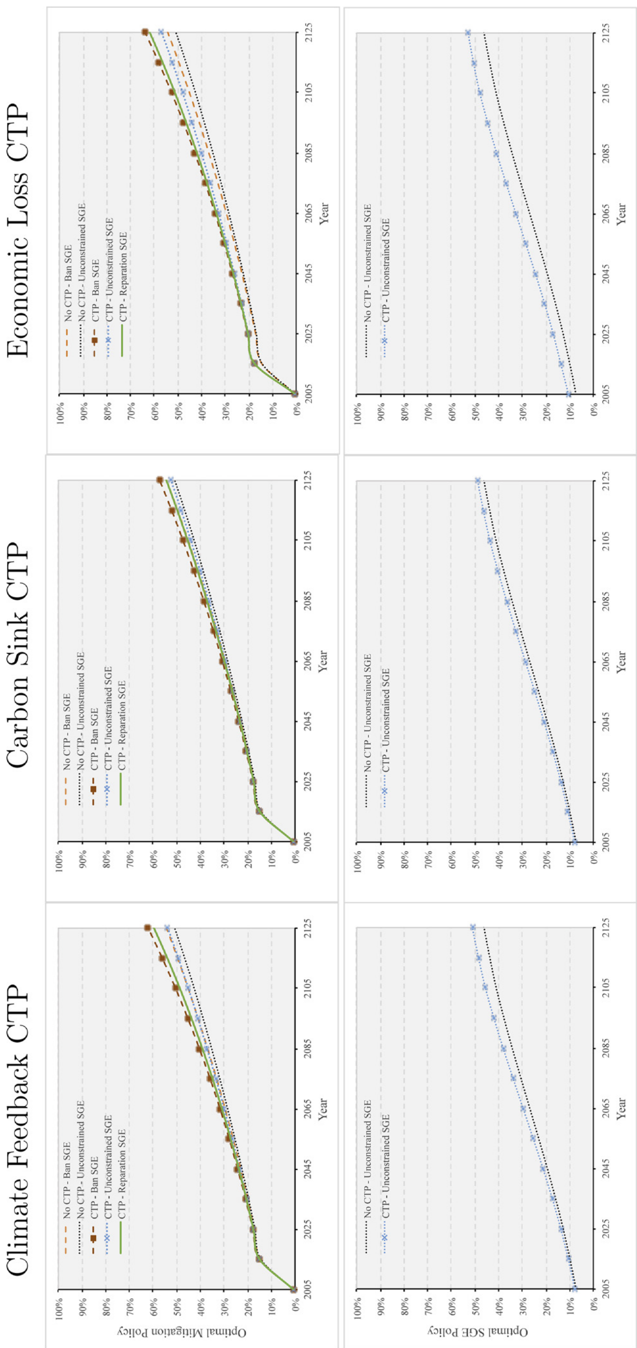


**Fig. 1.** Pre-CTP optimal policy. (For interpretation of the references to color in the text, the reader is referred to the web version of the article.)

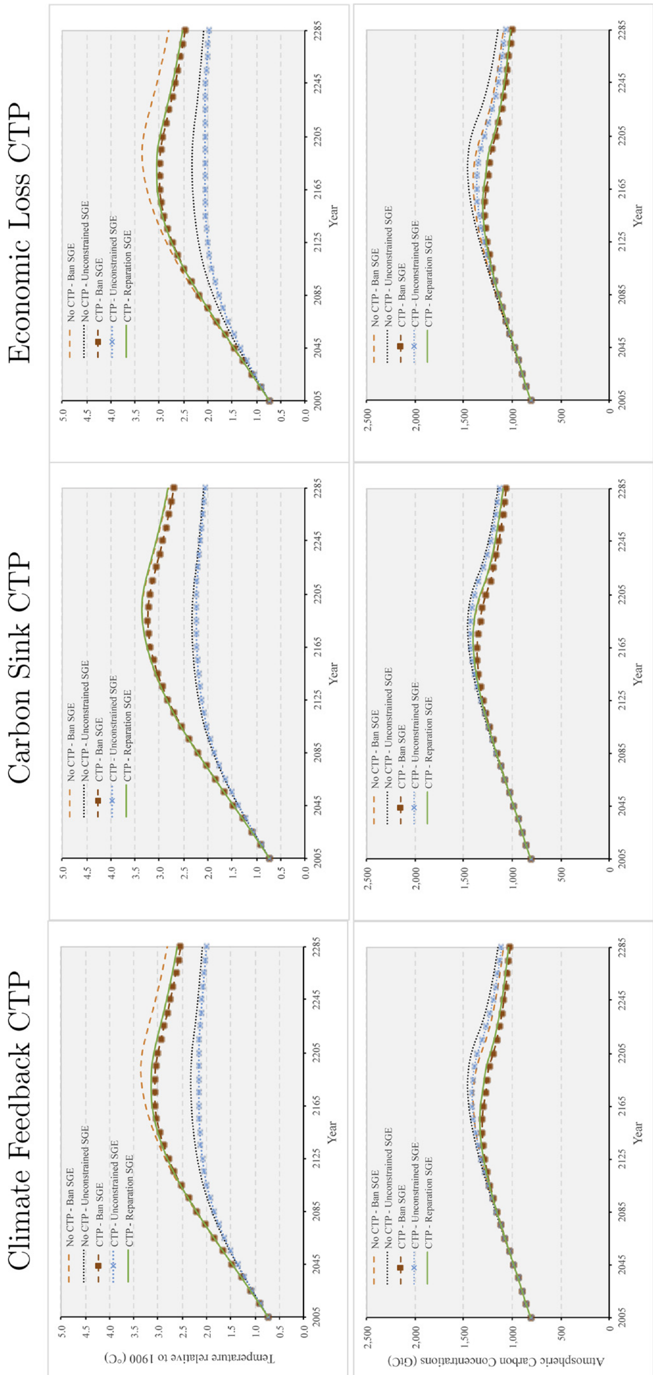


**Fig. 2.** Pre-CTP climate outcomes.

3.35 °C in 2195 and by the end of our simulation period in year 2285 it is down to 2.8 °C. Carbon concentrations peak at 1401 GtC in the year 2175 and come down to 1091 GtC in the year 2285.[18]

The no-CTP, Unconstrained SGE scenario is shown by the *black-dotted* line in Fig. 1. Mitigation levels start at around 15% in 2015 and reach 50% by 2125. Allowing SGE reduces the amount of mitigation by around 4% in 2125. This is equivalent to a reduction in the carbon price of 33$/tC, from 301$/tC to 268$/tC. SGE starts at around 10%; by 2125 it has reached 46%, reducing radiative forcing by about half. The resulting temperature peaks at 2.32 °C in 2175, one degree Celsius lower and 2 decades earlier than the case without SGE. Carbon concentrations peak in the year 2185 at 1459 GtC. SGE increases the peak by 58 GtC and delays it one more decade, relative to the baseline without SGE.[19]

Next, consider the scenarios that include the possibility of a CTP (though here in Case 1 the CTP is never actually crossed). The CTP, Ban SGE scenario is shown by the *brown-blocked-dashed* line in Fig. 1. For the Climate Feedback CTP (the first column in Fig. 1), the introduction of a CTP causes optimal mitigation to increase by about 10%, reaching 62% by 2125. Temperature peaks in 2175 at 3 °C, still quite high but 10% lower than in the no CTP, Ban SGE scenario. Carbon concentrations peak in 2145 at a value of 1308 GtC.

The CTP, Unconstrained SGE scenario is shown by the *blue-starred* line. The possibility of a CTP causes SGE to increase from 46% to 51%, relative to no-CTP, Unconstrained SGE. While this is small, it has a substantial impact on temperature, which now stays below 2.15 °C. The presence of SGE is associated with a substantial reduction in mitigation relative to the CTP, Ban SGE scenario, going from 62% down to 54%. Carbon concentrations increase relative to the CTP, Ban SGE scenario, reaching 1413 GtC, an increase of 100 GtC, 3 decades later.

Under the CTP, Reparation SGE scenario, shown by the *green-solid* line, there is some reduction of the amount of mitigation relative to CTP, Ban SGE, but it is not substantial. Both carbon concentrations and temperatures are slightly larger under the CTP, Reparation SGE scenario, relative to the CTP, Ban SGE scenario. Under both the *Ban* and *Reparation* rules, there is never any SGE used, since the CTP is not crossed. But, there is slightly more mitigation used under the *Ban* rule, since the planner knows that SGE will never be available.

Similar behavior is observed for the other two tipping points, in the second and third columns in Fig. 1. The qualitative effect of CTP and SGE rules on mitigation, temperature, and carbon levels is identical for all three CTPs. Relative to the *Ban* rule, the *Unconstrained* rule yields lower mitigation, higher carbon, and much lower temperature. The *Reparation* rule fits somewhere in between the *Ban* rule and the *Unconstrained* rule, but it is very close to the behavior of the system under the *Ban* rule.

