

Solar Geoengineering, Learning, and Experimentation*

David L. Kelly¹, Garth Heutel², Juan B. Moreno-Cruz³ and Soheil Shayegh⁴

¹Department of Economics, University of Miami

²Department of Economics, Georgia State University, and NBER

³School of Environment, Enterprise and Development, University of Waterloo

⁴RFF-CMCC European Institute on Economics and the Environment (EIEE),
Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy

July 31, 2023

Abstract

Solar geoengineering (SGE) can offset climate change by directly reducing temperatures. Both SGE and climate change itself are surrounded by great uncertainties. Implementing SGE affects learning about these uncertainties. We model endogenous learning over two uncertainties: the sensitivity of temperatures to carbon concentrations (the climate sensitivity), and the effectiveness of SGE in lowering temperatures. We present both theoretical and simulation results from an integrated assessment model, focusing on the informational value of SGE experimentation. Surprisingly, under current calibrated conditions, SGE deployment slows learning, causing a less informed decision. For any reasonably sized experimental SGE deployment, the temperature change becomes closer to zero, and thus more obscured by noisy weather shocks. Still, some SGE use is optimal despite, not because of, its informational value. The optimal amount of SGE is very sensitive to beliefs about both uncertainties.

*We thank David Anthoff, Simon Dietz, 3 anonymous referees, and conference and seminar participants at the AERE Summer Conference, the University of California at Berkeley, Colorado State University, the SEA annual conference, the WCERE, the 17th Occasional Workshop on Environmental and Resource Economics, and the University of Wisconsin Madison for valuable comments.

1 Introduction

Under current trajectories of greenhouse gas emissions, there exists only about a 5% chance of achieving the Paris Agreement's goal of limiting the global mean surface temperature increase to "well below" 2°C above pre-industrial levels by 2100 (Raftery et al., 2017). Achieving this goal through emissions abatement alone is a formidable challenge (Fitzpatrick and Kelly, 2017; Brown et al., 2019). Hence solar geoengineering (SGE), which has the potential to slow global warming relatively quickly, has gained popularity among some policy makers and researchers. Some SGE techniques mimic the climate impact of major volcanic eruptions by injecting reflective particles into the upper atmosphere to block and reflect solar radiation back into the space (Ramaswamy et al., 2018). SGE presents a number of difficulties.¹ Nonetheless, SGE is inexpensive and conceptually attractive, and a significant part of the optimal policy in many integrated assessment models.²

Here we consider an aspect of SGE not previously studied in integrated assessment modeling: the potential for endogenous learning over uncertainties surrounding SGE and the climate. We focus on two sources of uncertainty. The first is uncertainty about SGE effectiveness, which determines the temperature reduction from a given amount of SGE. The second is uncertainty about climate sensitivity, which is the equilibrium temperature increase resulting from a doubling of atmospheric carbon concentration. We model this as uncertainty in the climate feedback parameter. Both of these uncertainties are endogenous, so active learning or experimentation is possible. We study optimal experimentation with SGE, both theoretically and computationally with an integrated assessment model.

Intuition suggests that implementing SGE can increase our knowledge about it, since implementation is a type of field experiment. However, given the calibrated current initial conditions, we find a surprising result: implementing SGE actually slows learning about both SGE effectiveness and climate sensitivity, resulting in a less-informed decision. We decompose SGE's impact on learning into two effects. A *signal strength effect* comes from the effect of SGE on the temperature signal. If SGE makes the temperature change larger in magnitude and thus more visible amidst the noisy weather shocks, then the signal strength effect is positive and SGE speeds learning. However, a

¹For example, SGE performed via injecting sulfur particles causes harmful side effects (Keller et al., 2014; Proctor et al., 2018; Abatayo et al., 2020), does not reduce impacts like ocean acidification that are caused directly by CO₂, and has a relatively short half life, necessitating frequent injections. See Klepper and Rickels (2014) and Heutel et al. (2016a) for reviews of the economics of solar geoengineering. For a review of the science of SGE, see National Research Council, et. al. (2015). For a broader review of the social science of SGE, see Flegel et al. (2019).

²For example, Heutel et al. (2018) shows that optimal SGE policy reduces up to 50% of the radiative forcing (instantaneous warming effect) from CO₂.

small or medium-scale SGE deployment moves the temperature change closer to zero, and so the signal strength effect is negative, slowing learning about both SGE effectiveness and climate sensitivity. Only an unrealistically large SGE deployment causes a temperature change large enough in absolute value to create a positive signal strength effect which speeds learning. Second, a *noise amplification effect* causes the climate system to be more variable and slows learning about climate sensitivity (though does not affect learning about SGE effectiveness). Because the effectiveness of SGE deployment in reducing the temperature is uncertain, SGE causes the climate system to be more subjectively variable, slowing learning.

We then compute optimal SGE and abatement policies in an integrated assessment model (IAM) with endogenous learning. By “optimal,” we mean the policy path that maximizes the sum of discounted utility, given the costs and benefits of SGE, abatement, and climate change. Our IAM extends the well-known DICE model (Nordhaus, 2016) in several ways. First, in our model SGE is available as an additional policy tool. Second, we model uncertainty over two features of the model: climate sensitivity and SGE effectiveness. We calibrate SGE effectiveness and its uncertainty using volcano eruption data. Third, the planner learns endogenously and updates beliefs over time as a result of SGE deployment and greenhouse gas (GHG) abatement.³ The computational results show that optimal SGE deployment is never large enough to produce informational gains, and in fact the information loss acts as a disincentive to using SGE. Some SGE is still optimal because of its other benefits. Thus, SGE is used despite of, not because of, its effect on learning. Optimal SGE is highly sensitive to beliefs about climate sensitivity and SGE effectiveness. In the mean optimal simulation, SGE usage offsets at most 4% of the radiative forcing from CO₂. However, when beliefs are that climate change is severe and can be effectively reduced via SGE, deployment increases to 14%. This is consistent with the idea of using SGE as a last resort if climate change becomes severe, although the planner still uses a small amount of SGE even when beliefs are that the climate sensitivity is low.