This does not imply that optimal policy is unaffected by the type of tipping point. To see this, we present in Fig. 3 how the different effects identified in the theory section contribute to the optimal amount of mitigation and SGE under the different CTP types. In Eqs. (5) and (6) we identified two CTP-induced effects, *DWI* and *MHE*, that increase both mitigation and SGE. There is also a contemporaneous *IDR* effect that only affects, and decreases, SGE. Fig. 3 presents this decomposition of optimal mitigation and SGE policies, for optimal policy in 2050. The blue-colored section shows the optimal policy in the no-CTP scenarios. The effect of changing the pre-CTP value function is shown in a light shade of orange. This is the largest contribution for mitigation policy and SGE policy across all CTPs, except for SGE policy under the Economic loss CTP, where the marginal hazard effect dominates. The Economic loss CTP represents the highest risk, not because it causes larger overall damages, but because it cannot be dealt with after the threshold is crossed. Hence, more effort is exerted to avoid a CTP altogether. The *IDR* effect appears negative because its positive aspects are already included in the policy without CTP. That is, the reduction in the contemporaneous temperature damages is not a function of CTPs, but the extra amount of SGE used to deal with the CTP increases contemporaneous SGE damages.

### 4.2. Case 2: post-CTP regime policy

In Fig. 4 we analyze policy in the case where the CTP threshold is crossed in 2085.[20] We observe how the option of SGE affects the optimal post-CTP policy. The panels are organized in the same way as those in the previous section. Because the CTP is reached in these simulations, there are no no-CTP scenarios. Unlike in the pre-CTP regime, here there are important qualitative differences across CTPs.

Optimal policy for the *Ban* rule is shown in the *brown-square-dashed* line in all panels. For the Climate feedback CTP, the mitigation trajectory becomes slightly steeper after 2085, which in turn results in a faster decline in concentrations. Temperature, however, increases to a peak of 4.4 °C above preindustrial in 2185, one full century after the CTP is crossed. It eventually comes down to 3.92 °C. Recall that if the CTP is not reached, the maximum temperature increase under the *Ban* rule was 3 °C.

[18] The first two rows present the simulation results just through 2125, while the bottom two rows present them through 2285.

[19] All of the no-CTP outcomes are identical across the three columns, since the three columns differ only by type of CTP, which is irrelevant in the no-CTP scenarios.

[20] In order for each scenario to reach the CTP in the same year (2085), the actual realized CTP threshold temperature differs across scenarios. The third row of Fig. 4 indicates the realized CTP threshold across scenarios.

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002

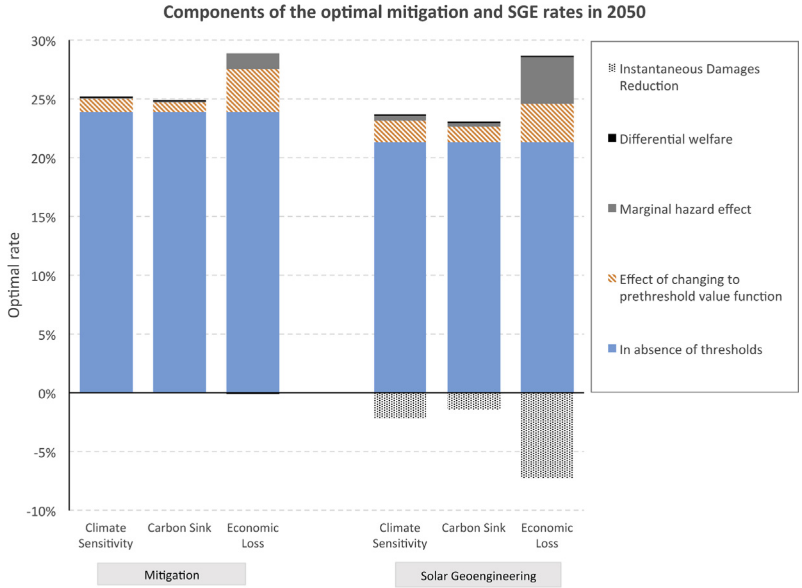


**Fig. 3.** Pre-CTP contribution of different effects. In each case we have numerically simulated the model under deterministic and stochastic optimal policies and calculated the difference in the probability of tipping points and the corresponding utilities when switching from the deterministic case to the stochastic case. These utilities are then translated into corresponding optimal control rates. (For interpretation of the references to color in the text, the reader is referred to the web version of the article.)

For the *Unconstrained* rule, shown by the *blue-starred* line, we see an increase in SGE and a reduction in mitigation immediately after the CTP is crossed, caused by the re-optimization of the program. The increase in carbon concentrations is not substantial. Temperatures, on the other hand, are kept below 2.5 °C throughout the planning horizon.[21]

Under the *Reparation* rule, shown by the *green-solid* line, SGE is deployed in higher rates than it is deployed under the *Unconstrained* rule, to compensate for the extra warming committed by the time the CTP is reached. Temperatures are nonetheless kept at slightly above 2.5 °C before they come down and asymptotically approach the unconstrained behavior.

The second column in Fig. 4 simulates the Carbon sink CTP. Under the *Ban* rule, carbon concentrations increase to 1977 GtC, while temperature reaches 4.6 °C above preindustrial and stays at that level for several decades. Mitigation picks up after the threshold is reached, but is slow and its effect limited. When SGE is used under the *Unconstrained* rule, temperatures are kept below 2.5 °C. Concentrations increase to 2087 GtC before starting to decline. Under the *Reparation* rule, SGE jumps to react to the crossing of the CTP, overshooting as before the amount of SGE under the *Unconstrained* rule. Temperatures and concentrations are almost exactly those under the *Unconstrained* rule.

The results in the last column of Fig. 4, for the Economic loss CTP, are much different than those for the other two CTPs. Once the CTP is crossed, a substantial Economic loss is felt, reducing the incentives to protect the economy and thus reducing both mitigation and SGE. Nonetheless, under the *Unconstrained* and the *Reparation* rules temperatures are kept below 2.5 °C and carbon concentrations are not much higher than if the CTP was not crossed (Fig. 5).

### 4.3. Summary

These simulations demonstrate that SGE can be used as a substitute, albeit an imperfect substitute, for mitigation in managing the risks of CTPs. Without the availability of SGE (the *Ban* rule), the presence of CTPs causes more mitigation to be used. Depending on the type of CTP, temperatures and carbon stocks may be higher or lower with the CTP than without it.

Under the optimal policy portfolio (the *Unconstrained* rule), the risk of a CTP increases the use of SGE but does not substantially affect mitigation. Thus, nearly all of the risk of CTPs is managed by SGE rather than by mitigation.

When SGE is restricted to only be allowed after the CTP is reached (the *Reparation* rule), mitigation is used much more intensively before the CTP is crossed, since it is the only policy option that can manage that risk. Once the CTP is crossed and SGE is allowed, SGE is used less intensively than under the *Unconstrained* rule, since there is no benefit in terms of reduced probability of CTP risk (no marginal hazard effect).[22]

[21] Recall the expected value of the temperature threshold location is 2.5 °C.
[22] In the appendix, we discuss how the different scenarios affect the carbon price and social welfare.