The type of SGE experiment we consider is a planet-wide implementation. Our results show that a large-scale, planet-wide implementation experiment speeds learning but is not optimal because of the deployment costs and harmful side effects. Our results also show that a small-scale, planet-wide implementation experiment slows learning but is nonetheless optimal. Other types of experiments

³In the terminology of LaRiviere et al. (2018), we model active adaptive management, since learning is taken into account when making decisions.

are also possible. For example, researchers may glean information from laboratory and/or localized, rather than planet-wide, field experiments of the kind suggested by Keith et al. (2014). In reality both types of experiments can be valuable. Laboratory experiments generate important predictions, that ultimately must be tested in the uncontrolled, planet-wide environment. We model this process by using volcano data to calibrate a distribution of prior uncertainty, which is then updated by observing the temperature following the planet-wide implementation. Although our model does not consider other types of experiments and thus reaches no conclusions about their optimality or effect on learning, such experiments could be incorporated into our framework as modifications of the prior distribution. In turn, we show how the prior distribution affects optimal planet-wide implementation and learning through implementation.

A large body of literature examines optimal GHG abatement policy when climate processes are uncertain and learning reduces uncertainty over time.⁴ Lemoine and Rudik (2017) provide an overview of the literature and methods used in dealing with uncertainty and learning in IAMs. For example, Kelly and Kolstad (1999) model Bayesian learning about climate sensitivity and find that learning takes up to 100 years to become informative. Leach (2007) considers uncertainty over a parameter governing the persistence of temperature, also finding that learning is very slow. Jensen and Traeger (2013) incorporate Bayesian learning in the DICE model, and Lemoine and Traeger (2014) add climate tipping points and allow for learning about the tipping point temperature threshold. Other studies have extended the model to include fat-tailed uncertainty (Kelly and Tan, 2015; Shayegh and Thomas, 2015; Hwang et al., 2017).

In general, the literature finds learning is relatively slow.⁵ Therefore, near-term abatement policy with learning is similar to policy without learning. Our paper builds upon this literature by showing that, in an environment with multiple uncertainties, SGE experimentation can actually slow learning by magnifying uncertainty in the climate system. Overall learning about SGE effectiveness is slow, because optimal SGE levels are relatively small.

Several studies introduce SGE into IAMs. Bickel and Lane (2009), Gramstad and Tjøtta (2010), Goes et al. (2011), and Bickel and Agrawal (2013) modify the DICE model to incorporate SGE, but without allowing for epistemic uncertainty. Emmerling and Tavoni (2018) use a different IAM,

⁴See Pastor et al. (2020); LaRiviere et al. (2018) for literature reviews of uncertainty in climate economic modeling.

⁵Hwang et al. (2019) shows that active learning through knowledge acquisition (e.g. investment in more accurate climatic observations) can reduce uncertainty more rapidly than passive learning.

WITCH, and incorporate SGE, with uncertainty modeled as a binary probability outcome.⁶ Heutel et al. (2016b) analyze the use of SGE with a risk of tipping points, and Helweggen et al. (2019) add the possibility of solar radiation management failure in their analysis. Most closely related to this paper, Heutel et al. (2018) modify the DICE model and allow for general forms of epistemic uncertainty. They demonstrate how uncertainty over climate sensitivity and SGE damages affects optimal policy. Our paper advances this literature by (1) considering a previously unexplored uncertainty, SGE effectiveness,⁷ in addition to climate sensitivity, (2) calibrating SGE effectiveness and uncertainty using volcano data, (3) considering the effect of endogenous learning on SGE policy, and (4) considering SGE when the climate sensitivity is uncertain. Under these extensions, optimal SGE deployment falls considerably relative to the prior literature, from a maximum of about 50% of radiative forcing to where the mean simulation has a maximum SGE of about 10% of radiative forcing over time.

Our paper is also related to an emerging literature on learning, experimentation, and R&D on SGE. One can learn about SGE through costly R&D (e.g. Quaas et al., 2017), small-scale field experiments (e.g. Keith et al., 2014), and/or learning via planetary-wide experimentation considered here. Much of this literature explores the moral hazard issue of whether or not learning about SGE will reduce the likelihood of abatement (e.g. Barrett et al., 2014; Quaas et al., 2017). Like Quaas et al. (2017), our results show that experimentation with SGE may reduce future abatement, but is nonetheless optimal. However, our model advances this literature by considering endogenous learning about multiple uncertainties within an IAM, yielding important new insights on how the size of the experiment affects learning and the interplay between SGE experimentation and learning about the climate sensitivity. For example, in Quaas et al. (2017) R&D about SGE exogenously speeds learning about the climate sensitivity, whereas here a planetary wide experiment with SGE endogenously slows learning about the climate sensitivity.

⁶Other papers study uncertainty over SGE but not in the framework of IAMs. Moreno-Cruz and Keith (2013) model uncertainty with a binary outcome using an analytical model. See Kravitz and MacMartin (2020) for a review of key uncertainties in SGE.

⁷Some research (e.g. Quaas et al., 2017) refer to effectiveness as a general theoretical uncertainty which could represent SGE damage uncertainty, uncertain ability to reduce temperatures, or other uncertainties. Here effectiveness is calibrated specifically as the ability of SGE to reduce temperatures.

2 Model

Our IAM is based on the Dynamic Integrated Climate-Economy (DICE) model (Nordhaus, 2016). Heutel et al. (2018) modified DICE to include SGE; here we extend that modification of DICE to also include endogenous learning about climate sensitivity and SGE effectiveness. We adopt the methodology of Traeger (2014), which reduces the dimension of the state space. In particular, we update the parameters given in Traeger (2014) to match the 2016 version of DICE. We further modify this calibration to match the speed at which the climate adjusts to CO₂ shocks emphasized by Dietz et al. (2020). More details about the original DICE model and its modification to include SGE are included in the above-cited papers. The full model is detailed in the online appendix. In this section we focus on the introduction of endogenous learning to the model.