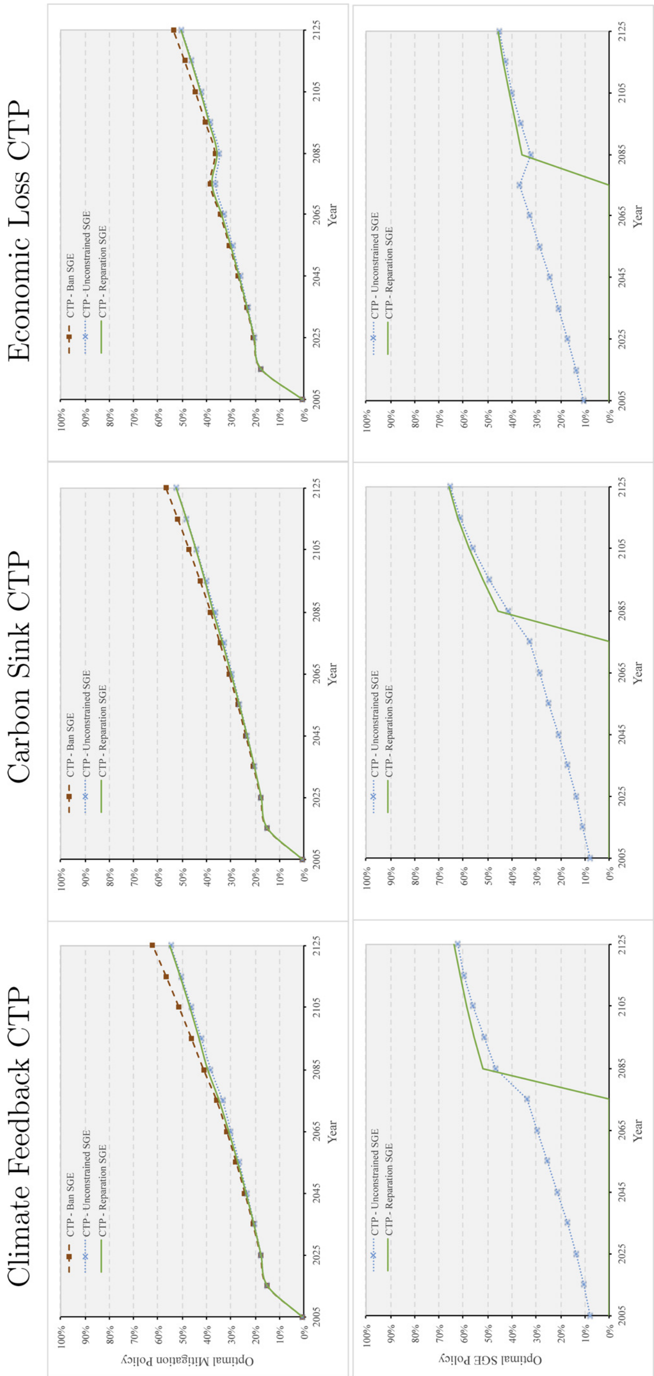


**Fig. 4.** Post-CTP optimal policy.

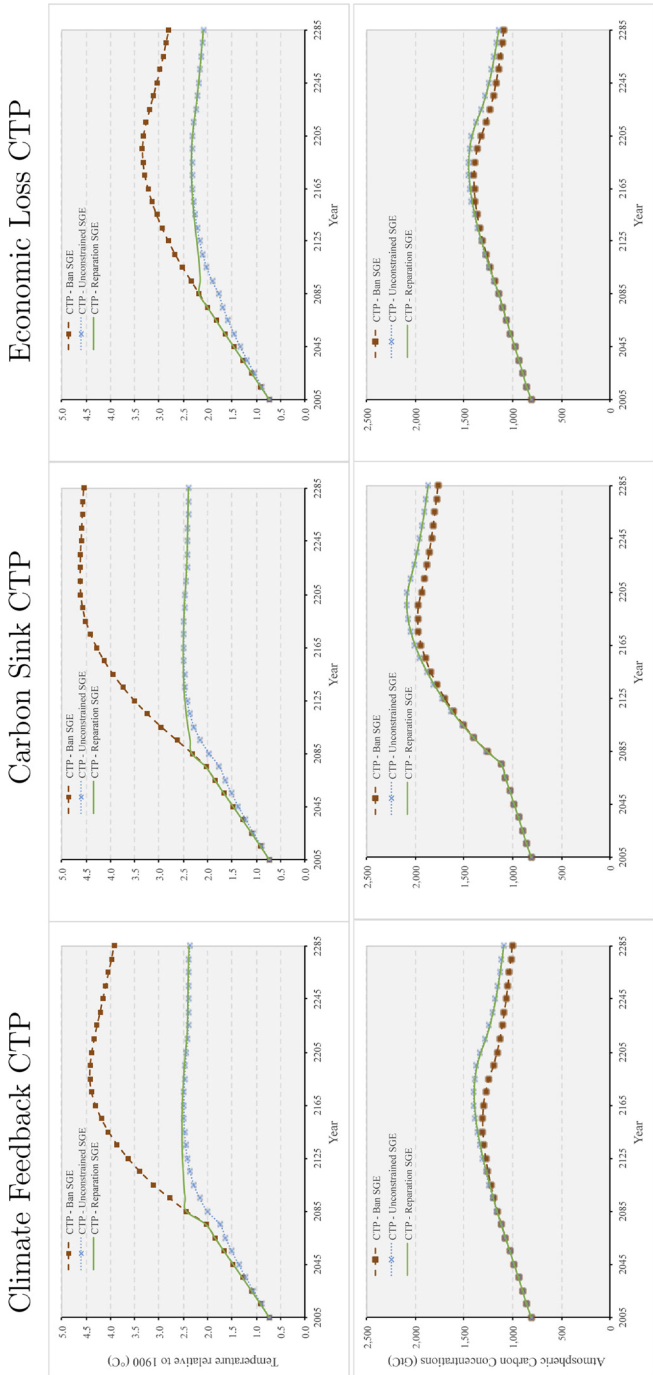


**Fig. 5.** Post-CTP climate outcomes.

## 5. Conclusion

We consider optimal climate policy when solar geoengineering is included as a policy option and tipping points are potential threats. Solar geoengineering is part of the optimal policy portfolio for two reasons. First, it provides a means to control temperature at (potentially) a lower cost than mitigation. Second, it can be used to reduce the risk of reaching a climate tipping point. Thus, refraining from using SGE only until a tipping point has been reached (our *Reparation* rule) is not a welfare-maximizing policy. The relatively fast nature of the tipping points (the parameters jump immediately) is part of what makes SGE such an attractive option for managing their risk – since SGE can lower temperatures much more quickly than can mitigation.

Our analytical results were reached using a simple model; we have done so to concentrate on the importance of SGE in dealing with mega-disasters caused by CTPs. Our numerical approach modifies the DICE model to incorporate SGE, three rules governing its use, and three types of tipping points. The simulation results confirm our predictions from the analytical model, but they also provide us with a quantitative characterization of alternative policy scenarios. As with any integrated assessment model, results depend on the parametrization and calibration of the model, much of which is highly speculative.[23]

In particular, in the case of solar geoengineering there are a number of potential caveats that we do not address but that could be incorporated into future work. Geoengineering may impose heterogeneous costs or benefits.[24] There are asymmetric uncertainties associated with SGE that are not modeled here.

We find that tipping points call for more action, but this action can take the form of a combination of mitigation and SGE, rather than mitigation alone. This allows for a climate policy with less mitigation and lower temperatures, relative to a world without SGE. At the levels suggested by our simulations, SGE does not eliminate the risk from CTPs altogether but substantially reduces it.

## Appendix A. Details on the model and solution algorithm

In the paper, we briefly summarize the DICE model, our modifications to it, and our solution algorithm. Here, we provide more details.

### *A.1. The stochastic DICE model with tipping points*

We modify the DICE model, first introduced by Nordhaus (1993). The model parameters and equations are from Nordhaus (2008). We have modified the DICE 2007 version of the model in order to include a probability of the tipping points and the SGE action. This is a finite horizon model with 60 time steps. Each time step is a decade, and the starting year is 2005. We model the stochastic DICE as a Markov decision process with a state space, an action space, an information space, a transition function, and a reward function.

- **State space**
  The global climate-economy system can be defined as a state with seven continuous variables: $T_t^{at}$ is atmospheric temperature (degrees Celsius above preindustrial), $T_t^{lo}$ is lower ocean temperature (degrees Celsius above preindustrial), $M_t^{at}$ is the atmospheric concentration of carbon (Giga Tons of Carbon, GTC), $M_t^{up}$ is the concentration in the biosphere and upper oceans (GTC), $M_t^{lo}$ is the concentration in deep oceans (GTC), $K_t$ is capital ($trillions), and $F_t$ is radiative forcing (W/m$^2$). In addition, there is an eighth, binary variable, $\nu_t$, representing whether or not the CTP has been reached. We define the state space as $S_t = \{T_t^{at}, T_t^{lo}, M_t^{at}, M_t^{up}, M_t^{lo}, K_t, F_t, \nu_t\}$.
- **Action space**
  At each time step, a mitigation action (control rate) $a_t$ and a SGE action $g_t$ are taken, which indicate the percentage reduction of GHG emissions and the percentage reduction of radiative forcing, respectively. Both actions impose immediate costs but prevent the future damages of higher temperature. Taking actions $a_t$ and $g_t$ at any given state will determine the next state deterministically. Therefore the action space is defined as $a_t \in [0, 1]$ and $g_t \geq 0$. As in the original DICE model, savings is fixed as a fraction of gross output and thus is not a choice variable.
- **Information space**
  The only parameter over which there is uncertainty is the temperature threshold defining the CTP, before the CTP is reached. It is a distributed uniformly between the current temperature and the upper limit temperature $\overline{T}$.[25] The probability of crossing the tipping point in the next time step is calculated as

$$h_{t+1} = \max\left\{0, \frac{\min(T_{t+1}^{at}, \overline{T})}{\overline{T} - T_{t+1}^{at}}\right\} \tag{A.1}$$

  We use $\overline{T} = 4.27^{\circ}$ as the upper bound for the tipping point threshold.

[23] See Pindyck (2013) for a critique of integrated assessment models.
[24] See Robock et al. (2008).
[25] Alternatively, we could introduce uncertainty into this system via atmospheric temperature shocks, climate sensitivity, or SGE damages (Heutel et al., 2015).

- **Transition functions**

The initial value of the state variable $\nu_t$ is zero (pre tipping point). If the tipping point is reached, this value changes to one, otherwise it stays at zero. The gross economic output, $Y_t$, is calculated from the given level of technology, capital, and labor in the current state:

$$Y_t = \Gamma_t \times K_t^{\beta} \times L_t^{1-\beta} \tag{A.2}$$

where $\Gamma_t$ is technology and $L_t$ is labor at time $t$. $\beta$ is the output elasticity of capital. The net output, $Q_t$, is calculated after subtracting climate change and SGE damages, and mitigation and SGE costs from gross output[26]:

$$Q_t = (\Delta_t - A(a_t) - G(g_t)) \times Y_t \tag{A.3}$$

$$\Delta_t = \frac{(1-\mathbf{u}_3)(1+\nu_t^g g_t^2)^{-1}}{1+\xi_1(W_t \times (T_t^{at})^2) + \xi_2(M_t^{at} - M_0^{at})^2 + \xi_3(M_t^{up} - M_0^{up})^2} \tag{A.4}$$

$$A(a_t) = \theta_1 \times a_t^{\theta_2} \tag{A.5}$$

$$G(g_t) = \theta_1^g \times g_t^{\theta_2^g} \tag{A.6}$$

where $\Delta_t$ includes damages from both climate change and from SGE. The denominator in the expression for $\Delta_t$ represents climate change damages; these depend on the atmospheric temperature, atmospheric carbon concentration, and upper ocean carbon concentration. The numerator in the expression for $\Delta_t$ contains the damages from SGE: $(1+\nu_t^g g_t^2)$. The term $\mathbf{u}_3$ represents the economic tipping point: if the economic tipping point has not passed yet, $\mathbf{u}_3 = 0$, and if the tipping point is passed $\mathbf{u}_3 = 10\%$. The parameters $\xi_1, \xi_2$, and $\xi_3$ are the damage cost coefficients and are adjusted to replicate the damage cost of the original DICE model for the year 2005. The parameters $\theta_1$ and $\theta_2$ are the coefficients of the mitigation cost function $A(a_t)$ and $\theta_1^g$ and $\theta_2^g$ are the coefficients of SGE cost function $G(g_t)$.