2.1 Overview

DICE is a representative agent model of the global economy and climate. The economic portion of the model is a standard aggregate growth model, with an endogenously evolving capital stock, exogenous growth in total factor productivity and in population, and an endogenous choice of the optimal amount of consumption and investment to maximize the net present value of the sum of utility based on per-capita consumption. This economic portion is integrated with a climate model, in which the carbon emissions generated from economic production affect the carbon stock of the atmosphere and oceans, and the atmospheric carbon stock in turn affects average temperatures. Finally, temperature affects the economy through a damage function, in which net output deviates from gross output based on the loss from temperature impacts. The climate policy choice variable in the original DICE model is abatement α , which is a number between 0 and 1 representing the fraction of total carbon emissions abated.⁸

The extension of DICE to allow SGE adds a second climate policy variable, solar geoengineering g . This policy variable is defined so that g takes values between 0 and 1 and represents the intensity of SGE use. Radiative forcing is the difference between the incoming solar heat energy and outgoing energy radiated from Earth. When radiative forcing is positive, the planet will warm because more

⁸In DICE, α is allowed to be greater than one only after the year 2160, reflecting the availability of carbon dioxide removal (CDR) technologies. Other studies, in line with the IPCC Special Report, have proposed a much earlier timeline for CDR deployment (Hänsel et al., 2020). As speculation about CDR timeline would have a direct impact on the optimal levels of abatement and SGE deployment, we impose $\alpha \leq 1$ to avoid unnecessary distractions from the main objective of our study. We use α to denote abatement, rather than μ , which was used in the original DICE model.

energy is entering than is leaving. When $g = 0$, the radiative forcing is unchanged by SGE. When g is positive, radiative forcing is reduced by a fraction ϕg , where ϕ measures how effective SGE is at reducing radiative forcing. Unlike the abatement policy variable α which must be between 0 and 1, SGE g can take any value greater than 0. If $g > 1/\phi$, then SGE more than offsets the warming effect caused by greenhouse gases, and the planet will cool even faster than if GHG concentrations immediately return to pre-industrial levels.

2.2 SGE Costs and Damages

Following Heutel et al. (2018), we posit the following model of costs and damages:

$$Y = (1 - \Lambda - D) Q, \quad (1)$$

$$\Lambda = \theta_1(t) \alpha^{\theta_2} + \theta_{GE} g^{\theta_3}, \quad (2)$$

$$D = \pi_T T_{AT}^2 + \pi_m (M - M_{1750})^2 + \pi_g g^2, \quad (3)$$

where Q is gross world output, Y is world output net of damages, D , and spending on emissions abatement and SGE, Λ . The cost function Λ gives the percentage of output spent on emissions abatement and SGE. The cost of a backstop technology which eliminates all carbon emissions, $\alpha = 1$, is $\theta_1(t)$, which falls exogenously over time. In our notation, functions of time are variables that evolve exogenously over time. We suppress the time dependence of endogenous variables, and use primes to denote next-period values.

Nordhaus (2016) combines damages associated with higher carbon stocks and temperatures into a single function of atmospheric temperature. However, since SGE directly affects temperature while abatement directly affects the carbon stock, we use the damage function (3), which is a function of both the deviation of atmospheric temperature from recent averages, T_{AT} , and the stock of atmospheric carbon M above preindustrial levels M_{1750} .⁹ Following Heutel et al. (2018), we also include damages caused directly by SGE g in the damage function (for example, warming of the tropical cold-point tropopause leading to the depletion of the ozone layer Heckendorn et al., 2009).

⁹Traeger (2014) replaces the endogenous stock of carbon in the ocean with an exogenous function that determines carbon absorption by the ocean and biosphere. Hence, we combine the damages from ocean and atmospheric carbon into just a function of atmospheric carbon.

2.3 Carbon Cycle, Radiative Forcing, and Temperature

Following Traeger (2014), the carbon cycle evolves according to:

$$E = \sigma(t)(1 - \alpha)Q + E_{LAND}(t), \quad (4)$$

$$M' = M_{1750} + (1 - \delta_m(t))(M - M_{1750}) + E, \quad (5)$$

Here E is carbon emissions, $\sigma(t)$ is the emissions intensity of output, and $E_{LAND}(t)$ are exogenous emissions from land use changes. In equation (5), $\delta_m(t)$ is the fraction of atmospheric carbon absorbed by the biosphere and deep ocean. The exogenous absorption rate declines over time as the oceans become more saturated with carbon.

Radiative forcing F is given by:

$$F' = \left(\eta \log_2 \left[\frac{M'}{M_{1750}} \right] + F_{EX}(t+1) \right) (1 - \phi g), \quad (6)$$

Here η is the forcing coefficient for carbon and $F_{EX}(t)$ is exogenous forcing from other greenhouse gasses and aerosols. Let $\phi \in [0, 1]$ denote the parameter which determines how effective SGE is at reducing radiative forcing.¹⁰

Atmospheric temperature evolves according to:

$$T'_{AT} = T_{AT} + \xi_1 \left\{ F' - \xi_2 T_{AT} - \xi_3 [T_{AT} - T_{LO}](t) \right\} + \epsilon'. \quad (7)$$

Here ξ_1 governs the speed at which the temperature adjusts to changes in radiative forcing. The temperature differential between the ocean and atmosphere, $[T_{AT} - T_{LO}](t)$ is exogenous as in Traeger (2014). The parameter ξ_2 measures the magnitude of feedback effects. For example, as the temperature rises, ice melts, reducing the albedo effect, causing still higher temperatures (a positive feedback). We thus call ξ_2 the *feedback parameter*. Finally, ϵ is an iid random weather shock.