Part of the net output at each time step is saved and invested and the rest is consumed:

$$K_{t+1} = (1-\delta) \times K_t + \theta_3 \times Q_t \tag{A.7}$$

where $\delta$ is the capital depreciation rate and $\theta_3$ is the saving rate. The industrial emissions $E_t$ are found from the carbon intensity of output $\sigma_t$, taking into account the abatement decision:

$$E_t = \sigma_t \times (1-a_t) \times Y_t \tag{A.8}$$

The atmospheric and ocean carbon state variables in the next time period are:

$$M_{t+1}^{at} = E_t + (1-\mathbf{u}_2) \times M_t^{at} + \phi_{21} \times M_t^{up} \tag{A.9}$$

$$M_{t+1}^{up} = \mathbf{u}_2 \times M_t^{at} + \phi_{22} \times M_t^{up} + \phi_{32} \times M_t^{lo} \tag{A.10}$$

$$M_{t+1}^{lo} = \phi_{23} \times M_t^{up} + \phi_{33} \times M_t^{lo} \tag{A.11}$$

where $\phi_{21}, \ldots, \phi_{33}$ are carbon cycle transition coefficients. The parameter $\mathbf{u}_2$ indicates the carbon sink tipping point. When the tipping point is crossed it drops to half of its initial value.

The temperature equations for the next state are:

$$T_{t+1}^{at} = T_t^{at} + \eta_1 \times \{F_{t+1} - \eta_2 T_t^{at} - \eta_3 \times \{T_t^{at} - T_t^{lo}\}\} \tag{A.12}$$

$$T_{t+1}^{lo} = T_t^{lo} + \eta_4 \times \{T_t^{at} - T_t^{lo}\} \tag{A.13}$$

$$F_{t+1} = \eta_2 \mathbf{u}_1 (\log_2(M_t^{at}/M_0^{at})) \times (1-g_t) \tag{A.14}$$

where $\eta_1, \ldots, \eta_4$ are temperature coefficients and $\mathbf{u}_1$ is the climate sensitivity tipping point indicator. If the tipping point is crossed, $\mathbf{u}_1$ will go up from 3 °C to 5 °C.

- **Reward function**

The reward is calculated as the social utility of consumption at each time epoch:

$$U_t = \frac{\{(1-\theta_3) \times Q_t\}^{1-\alpha}}{1-\alpha} \tag{A.15}$$

where $\alpha$ is the elasticity of marginal utility of consumption. The objective is to maximize the sum of discounted expected social utilities over the modeling horizon given uncertainty in climate sensitivity:

$$\max_{a_t \in A(S_t)} \mathbb{E} \left\{ \sum_{t=0}^{T} \gamma^t U_t(S_t, a_t, W_t) \right\} \tag{A.16}$$

[26] In this model, abatement costs (and SGE costs) are expressed as a fraction of gross output, while in DICE abatement costs are a fraction of net output.

### A.2. Solution method

Here we provide more details on the solution algorithm, which is developed in Shayegh and Thomas (2015).

Approximate dynamic programming is a fast growing subject in the field of stochastic optimization. There is a considerably large set of theories and approximation techniques, and the performance of each method depends on the application. Several classes of approximation approaches are used in sequential decision-making problems. Look-ahead policies (used in rolling horizon techniques) and value function approximation (used in value iteration techniques) are among the most popular approaches with a wide range of techniques and algorithms associated with them.

In look-ahead policies, the optimal action is found by solving a set of smaller problems over a limited horizon. When the state and action spaces are large, roll-out heuristics are deployed that expand the decision horizon for a few steps into the future. The value of this expanded horizon is then used to find the optimal action at the current time. The advantage of using these heuristics is that they are usually very fast and simple and provide a better alternative than a pure myopic policy (Truong, 2014). The shortcomings of these methods are their inability to update their forecast of the future and their inflexibility in adjusting to different realizations of the uncertainty in the model. More formally, the look-ahead algorithms utilize what is called "online" learning. They learn about the value of the current state in real time and on the go (Goodson et al., 2015).

On the other hand, value function approximation techniques are more flexible and learn "offline". There are many techniques from parametric regression to nonparametric methods to approximate a value function. Using polynomial functions to approximate the value functions falls under the parametric approximation of value function. By definition, the value function approximation is used to estimate the cost of all future steps that are combined in a value function of the next state. Therefore, this method can be viewed in fact as a one-step look-ahead approximation (Bertsekas et al., 1995). The main difference is that there is a set of tunable parameters or weights that are being updated through an approximate value iteration algorithm.

In this paper, we use a novel method to combine these two approaches. We introduce a look-ahead scheme as an alternative to the parametric structure of a value function approximation, but we keep the approximate value iteration element. This way, we keep the advantages of the online look-ahead method (it is fast and simple), but we add flexibility to our approximation and ensure convergence.

Note that our use of the two-step-ahead utility as the basis for the value function approximation is non-conventional. The common way of approximating a smooth value function is to represent it as a linear combination of polynomials which are functions of the state variables. By contrast, our basis functions are not directly functions of the state variables. Instead, our basis function is proportional to expected utility two periods in the future, which is just indirectly a function of state variables.[27]

To demonstrate the algorithm, consider a simple transition between two states $S_t$ to $S_{t+1}$ as shown in Fig. A.1.[28] In the top panel of the figure, $S_t$ is a state where the tipping point has not yet been reached. The uncertainty about reaching the tipping point in the next state is shown as $h$. After observing the status of the current state (i.e. if the tipping point has reached or not) and taking the action $a_t$ at state $S_t$, we will be able to calculate the next state $S_{t+1}$. In order to find the optimal action $a^*$ we deploy our two-step-ahead approximation algorithm. First, the value of the current state $S_t$ will be calculated by taking a candidate action $a_t$. We build a decision tree based on the probability of reaching the tipping point in the next state and will calculate the expected value of the future states by assigning a set of fixed mitigation and SGE actions to them.[29] There are two possible next states (post-decision states) from taking the action $a_t$ in state $S_t$: either the tipping point is reached ($S_t^{a1}$) and therefore we stay in this state for the consecutive time step ($S_{t+1}^{a11}$), or the tipping point is not reached ($S_t^{a0}$) and we will once again face two possibilities for the next time step (No tipping point state $S_{t+1}^{a00}$ and tipping point state $S_{t+1}^{a01}$). The rewards are calculated as the immediate utility of the current state $U_t^{a_t}$ and the expected utilities in following two stages ($U_{t+1}^0$, and $U_{t+2}^0$).

The bottom panel of Fig. A.1 demonstrates the algorithm in the post-threshold state. In the case of pre-threshold states, we use a combined pre- and post-threshold values in an expectation form in the approximation value function as it is pictured in the top panel of Fig. A.1. For post-threshold states, as the bottom panel in this figure shows, only post-threshold states are feasible, and therefore the value function approximation is calculated differently, without taking the expected values of future utilities. As this figure shows, the expected utility from the next two look-ahead states are used to calculate $\overline{V}_t$, the value function approximation (orange boxes in Fig. A.1). This value is then used in the Bellman equation to calculate $\hat{V}_t$, the value of the current state (green boxes in Fig. A.1):

$$\hat{V}_t = \max(U_t + \gamma \overline{V}_t) \tag{A.17}$$

The value $\hat{V}_t$ is then used to update the approximation in the previous state $\overline{V}_{t-1}$. In this fashion, we create a link from future states back to the initial ones (Powell, 2007). Therefore, while the pre- and post-threshold values diverge over time, the

[27] We thank the referee for clarifying this important distinction.

[28] The algorithm and its applications in the DICE framework are based on the previous work by Shayegh and Thomas (2015).

[29] These fixed actions are the optimal actions maximizing the immediate utility function. In the DICE case, the immediate optimal mitigation and SGE actions are zero because any non-zero action imposes an immediate cost to the system while its benefit in reducing damages will be felt only in the future.