The *climate sensitivity* $\Delta T_{2\times}$ is defined as the equilibrium increase in temperature that would arise from a doubling of the pre-industrial carbon stock, without any SGE. Combining equations (6) and (7) at the steady state implies the climate sensitivity is $\Delta T_{2\times} = \eta/\xi_2$. As described below, we

¹⁰Here SGE affects radiative forcing through a multiplicative effect, as modeled in Heutel et al. (2018). Other papers instead model SGE as affecting radiative forcing through an additive term. Heutel et al. (2018) demonstrate that this modeling choice has little effect on optimal policy outcomes (see their Figure 6 and Appendix Figure A12).

model uncertainty in the climate sensitivity through uncertainty in the feedback parameter.

The timing is such that both CO_2 and SGE affect radiative forcing in $t + 1$, and therefore the temperature in $t + 1$.

2.4 Uncertainty

We model two sources of uncertainty, one over SGE and one over the climate system. First, we allow the SGE effectiveness parameter, ϕ , to be uncertain. The planner treats ϕ as a random variable $\tilde{\phi}$. A given quantity of SGE g has an uncertain effect on forcing due to uncertainties in how long the sulfur particles will remain in the stratosphere, how well sulfur particles mix across the globe, and effects on cloud formation. We do not explicitly model uncertainty in the parameter representing damages from SGE, π_g . This is analogous to the choice of learning over the climate sensitivity rather than over the climate damages.¹¹

Second, we model uncertainty in the climate sensitivity by allowing the feedback parameter, ξ_2 , to be uncertain. The planner treats ξ_2 as a random variable $\tilde{\xi}_2$. Because climate sensitivity $\Delta T_{2\times}$ is η/ξ_2 , which is a function of ξ_2 , treating the feedback parameter ξ_2 as uncertain implies that climate sensitivity $\Delta T_{2\times}$ is also uncertain. While much of the literature models climate sensitivity uncertainty through uncertain feedbacks (e.g. Roe and Baker, 2007; Kelly and Tan, 2015), this is not the only way to model uncertainty in climate sensitivity. Kelly and Kolstad (1999) and Heutel et al. (2018) model climate sensitivity uncertainty by assuming that the forcing parameter η is uncertain, not the feedback parameter ξ_2 . When the feedback parameter is distributed normally, the climate sensitivity is a function of the reciprocal of a normally-distributed random variable and therefore has a reciprocal normal distribution. Roe and Baker (2007) and Kelly and Tan (2015) derive analytical approximations of this distribution which are asymptotically equivalent to a Pareto distribution with infinite mean and variance (see Kelly and Tan, 2015, for a detailed discussion). A common alternative is to treat the climate sensitivity itself as the uncertain parameter, and assume a log-normal distribution (Spafford and MacDougall, 2020; Wagner and Weitzman, 2018). Quantitatively, the Pareto has a fatter tail and thus the tail of the distribution is more important for the results. Further, an upper limit to the climate sensitivity must be specified to make the problem finite (see

¹¹See the 2013 IPCC report for a discussion of the extensive uncertainty with respect to SGE effectiveness. To our knowledge, this paper is the first to analyze uncertainty over SGE effectiveness. It is possible that SGE effectiveness and damages are linked in that if SGE is more (less) effective, then SGE is more (less) damaging. We do not consider this possibility for computational reasons.

the online appendix).

We assume that radiative forcing net of SGE, $\frac{F'}{(1-\phi g)}$, is not observed, which prevents the planner from instantly learning the effectiveness of SGE by inverting the forcing equation (6).¹² Furthermore, the random weather shock prevents the planner from immediately learning the true feedback parameter by inverting the temperature equation (7). Equation (6) shows that the uncertainty over SGE effectiveness ϕ affects forcing from carbon (the first term in parentheses) in the same way that it affects exogenous forcing F_{EX} .¹³

A change of variables simplifies the above equations. Let H' denote the residual after removing the known influences on next-period temperature T'_{AT} .

$$\begin{aligned} H' &\equiv T'_{AT} - \xi_1 \left(\eta \log_2 \left[\frac{M'}{M_{1750}} \right] + F_{EX}(t+1) \right) + \xi_1 \xi_3 [T_{AT} - T_{LO}](t) \\ &= -\xi_1 \left(\eta \log_2 \left[\frac{M'}{M_{1750}} \right] + F_{EX}(t+1) \right) g\tilde{\phi} + (1 - \xi_1 \tilde{\xi}_2) T_{AT} + \epsilon'. \end{aligned} \quad (8)$$

Let \hat{F}' denote gross forcing before SGE

$$\hat{F}' \equiv \xi_1 \left(\eta \log_2 \left[\frac{M'}{M_{1750}} \right] + F_{EX}(t+1) \right) \quad (9)$$

and $\tilde{\beta}$ denote the inertia of the temperature, which we call the *net feedback parameter*:

$$\tilde{\beta} \equiv 1 - \xi_1 \tilde{\xi}_2. \quad (10)$$

Equation (7) then simplifies to:

$$H' = \tilde{\beta} T_{AT} - \tilde{\phi} \hat{F}' g + \epsilon'. \quad (11)$$

This implies that the climate sensitivity is

$$\Delta T_{2\times} = \frac{\eta \xi_1}{1 - \tilde{\beta}}. \quad (12)$$

Equation (11) is a linear regression equation in which the planner tries to determine for example,

¹²It is possible that in the future that governments will measure radiative forcing if SGE is deployed. In this case, learning speed could be increased by using equation (6) to estimate ϕ and equation (7) to estimate feedbacks separately. However, this presumes that all other parameters in equation (6) are known with certainty. If not, then the problem of SGE slowing learning about other uncertainties remains.

¹³Strictly speaking, F_{EX} includes the radiative forcing from aerosols such as SO_2 , which is governed by the uncertain parameter ϕ . However, in the interests of simplification, we assume all components of F_{EX} are certain.

whether an unusually high temperature (relative to the known underlying carbon forcing and ocean effects) occurs because the feedback effects are stronger than expected, SGE is less effective than expected, or the weather shock was unusually high.

An assumption about learning here is that the average global temperature (contained in H' and T_{AT}) and forcing (contained in \hat{F}') are the only data that can be used to learn about either type of uncertainty. More generally, it may be possible to learn using other information, for example, regional temperature changes in response to SGE (Ricke et al., 2012). However, the use of regional data does not necessarily imply faster learning as the advantage of a larger panel data set is offset by spatial correlation, greater variability of regional shocks, and other issues.

2.5 Learning

The weather shock is normally distributed, $\epsilon' \sim N(0, \rho_\epsilon^{-1})$, where ρ_ϵ is the precision (inverse of the variance). The planner has prior beliefs that the true net feedback parameter is drawn from a normal distribution (Roe and Baker, 2007), and the same for the SGE effectiveness parameter:

$$\begin{bmatrix} \tilde{\beta} \\ \tilde{\phi} \end{bmatrix} \sim N[\mu, P^{-1}], \quad \mu \equiv \begin{bmatrix} \beta \\ \phi \end{bmatrix}, \quad P \equiv \begin{bmatrix} P_1 & P_2 \\ P_2 & P_3 \end{bmatrix}. \quad (13)$$

Here β and ϕ are the means of the prior distributions and P is the precision of the prior beliefs (the inverse of the variance-covariance). We assume P is positive definite, which is the case if the priors are formed using a variance-covariance matrix.

Bayes' theorem implies that, after observing H' , the posterior distribution remains bivariate normal, with mean and precision:

$$P' = P + \rho_\epsilon X \cdot X^{tr}, \quad (14)$$

$$\mu' = (P')^{-1} (P\mu + \rho_\epsilon X H') \quad (15)$$

$$X \equiv \begin{bmatrix} T_{AT} \\ -\hat{F}'g \end{bmatrix} \quad (16)$$

Here tr denotes the transpose operator. The planner can manipulate the second element of X directly through the SGE policy choice, g and indirectly through the abatement policy choice, α .¹⁴

¹⁴ \hat{F}' is a function of M' via equation (6) which is a function of emissions E via equation (5) and abatement α via

Thus, learning is endogenous. However, the first element of X , T_{AT} , is already determined at the start of period t and thus unaffected by the choice of SGE and abatement in period t . It is natural to assume $P_1 > 0$, $P_3 > 0$, and $P_2 \leq 0$ given that P_2 evolves over time as the sum of observations on $-T_{AT}\hat{F}g$, which is negative.

Finally, a critical feature of the model is that T'_{AT} and H' are sums of normal random variables and are therefore also normally distributed.

$$H' \sim N [\mu_H, P_H^{-1}] \quad (17)$$

$$\mu_H = X^{tr} \cdot \mu \quad (18)$$

$$P_H^{-1} = X^{tr} \cdot P^{-1} \cdot X + \rho_\epsilon^{-1} \quad (19)$$

Therefore, a single numerical integration with respect to the normal random variable H is sufficient for computing the expectation of the value function. More importantly, the variability of the climate as perceived by the planner (i.e. considering both stochasticity and uncertainty) is a function of SGE deployment. We will show that this magnification of variability caused by SGE can slow learning.

3 Theoretical Predictions

Before presenting the numerical simulation results, we present some theoretical propositions about learning that illuminate the model's predictions. Proofs of all propositions are in the online appendix.

While the specification of learning in the previous section is most clearly presented by using the precision matrix P , the predictions in this section are more conveniently expressed with the prior variance-covariance, W . Let:

$$W = P^{-1} \equiv \begin{bmatrix} W_1 & W_2 \\ W_2 & W_3 \end{bmatrix}. \quad (20)$$

Evaluating the posterior variance-covariance, W' , of the belief distribution over the uncertain vari-
equation (4).

ables $\tilde{\beta}$ and $\tilde{\phi}$ after observing H' reveals that:¹⁵

$$W' = \frac{1}{1 + \rho_\epsilon X^{tr} W X} \begin{bmatrix} W_1 + \rho_\epsilon |W| X_2^2 & W_2 - \rho_\epsilon |W| X_1 X_2 \\ W_2 - \rho_\epsilon |W| X_1 X_2 & W_3 + \rho_\epsilon |W| X_1^2 \end{bmatrix}. \quad (21)$$

SGE, g , directly affects $X_2 = -\hat{F}'g$, and abatement, α , indirectly affects X_2 . Therefore equation (21) shows both current-period policy variables affect the posterior variance of the estimates for both $\tilde{\beta}$ and $\tilde{\phi}$ through their effect on X_2 . By contrast, the policy choices g and α do not affect $X_1 = T_{AT}$, which is already determined at the start of the period.

We want to identify conditions under which SGE use unambiguously increases information, and likewise for the use of abatement. To do so, we first establish conditions under which the terms in the variance-covariance matrix (21) are decreasing over time. This result also creates natural bounds for the variance states in the computational algorithm and signs the elements of P and W for the later propositions. Replace the prime notation W and W' with W_t and W_{t+1} subscripts, respectively, where W_t represents the variance-covariance matrix of beliefs at time t , and W_0 represents it at the initial period. Let $W_{i,t}$ represent the value of the i th element of W_t . We restrict the prior variance-covariance matrix W_0 to be positive-definite. Starting from a diffuse prior, $P_0 = 0$, (as would be the case in an OLS regression), equation (14) implies that P_t and therefore W_t are positive-definite for all t . Without this restriction, it is possible to construct examples where the variances can become negative or unbounded. With this assumption, we can show:

PROPOSITION 1 *Let $W_{1,0} > 0$, $W_{2,0} > 0$, $W_{3,0} > 0$, and $|W_0| > 0$. Then: $0 \leq W_{1,t} \leq W_{1,0}$, $0 \leq W_{3,t} \leq W_{3,0}$, $0 \leq W_{2,t} \leq (W_{1,0}W_{3,0})^{\frac{1}{2}}$ and $|W_t| > 0$ for all finite t . Further, $W_{i,t+1} \leq W_{i,t}$ for $i = 1, 3$. That is, the planner gains information after each observation.*

Proposition 1 shows the variance-covariance matrix terms remain on finite intervals, which is necessary for convergence of the computational algorithm. More importantly, Proposition 1 shows that, regardless of the decisions, the planner gains information each period ($W_{1,t+1} < W_{1,t}$ and $W_{3,t+1} < W_{3,t}$).