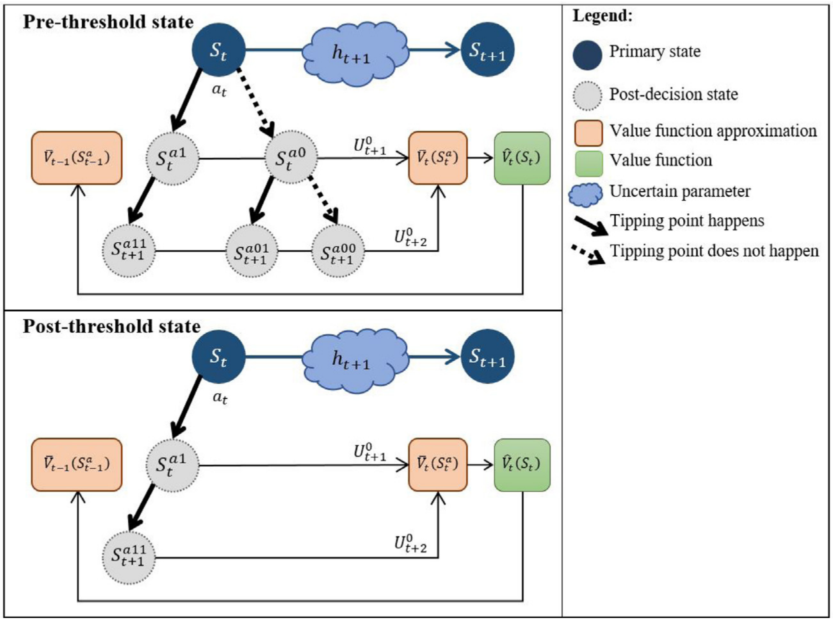


**Fig. A.1.** An example of the two-step-ahead algorithm for the DICE model with tipping point. The value of a pre-threshold state (top) is calculated by approximating the future values using the expected utilities of the states in two-step-ahead. The value of a post-threshold state (bottom) is calculated by approximating the future values using the utilities of the states in two-step-ahead. (For interpretation of the references to color in the text, the reader is referred to the web version of the article.)

updating algorithm ensures that the correct values of future states carried back to the current state are used to update the approximation coefficient.

This method does not deviate from the standard approximate dynamic programming methodology. It finds a simple and useful value function approximation. In the literature, this is usually done by constructing a basis function. The choice of a basis function is not arbitrary. Our point is not to compare all basis functions and pick the best one. Here, we present a method of constructing the value function that works much faster than the standard methods and generates a comparable set of results. Of course, each method has its limitations and cannot be used for every application.

Our method involves an alternative: it constructs a value function approximation using some projection of the model into the future. That is, it fixes the action and builds a decision tree in two levels. We calculate the utilities of the end nodes of the tree and combine them in a simple linear function. This is a function of the current state variables (the decision tree's starting node) and therefore works similarly to a basis function for approximating the future values of the model. The optimal action is the one that maximizes the value of the current state as a function of these utilities:

$$a_t^*(S_t) = \underset{a_t}{argmax}(U_t^{a_t} + \gamma \overline{V}_t(S_t^a)) = \underset{a_t}{argmax}(U_t^{a_t} + \gamma \overline{V}_t(U_{t+1}^0, U_{t+2}^0)) \tag{A.18}$$

As before, the value function $\overline{V}_t$ is the value approximation function that estimates the value of all future states based on the expected value of the two-step-ahead post-decision states. Therefore, the goal of this algorithm is to find a set of tunable parameters of a linear function that can approximate the value of all future states. As shown in Eq. (A.18), we use the utilities of the post-decision states to construct the approximation function $\overline{V}_t$. We find the optimal action that maximizes the sum of the current state utility and the approximated value of future states.

The choice of the number of steps ahead is merely a heuristic convenience. In principle, as the number of steps in a rolling horizon method increases, the decision tree grows exponentially and with it the complexity of the approximation. We choose a two-step algorithm based on earlier results from using a similar approximation technique for addressing uncertainty in climate sensitivity in the DICE model (Shayegh and Thomas, 2015).

For this problem we consider a very simple function approximation with only one parameter $\overline{V}_t(S_t^a) = \theta_t^m \times \mathbb{E}(U_{t+2}^0)$, where $\theta_t^m$ is the tunable parameter of the value function approximation at time $t$ and iteration $m$. This parameter defines the "policy" in our model. The initial value of this parameter is assumed to be zero, and it is updated at the end of each iteration. To find the expected utility $\mathbb{E}(U_{t+2}^0)$ we need to calculate the corresponding probabilities on each branch of the decision tree using Eq. (A.1). The idea of using $\overline{V}$ to approximate the value of future states is coming from conventional rolling horizon algorithms where due to discounting, the future utilities have a smaller contribution in the Bellman equation and therefore their contribution can be estimated by utilities of few steps ahead. However, unlike these techniques, we are not solving a series of finite-horizon models with approximated scrap values. We use the two-step-ahead approximation to build our

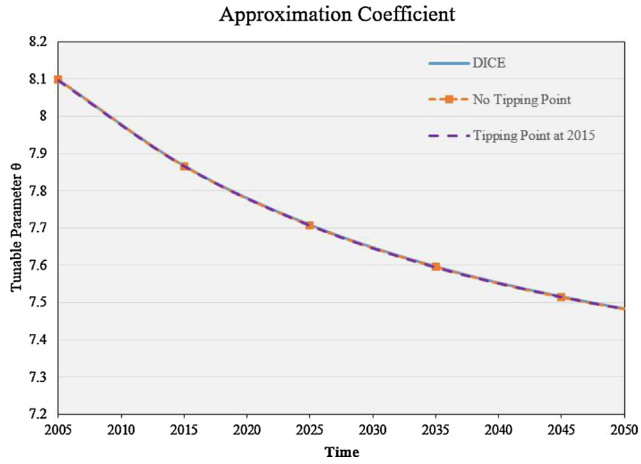


**Fig. A.2.** Optimal value of the tunable parameter $\theta$ changes over time but stays robust under different scenarios. The reason is that the variation in the value of future states is being captured in the utility functions of the future that we use our basis function. Therefore, as we expected, the optimal choice of the approximation parameter is invariant to different realizations of uncertain parameters in the model.

approximate value function, and then we use this approximation in a standard value iteration algorithm to find the optimal policy. We verify that across simulations, the optimal value of the tunable parameter of the value function approximation parameter is invariant to policy regime.

The value of the current state is calculated from $\hat{V}_t(S_t) = \max_{a_t}(U_t^{a_t} + \gamma \overline{V}_t(S_t^a))$. Once the optimal action $a^*$ is found, the next state $S_{t+1}$ can be constructed accordingly. The optimal value of $\hat{V}_t(S_t)$ will be then used to update the approximation function $\overline{V}_{t-1}(S_{t-1}^a)$ from the previous time step. This way the linkage between all future states and the current state is established. The algorithm works by moving forward in time, estimating the value function approximation ($\overline{V}$) and calculating the optimal value ($\hat{V}$) for every time step.

Once the algorithm reaches the last time step, ($t=60$), we use its value, consisting of the immediate utility (i.e. there are no future states), to update the value function approximation $\overline{V}_{59}$ in the previous time step. This process continues until all value function approximations are updated. The updating mechanism is based on the following stochastic gradient algorithm:

$$\theta_t^{m+1} = \theta_t^m - \alpha \times (\overline{V}_t - \hat{V}_{t+1}) \times U_{t+2}^0 \quad \text{(A.19)}$$

The step size $\alpha$ is chosen as $[U_{t+2}^0]^{-2}$ to simplify the updating equation and guarantees convergence. The new coefficient for the $m+1$ iteration is calculated as

$$\theta_t^{m+1} = \theta_t^m - \frac{\overline{V}_1 - \hat{V}_2}{U_1(S_2^0, 0)} \quad \text{(A.20)}$$

The newly updated value of $\theta_t^{m+1}$ will be used in the next iteration. In each iteration, moving from $t=0$ to $t=60$, we construct a forward looking value approximation for each time step. These values are then being updated at the next iteration using the new coefficient $\theta_t^{m+1}$. We continue iterating over this algorithm until the error, defined as the difference between $\theta_t^{m+1}$ and $\theta_t^m$ converges to zero within the maximum allowable tolerance. Fig. A.2 shows the optimal value of the tunable parameter $\theta$ over time and across three scenarios: Deterministic DICE, Stochastic model when tipping point is not reached, Stochastic model when tipping point is reached at 2015. As we expect, in all these scenarios the optimal policy remains the same and only varies over time.

- **Discounting**

  Conventional look-ahead algorithms only consider a limited horizon ahead while finding the optimal action in the current state and discard the rest of the states. Therefore, their performance is very dependent on the discounted value of discarded states. However in our algorithm, we do not solve a sequence of small finite horizon problems. We construct two deterministic steps ahead using a predefined control variable. The utilities of these two states are indeed functions of our current state variables and action and therefore, any linear combination of them can be considered a basis function. Such a novel basis function indeed preserve the key elements of the future values since it is using the endogenous mechanisms and parameters that are used to calculate the utilities of future states. In conventional look-ahead algorithms these utilities alone are estimated to capture the rest of states' utilities due to discounting. However, in our method we do not rely on

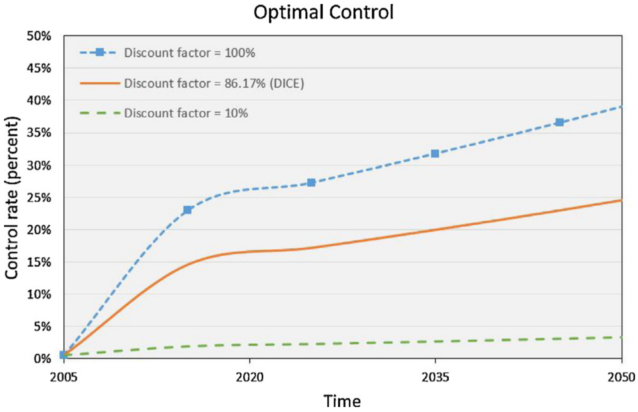


**Fig. A.3.** Optimal control rate under different discounting regimes. The continuous line in the middle shows the original DICE model. The box-dashed line at the top is the case with no discounting. The dashed line at the bottom is the case with a very high discount rate. In all cases the algorithm was able to find the optimal solution regardless of the discount rate. The performance of the proposed algorithm is independent of the choice of the discount rate.

such assumption. If this was the case then we would not have any tunable coefficient or constant action for the look-ahead states in the model. Instead, we would "solve" this smaller finite horizon problem and would find the optimal actions not only for the current state but for the next two states ahead as well. However, as described before, we are not using a rolling horizon algorithm. The two-steps-ahead states are constructed merely to use their utility functions as our basis in the approximation, and therefore this technique works regardless of discounting.

To verify this, here we solve the model under two extreme discount factors and compare with the standard DICE model discount factor (86.17%). First, we eliminate the discount rate and give equal weights to future values (discount factor = 100%). Second, we increase the discount rate so that the future states carry very low weight into the evaluation of the current state (discount factor = 10%). As shown in Fig. A.3, in all three cases our algorithm successfully solves the optimization problem and finds the optimal mitigation rate (control rate) as expected.

## Appendix B. Additional simulation results

Figs. B.1–B.4 present the simulation results of the optimal carbon price and the time path of utility under the different scenarios described in the text.

Fig. B.1 shows the carbon price under the pre-CTP scenario. The carbon price follows a path that qualitatively mimics the time path of mitigation. The carbon price under the No CTP—Ban SGE scenario, shown in Fig. B.1, starts at around 37$/tC in 2015 and reaches 301$/tC in 2125. When SGE is introduced, it reduces the carbon price of 33$/tC, from 301$/tC to 268$/tC. The carbon price under the unconstrained-rule is about 296$/tC by 2125. The carbon price under the reparation rule reaches 356$/tC by 2125. This value is lower than the carbon price under the Ban-rule, but much larger if compared to the Unconstrained rule.

Fig. B.2 shows welfare under the pre-CTP scenario. The welfare panels show how the utility changes over time, relative to the utility in the Unconstrained SGE scenario. Initially, the Unconstrained SGE scenario delivers less utility than the two alternative scenarios with a CTP. This is the case because at the beginning of the simulation temperature is very similar across rules and therefore, SGE is used at a cost without much of a benefit. But as temperatures start to diverge across scenarios, SGE use generates substantial gains relative to the alternative rules. But as time progresses, all policies look similar as mitigation reaches it maximum level and technological progress starts to dominate the story.

Fig. B.3 shows the carbon price under the post-CTP scenario. The Ban SGE rule exhibits the highest carbon price and the trajectory becomes slightly more inclined once the CTP is reached. The Unconstrained SGE rule shows the lowest value for the carbon price. The carbon price for the Remediation SGE rule lies in between the two other rules. It is always lower than the Ban SGE rule because the possibility of SGE in the future creates inter-temporal substitution away from mitigation and toward SGE, hence a lower carbon tax. Once SGE can be used after the CTP, the carbon price falls and converges toward the Unconstrained SGE rule.

Fig. B.4 shows welfare under the post-CTP scenario. As before, the Unconstrained SGE scenario delivers less utility than the two alternative scenarios with a CTP. But as soon as the CTP is reached in 2085, the difference increases and becomes negative. Under the Ban SGE rule, this negative difference stays large for a long period. Under the Remediation SGE rule, the difference with the Unconstrained SGE rule is large and negative as SGE overshoots to compensate for past-temperatures, and then as it quickly converges toward the Unconstrained SGE rule, the difference disappears.