¹⁵See equation (1.49) in the online appendix.

3.1 Learning from SGE

Although the planner gains information with each observation over time, the policy decisions affect the amount of information gained. In fact, whether an increase in SGE speeds or slows learning about either of the two unknown variables $\tilde{\beta}$ and $\tilde{\phi}$ depends on the level of SGE relative to the current information state.¹⁶ We will show below that there are two threshold values of SGE, defined as I_1 and I_2 , that determine the direction of the effect on SGE on learning.

First, consider the effect of SGE on the posterior variance of belief over the SGE effectiveness parameter ϕ . This variance in period $t + 1$ is $W_{3,t+1}$, which we will now denote as W'_3 . An increase in SGE speeds learning about SGE effectiveness if the partial derivative of W'_3 with respect to g is negative. We show in Appendix section 2 that:

$$\frac{\partial W'_3}{\partial g} = \frac{\hat{F}\rho_\epsilon 2W'_3}{1 + \rho_\epsilon X^{tr}WX} (W_2 T_{AT} - W_3 \hat{F}'g) \quad (22)$$

$$\frac{\partial W'_3}{\partial g} < 0, \Leftrightarrow g > I_1 \equiv \frac{W_2}{W_3} \cdot \frac{T_{AT}}{\hat{F}'} \quad (23)$$

SGE g increases learning about SGE effectiveness if and only if the level of SGE exceeds the threshold value I_1 . To understand the intuition, it is helpful to view equation (11) ($H' = \tilde{\beta}T_{AT} - \tilde{\phi}\hat{F}'g + \epsilon'$) as an OLS regression equation.¹⁷ Suppose a given set of past decisions results in a data set which is used to estimate β and ϕ using equation (11), with associated variance-covariance matrix (20). Now the planner has the opportunity to alter the current SGE and thus the current X_2 ($-\hat{F}'g$), and then re-run the regression with the past and current data. How does changing g affect the variance of the parameter estimates?

Suppose SGE begins from a low baseline close to zero, or even at zero.¹⁸ At the baseline, little SGE is being used, and so changes in temperature are more easily attributed to changes in the previous temperature through the feedback parameter β . A warm year is unlikely to be caused by ineffective SGE, since SGE is minimal, but very likely to be caused by strong feedbacks. The observation thus provides relatively more information about the feedback parameter β . Further, if $W_2 > 0$ then being more certain about β means that all prior information about SGE effectiveness is

¹⁶Compare this to Quaas et al. (2017), which exogenously and unambiguously assumes SGE R&D increases learning over climate sensitivity (p. 10 in their paper).

¹⁷In fact, it is straightforward to show that (21) is identical to the variance-covariance matrix of the OLS parameter estimates with known σ_ϵ^2 , for appropriately chosen priors $P_1 = \sigma_\epsilon^2 \sum_{i=0}^{t-1} X_{1,i}$ and analogously for P_2 and P_3 .

¹⁸Even at zero SGE ($g = 0$), the distribution of beliefs is well-defined, as is the derivative given in equation (22). To see this, replace X_2 with zero in equations (21) and (22).

more useful, as prior information on temperature and SGE is more easily correlated given feedback effects are now more certain.¹⁹ Note that $I_1 = 0$ (any SGE speeds learning about effectiveness) if the prior is diffuse and no prior information about effectiveness exists to improve upon, $W_3 = \infty$.

Now suppose SGE is increased from the low/zero baseline. The increase in SGE has a negative effect on temperature which cancels with the positive feedback effect (βX_1). The resulting signal H' is smaller relative to the noisy weather shock. The planner has low confidence that the temperature change resulted from SGE (and also low confidence the temperature change resulted from feedbacks). Hence, as SGE increases, the signal becomes smaller, the new information is less informative, and W'_3 changes less.

Suppose instead that SGE is increased from a high baseline. In this case, the planner expects a large temperature reduction due to SGE, and further increases in SGE start to make the temperature signal H' larger in magnitude (more negative). It is then increasingly unlikely that the temperature change is the result of a large negative weather shock. The planner attributes the change in temperature to SGE with high confidence, and thus the data point is highly informative, and the variance falls as SGE rises.

We denote the nonlinear effect of SGE on H' as the *signal strength effect*. The signal strength effect is quadratic in g : an increase in g from a small baseline cancels with the feedback effects weakening the signal. Increasing g from a large baseline (more than the threshold I_1) swamps the feedback effect, strengthening the signal.²⁰ Equation (22) shows that increasing SGE strengthens the signal for a baseline $g > I_1$. The cutoff I_1 depends on temperature since if the temperature is larger, so is the feedback effect which must be overcome. Conversely, if forcing is large, then SGE has a large effect on temperature which more easily swamps the feedback effect. The cutoff I_1 also depends on the pre-existing uncertainty, for example, if W_2 is small it is easier to separate the feedback and SGE effects, and so a smaller temperature change is more easily attributed to the change in SGE.

Next, consider the effect of SGE on the posterior variance W'_1 , which is the variance of beliefs

¹⁹Note that aerosols create uncertain feedbacks, especially through effects on cloud formation Forster et al. (2021), creating a correlation between the priors (W_2). It follows that a more precise estimate of CO₂ feedbacks improves prior information about SGE effectiveness.

²⁰Mathematically, in (21), $1 + \rho_\epsilon X^{tr} W X$ is quadratic in g , and so g has a quadratic effect on W' .

about the net feedback parameter β . We show in Appendix section 2 that:

$$\frac{\partial W'_1}{\partial g} = \frac{\hat{F}\rho_\epsilon 2W'_2}{1 + \rho_\epsilon X^{tr}WX}(W_1 T_{AT} - W_2 \hat{F}'g), \quad (24)$$

$$\frac{\partial W'_1}{\partial g} < 0, \Leftrightarrow g > I_2 \equiv \frac{W_1}{W_2} \cdot \frac{T_{AT}}{\hat{F}'}. \quad (25)$$

SGE g increases learning about the net feedback parameter if and only if the level of SGE exceeds the threshold value I_2 . SGE affects the signal strength H' in an identical, quadratic way. The signal strength effect on learning about feedbacks is proportional to the signal strength effect on learning about SGE effectiveness. However, SGE has an additional effect on learning about feedbacks, which we call the *noise amplification effect*. Increasing SGE from any baseline increases the subjective noise in the climate system when trying to estimate the feedback effect. In particular, the effect of feedbacks on temperature is now obscured not only by the noisy weather shock, but also by an uncertain SGE effect. The planner is now uncertain if (say) a warm year was caused by stronger than expected feedbacks, weather noise, or SGE being less effective than expected. This noise amplification effect slows learning about feedback effects, all other things equal.²¹

The appendix shows $I_2 > I_1$. That is, the noise amplification effect implies the baseline g must be larger for SGE to speed learning about feedbacks relative to SGE effectiveness, as the signal strength effect must not only be positive, but also overcome the noise amplification effect.

The following proposition summarizes the above discussion on the role of SGE on learning:

PROPOSITION 2 *Let $W_{1,0} > 0$, $W_{2,0} > 0$, $W_{3,0} > 0$, and $|W_0| > 0$. Then $I_1 \leq I_2$ for all finite $t \geq 0$ and up to three regions exist:*

- *Region 1: $g < I_1$. In this region an increase in SGE increases posterior uncertainty for both $\tilde{\beta}$ and $\tilde{\phi}$.*
- *Region 2: $I_1 \leq g < I_2$. In this region an increase in SGE decreases uncertainty over $\tilde{\phi}$ but increases uncertainty over $\tilde{\beta}$.²²*
- *Region 3: $g \geq I_2$. In this region an increase in SGE decreases uncertainty for both $\tilde{\beta}$ and $\tilde{\phi}$.*

²¹Mathematically, the noise amplification effect is the second term in the numerator of W'_1 given in equation (22).

²²Region 2 can be further decomposed into two subregions where SGE has differential effects on the covariance W'_2 . See the appendix.

Proposition 2 shows that sufficiently small SGE experiments slow learning, whereas sufficiently large SGE experiments speed learning. To understand the magnitude of the SGE experiment required to speed learning, we can use the calibrated initial values to recover the current period I_1 and I_2 . Details on the calibration of the initial conditions of the model are presented in Section 4 and Appendix section 4. Given those calibrated parameters, we calculate the initial values of the variables defining the cutoff regions: $I_1 = 7.65$, and $I_2 = 17.85$.²³ Hence only for $g > I_1 = 7.65$ will experimentation with SGE speed learning about the net feedback parameter. Since g is normalized so that $g = 1$ corresponds to about 5MT of sulfur, to speed learning would require a massive experiment of about 38.25MT of sulfur.²⁴ Thus, for any reasonably-sized implementation of SGE experiment, on the scale of what is being proposed by policy makers and IAMs, SGE will slow learning about both the net feedback parameter, $\tilde{\beta}$, and the effectiveness of SGE, ϕ . There is no justification for small-scale, planet-wide SGE experimentation based on reducing uncertainties, at least at the current calibrated conditions. Later, in the numerical simulation results, we explore whether this result holds over time as conditions change.²⁵

3.2 Learning from Abatement

Consider now the effect of abatement on learning. The next proposition shows that abatement policy, α and SGE policy, g have opposite effects on information gain – whenever more SGE decreases posterior uncertainty over an unknown variable, more abatement increases uncertainty over that variable. That is, whenever $\frac{\partial W'_i}{\partial g} < 0$, we will have $\frac{\partial W'_i}{\partial \alpha} > 0$.

PROPOSITION 3 *Let $W_{1,0} > 0$, $W_{2,0} > 0$, $W_{3,0} > 0$, and $|W_0| > 0$. Then*

- *Region 1: $g < I_1$. In this region an increase in abatement decreases uncertainty for both $\tilde{\beta}$ and ϕ .*

²³This comes from the fact that $\frac{T}{F'} = \frac{0.85}{0.54} = 1.56$, and $W_1 = 0.13^2$, $W_2 = 0.0385^2$, and $W_3 = 0.0174^2$. Thus $\frac{W_2}{W_3} = 4.89$ and $\frac{W_1}{W_2} = 11.41$.

²⁴For reference, the Mount Pinatubo eruption in 1982 injected an estimated 20MT of sulfur and significantly cooled global temperatures.

²⁵We also consider alternative calibrations for the initial conditions. First, we use the calibration based on Heutel et al. (2018), which starts in an earlier period with a lower temperature $T_0 = 0.8$ and a smaller initial forcing of $\hat{F}' = 0.21$. Both changes increase the thresholds so that $I_1 = 18.35$ and $I_2 = 42.82$. Second, we begin with a higher initial state of atmospheric carbon concentrations, corresponding to about a 600 ppm (1278 GTC) world, yielding values of $I_1 = 3.76$ and $I_2 = 8.78$. Both alternative calibrations show how the threshold level of SGE for which SGE speeds learning decreases with atmospheric CO₂ concentrations. However, the alternative parameters are still such that any reasonably-sized implementation of SGE will slow learning.

- *Region 2: $I_1 \leq g < I_2$. In this region an increase in abatement decreases uncertainty over $\tilde{\beta}$ but increases uncertainty over $\tilde{\phi}$.*
- *Region 3: $g \geq I_2$. In this region an increase in abatement increases uncertainty for both $\tilde{\beta}$ and $\tilde{\phi}$.*

Increasing SGE decreases X_2 (makes $X_2 = -\hat{F}'g$ more negative), while increasing abatement increases X_2 (brings it closer to zero).