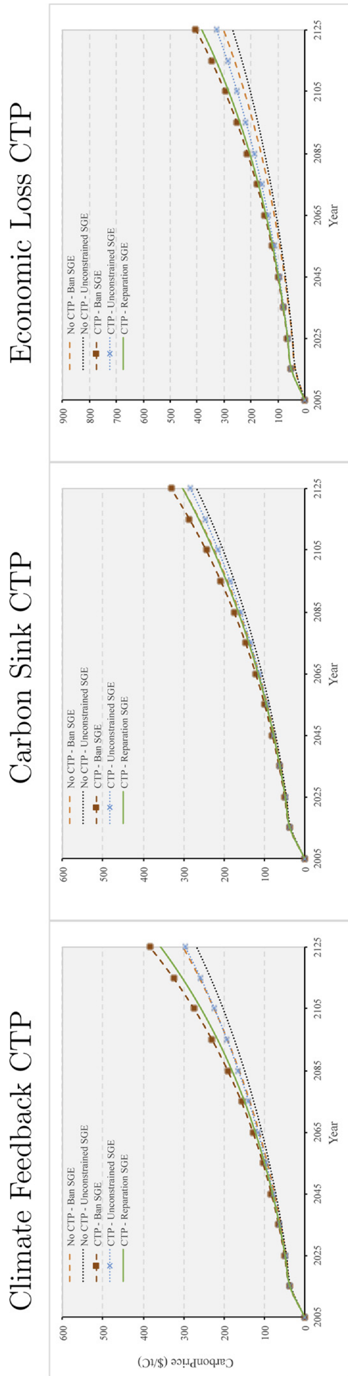


**Fig. B.1.** Pre-CTP carbon prices.

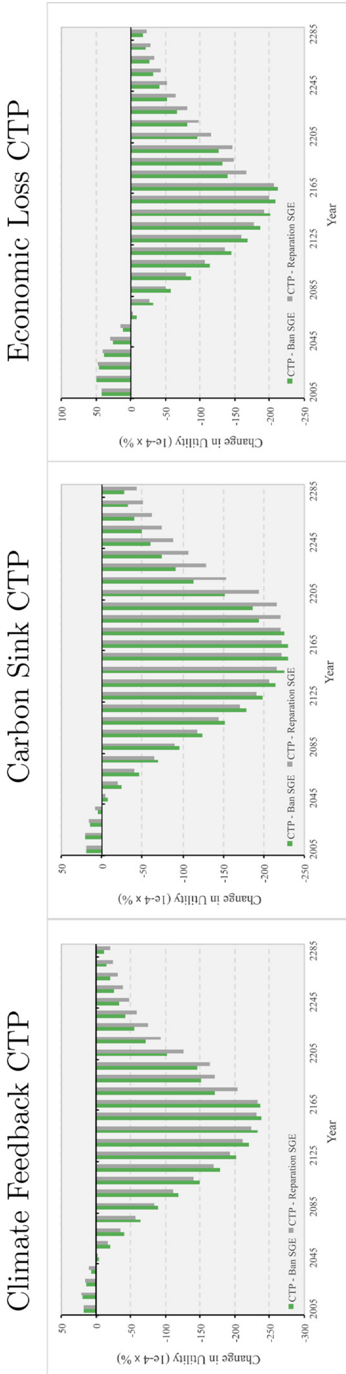


**Fig. B.2.** Pre-CTP utility.

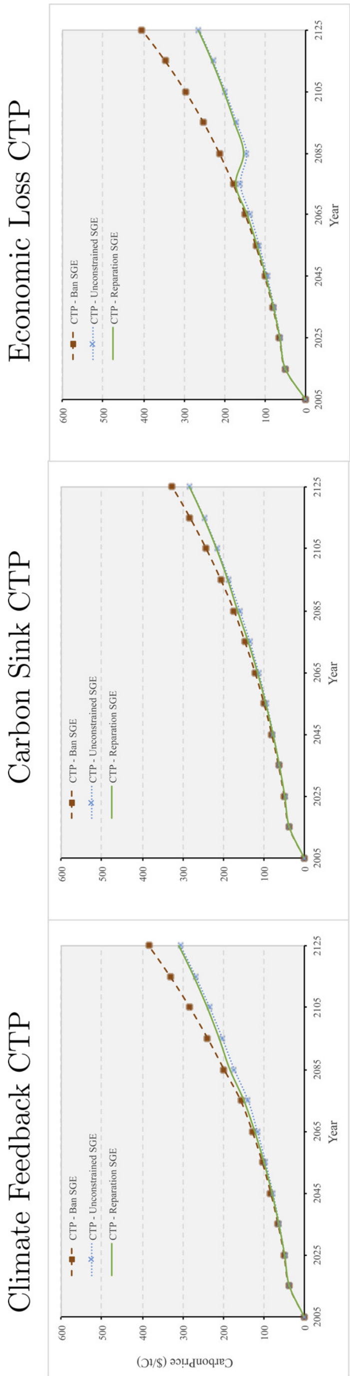


**Fig. B.3.** Post-CTP carbon price.

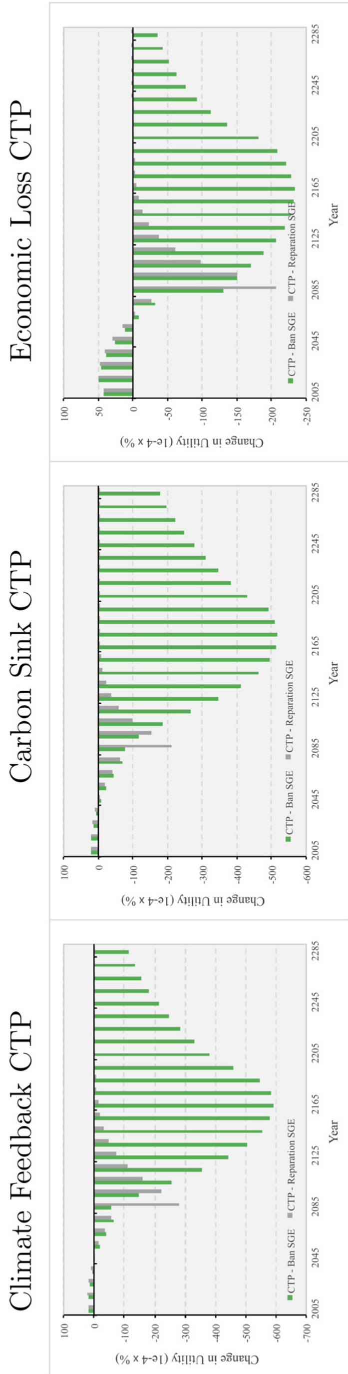


**Fig. B.4.** Post-CTP utility.

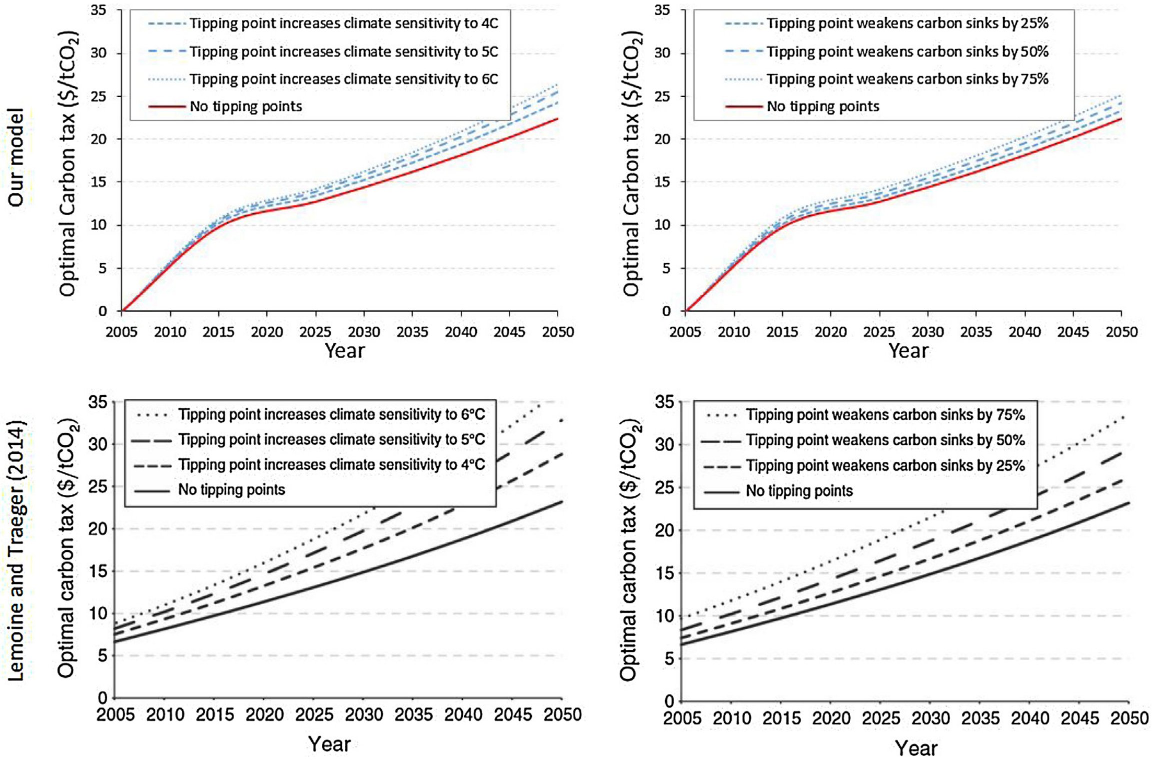


**Fig. C.1.** Comparison of temperature results in our model with the Lemoine and Traeger's model. We replicated the tipping points definition from their model but used the original DICE model's time step and carbon circulation structure. The initial Abatement rate used to calculate the optimal carbon tax in 2005 comes from the original DICE model's assumption in our model.

## Appendix C. Comparing our numerical solution to Lemoine and Traeger (2014)

In this appendix, we verify the two-step-ahead solution algorithm from Shayegh and Thomas (2015) by using it to solve the model in Lemoine and Traeger (2014). We demonstrate that our results are identical to their reported results, which were solved by them using a different solution algorithm. In that study, for numeric efficiency, they reformulated the DICE-2007 model to use effective labor units for capital and combined biosphere and shallow ocean stock for carbon dynamics. Moreover, they downscale the original decadal time steps in the DICE model to an annual step size. Their solution is based on approximating the value function using a $10^4$ basis of Chebychev polynomials.

Compared to Lemoine and Traeger (2014)'s solution method, our method is significantly simpler, is faster to converge, and uses only one tunable parameter for each approximation. We keep the original structure of the DICE model and use the three carbon circulation layers (atmosphere, upper ocean, and lower ocean). Furthermore, we keep the decadal structure of the DICE model. We use the same initial mitigation for the year 2005 that was used in the original DICE model. Despite these differences, we show that our results follow the Lemoine and Traeger (2014)'s results very closely.

We consider the two types of CTPs modeled by Lemoine and Traeger (2014) for this comparison: Climate Sensitivity and Carbon sink. We consider three levels of increased climate sensitivity and model them separately, and three levels of Carbon sink CTP intensity.[30] The Climate sensitivity tipping point changes the effect of emissions on temperature, and the Carbon sink tipping point changes the timing of such an effect. The results are shown in Figs. C.2 and C.3, for the pre-CTP policy regime, in order to see how the modeled policymaker adjusts to the possibility over time.

Fig. C.1 shows the optimal carbon tax pathway in both models. As mentioned above, we use the original DICE estimation for the initial value in the year 2005. Our model generates slightly lower optimal carbon taxes compared to the Lemoine and Traeger (2014)'s model. However such difference do not significantly affect the outcome in terms of carbon concentration and global mean temperature. C.2 shows the carbon concentrations, and Fig. C.3 shows temperature, both comparing the results of our model that has decadal time steps and a two-layer ocean with the results from Lemoine and Traeger (2014).[31]

[30] The Carbon sink CTP increases the lifetime of $CO_2$ in the atmosphere by reducing the fraction of atmospheric emissions that is transferred to the upper ocean layer at each time step. We reduce this fraction by either 25%, 50%, or 75%.

[31] The figures from Lemoine and Traeger (2014) are cut and pasted directly from their paper. Note that all the results in this section are without SGE and without the Economic Loss CTP, since we are solely concerned in replicating the model in Lemoine and Traeger (2014).

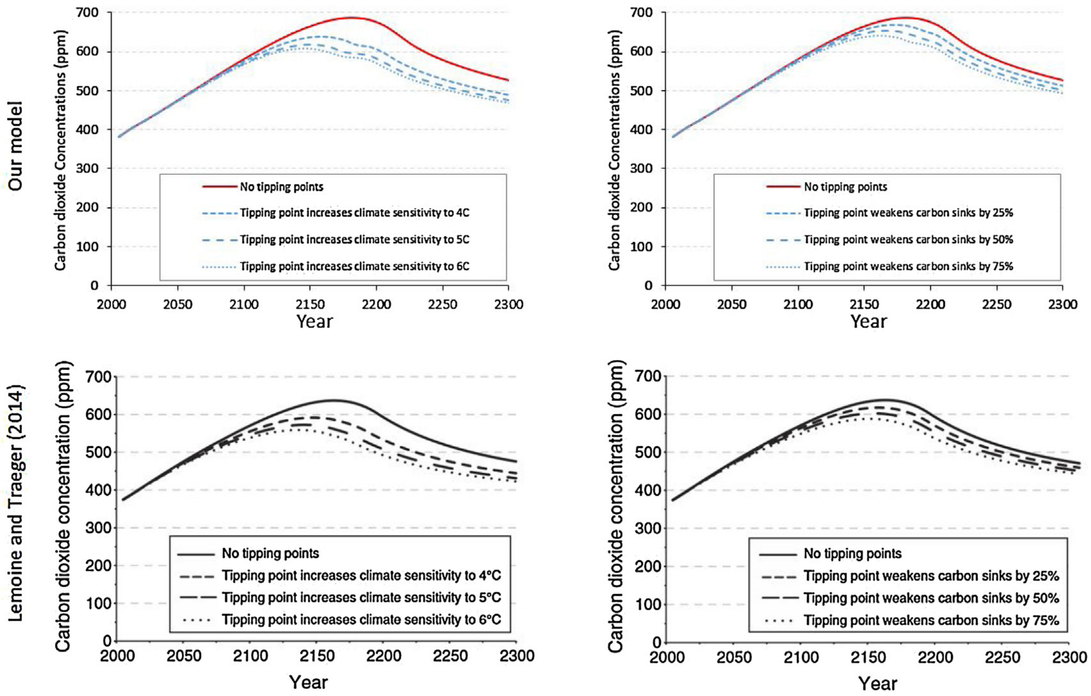


**Fig. C.2.** Comparison of carbon concentration results in our model with the Lemoine and Traeger's model. We replicated the tipping points definition from their model but used the original DICE model's time step and carbon circulation structure.

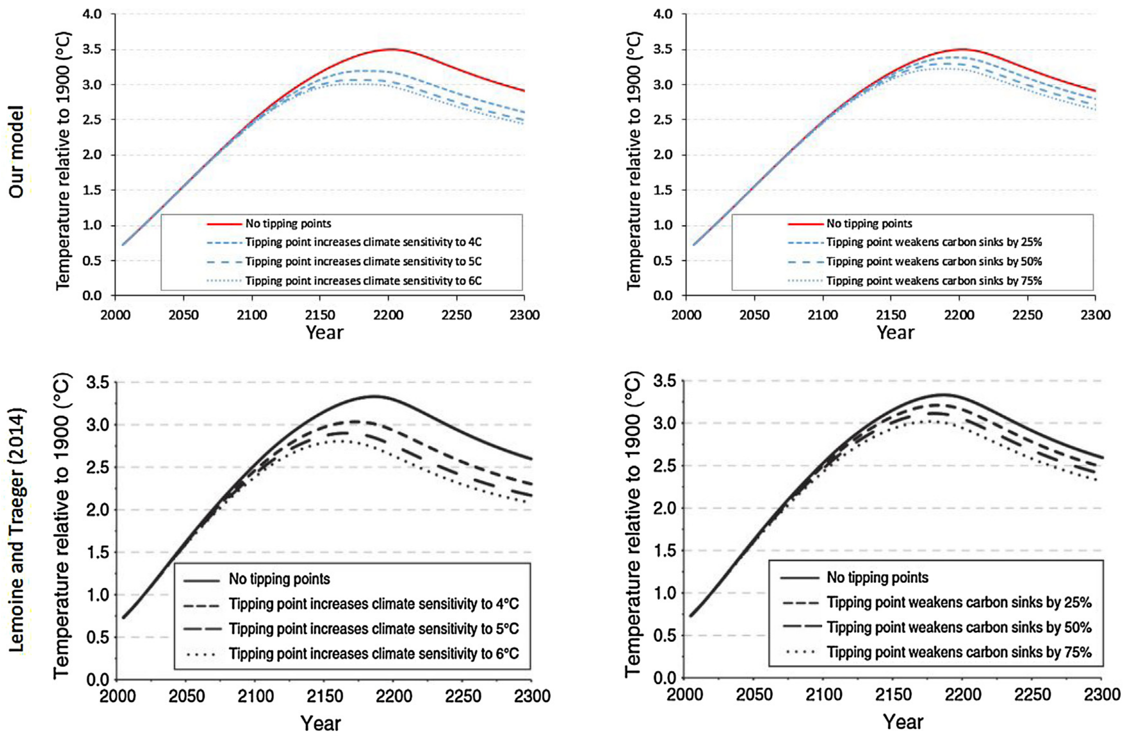


**Fig. C.3.** Comparison of temperature results in our model with the Lemoine and Traeger's model. We replicated the tipping points definition from their model but used the original DICE model's time step and carbon circulation structure.

The results are nearly identical.[32] This verifies that the solution method from Shayegh and Thomas (2015) is appropriate to use in this context.

## References

Baker, E., Solak, S., 2011. Climate change and optimal energy technology R&D policy. Eur. J. Oper. Res. 213 (2), 442–454.
Barrett, S., 2008. The incredible economics of geoengineering. Environ. Resour. Econ. 39, 45–54.
Barrett, S., Lenton, T.M., Millner, A., Tavoni, A., Carpenter, S., Anderies, J.M., Chapin III, F.S., Crépin, A.-S., Daily, G., Ehrlich, P., et al., 2014. Climate engineering reconsidered. Nat. Clim. Change 4 (7), 527–529.
Bertsekas, D.P., Bertsekas, D.P., Bertsekas, D.P., Bertsekas, D.P., 1995. Dynamic Programming and Optimal Control, vol. 1. Athena Scientific Belmont, MA.
Bickel, J., Agrawal, S., 2013. Reexamining the economics of aerosol geoengineering. Clim. Change 119 (3–4), 993–1006.
Bickel, J., Lane, L., 2009. An analysis of climate engineering as a response to climate change. Copenhagen Consensus Center Report 40.
Blackstock, J.J., Long, J.C.S., 2010. The politics of geoengineering. Science 327 (5965), 527.
Bony, S., Bellon, G., Klocke, D., Sherwood, S., Fermepin, S., Denvil, S., 2013. Robust direct effect of carbon dioxide on tropical circulation and regional precipitation. Nat. Geosci. 6 (6), 447–451.
Brander, L.M., Rehdanz, K., Tol, R.S., Van Beukering, P.J., 2012. The economic impact of ocean acidification on coral reefs. Clim. Change Econ. 3 (1).
Cai, Y., Judd, K., Lontzek, T., 2013. The social cost of stochastic and irreversible climate change. Working Paper 18704. National Bureau of Economic Research.
Crutzen, P.J., 2006. Albedo enhancement by stratospheric sulfur injections: a contribution to resolve a policy dilemma? Clim. Change 77, 211–219.
Emmerling, J., Tavoni, M., 2013. Geoengineering and abatement: a 'flat' relationship under uncertainty. Fondazione Eni Enrico Matei, Nota Di Lavoro 2013.031.
Goes, M., Tuana, N., Keller, K., 2011. The economics (or lack thereof) of aerosol geoengineering. Clim. Change 109 (3–4), 719–744.
Goodson, J.C., Thomas, B.W., Ohlmann, J.W., 2015. A rollout algorithm framework for heuristic solutions to finite-horizon stochastic dynamic programs. Technical report, Working paper. Saint Louis University http://www.slu.edu/goodson/papers/GoodsonRolloutFramework.pdf (accessed on June 18).
Gramstad, K., Tjøotta, S., 2010. Climate engineering: cost benefit and beyond. University of Bergen, Department of Economics, No. 05/10.
Heckendorn, P., Weisenstein, D., Fueglistaler, S., Luo, B.P., Rozanov, E., Schraner, M., Thomason, L.W., Peter, T., 2009. The impact of geoengineering aerosols on stratospheric temperature and ozone. Environ. Res. Lett. 4 (4), 045108.
Heutel, G., Moreno-Cruz, J., Shayegh, S., 2015. Solar geoengineering, uncertainty, and the price of carbon. NBER Working Paper 21355.
Keith, D.W., 2000. Geoengineering the climate: history and prospect. Annu. Rev. Energy Environ. 25, 245–284.
Keith, D.W., Parson, E., Morgan, M.G., 2010. Research on global sun block needed now. Nature 463, 426–427.
Kolstad, C.D., 1996. Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. J. Environ. Econ. Manag. 31 (1), 1–18.
Lemoine, D., Traeger, C., 2014. Watch your step: optimal policy in a tipping climate. Am. Econ. J.: Econ. Policy 6 (1), 137–166.
Lontzek, T.S., Cai, Y., Judd, K.L., Lenton, T.M., 2015. Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. Nat. Clim. Change 5, 441–444, http://dx.doi.org/10.1038/nclimate2570.
McClellan, J., Keith, D.W., J.A.R., 2012. Cost analysis of stratospheric albedo modification delivery systems. Environ. Res. Lett. 7, 034019.
Moreno-Cruz, J., 2015. Mitigation and the geoengineering threat. Resour. Energy Econ. 41, 248–263.
Moreno-Cruz, J., Keith, D., 2013. Climate policy under uncertainty: a case for solar geoengineering. Clim. Change 121, 431–444.
Naevdal, E., Oppenheimer, M., 2007. The economics of the thermohaline circulation – a problem with multiple thresholds of unknown locations. Resour. Energy Econ. 29 (4), 262–283.
Nordhaus, W.D., 1993. Optimal greenhouse-gas reductions and tax policy in the "dice" model. Am. Econ. Rev., 313–317.
Nordhaus, W.D., 2008. A Question of Balance: Weighing the Options on Global Warming Policies. Yale University Press.
Pindyck, R.S., 2013. Climate change policy: what do the models tell us? J. Econ. Lit. 51 (3), 860–872.
Pongratz, J., Lobell, D.B., Cao, L., Caldeira, K., 2012. Crop yields in a geoengineered climate. Nat. Clim. Change 2 (2), 101–105.
Powell, W.B., 2007. Approximate Dynamic Programming: Solving the Curses of Dimensionality, vol. 703. John Wiley & Sons.
Rasch, P., Tilmes, S., Turco, R., 2008. An overview of geoengineering of climate using stratospheric sulphate aerosols. Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci. 366, 4007–4037.
Robock, A., Oman, L., Stenchikov, G.L., 2008. Regional climate responses to geoengineering with tropical and arctic SO_2 injections. J. Geophys. Res.: Atmos. (1984–2012) 113 (D16).
Shayegh, S., Thomas, V.M., 2015. Adaptive stochastic integrated assessment modeling of optimal greenhouse gas emission reductions. Clim. Change 128 (1–2), 1–15.
Truong, V.-A., 2014. Approximation algorithm for the stochastic multiperiod inventory problem via a look-ahead optimization approach. Math. Oper. Res. 39 (4), 1039–1056.
Victor, D.G., 2008. On the regulation of geoengineering. Oxf. Rev. Econ. Policy 24, 322–336.
Wigley, T., 2006. A combined mitigation/geoengineering approach to climate stabilization. Science 314, 452–454.

[32] Our model generates slightly higher temperature that can be attributed to the longer time steps.

Please cite this article in press as: Heutel, G., et al., Climate tipping points and solar geoengineering. J. Econ. Behav. Organ. (2016), http://dx.doi.org/10.1016/j.jebo.2016.07.002