Increasing abatement reduces the amount of forcing available for a given level of SGE to act on. SGE becomes less efficient in this sense. Therefore, if SGE is high enough so that SGE speeds learning ($g > I_1$), then abatement weakens the positive signal strength effect of SGE, slowing learning. Conversely, if SGE is low enough so that SGE slows learning ($g < I_1$), then abatement weakens the negative signal strength effect of SGE, speeding learning.²⁶ It is useful to contrast Proposition 3 with the results of Kelly and Kolstad (1999). That paper had only one uncertainty, the forcing parameter, $\xi_1\eta$. From equations (9), (11), and (16), uncertainty about the forcing parameter and uncertainty about SGE effectiveness are both equivalent to uncertainty about the net forcing X_2 . (More gross forcing is equivalent to more net forcing after SGE.) In Kelly and Kolstad (1999), abatement always slowed learning. Here, we have uncertainty about the net feedback parameter in addition to uncertainty about forcing through SGE. Suppose $W_1 = W_2 = 0$, so that only net forcing is uncertain. Then it is straightforward to show that $W'_1 = W'_2 = 0$. Since the feedback effect is certain, no further learning takes place about it regardless of abatement. Further, $I_1 = 0$, so region 2 holds for all g . In this region, an increase in abatement α slows learning about SGE effectiveness or equivalently net forcing. Thus, the result of Kelly and Kolstad (1999) emerges as a special case when the prior uncertainty over the net feedback parameter is sufficiently small. When the net feedback parameter uncertainty is large, abatement may speed learning.

Quantitatively, experimentation with abatement is less fruitful than experimentation with SGE. Changes in abatement have only small effects on the current stock of CO_2 , which has built up over centuries. Therefore, changes in abatement have little effect on X_2 . Consider again the calibrated initial condition described in the previous section. If $\alpha = 0$, then, from equation (9), $\hat{F}' = 0.5496$, and if $\alpha = 1$, $\hat{F}' = 0.5333$. Thus, I_1 varies between 7.56 and 7.79, so the planner does not create

²⁶Note that abatement reduces forcing by a known amount contained in H' in equation (17), which does not matter for learning. Abatement has a second effect which is the unknown reduction in the ability of SGE to reduce forcing, X_2 , which is what matters for learning.

additional regions by changing α . Consider the effect of increasing abatement α from 0 to 1 for a given value of g . For example, when $g = 0.1$:

$$\frac{W'(g = 0.1, \alpha = 1) - W'(g = 0.1, \alpha = 0)}{W'(g = 0.1, \alpha = 0)} = \begin{bmatrix} -0.0171\% & -0.0394\% \\ -0.0394\% & -0.0168\% \end{bmatrix}, \quad (26)$$

With g in region 1, moving from $\alpha = 0$ to $\alpha = 1$ speeds learning as indicated by Proposition 3, but the effect is small: the variance terms decrease by less than two hundredth of one percent despite the large increase in abatement.

Thus, the justification for experimentation with abatement to reduce uncertainties is quantitatively small, though of course abatement yields other benefits. The numerical simulations below determine optimal abatement and SGE including all costs and benefits, informational and otherwise.

3.3 Value of Information

In region 2, SGE increases information over SGE effectiveness $\tilde{\phi}$ but decreases information over the net feedback parameter $\tilde{\beta}$. Ignore for the moment all of the other costs and benefits associated with an SGE experiment, and consider only the effect on uncertainty. Is the information gain worth the information loss? That is, is the value of information from SGE experimentation positive, and does it result in more accurate decision making and higher utility?

The concept of Blackwell informativeness answers this question. If an SGE experiment results in an uncertainty distribution such that temperature can be predicted more accurately than without such an experiment, then the more accurate prediction leads to higher utility. In this case, the SGE experiment is said to be more Blackwell informative, or has a positive value of information.²⁷

The next proposition demonstrates that for $g > I_1$, the information gain about SGE effectiveness is worth the information loss about the net feedback parameter, or that Blackwell informativeness is increasing in g .

PROPOSITION 4 *Let $W_{1,0} > 0$, $W_{2,0} > 0$, $W_{3,0} > 0$, and $|W_0| > 0$.*

1. *Consider two posterior distributions $\Phi_1 = N[\mu'(g_1), W'(g_1)]$ and $\Phi_2 = N[\mu'(g_2), W'(g_2)]$, produced from experiments g_1 and g_2 , with $g_2 > g_1$. Then experiment 2 is more Blackwell*

²⁷See for example de Oliveira (2018).

informative than experiment 1 for a given X if:

$$\frac{\left(\mu'(g_2)^{tr} X\right)^2}{X^{tr} W'(g_2) X + \rho_\epsilon^{-1}} > \frac{\left(\mu'(g_1)^{tr} X\right)^2}{X^{tr} W'(g_1) X + \rho_\epsilon^{-1}}. \quad (27)$$

2. Consider two posterior distributions $\Phi_1 = N[\mu', W'(g)]$ and $\Phi_2 = N[\mu', W'(g + \Delta g)]$, where $\Delta g > 0$. Then for Δg sufficiently small, experiment 2 is more Blackwell informative than experiment 1 for a given X if $g > I_1$.

Consider an experiment of increasing SGE policy, g . Given the normally-distributed prior uncertainty and weather shocks, the problem can be condensed into a single random variable (H' or T'_{AT}).²⁸ Thus, the posterior distribution is more Blackwell informative if it produces a more accurate prediction of T'' in the next period. Further, a more accurate prediction necessarily results in a more informed decision and therefore higher utility. An increase in SGE makes $X_2 (= -\hat{F}'g)$ more negative, which affects W' through equation (21). Considering this effect only, the second part of Proposition 4 shows that the information gain regarding $\tilde{\phi}$ is more valuable than the information loss regarding $\tilde{\beta}$ in region 2. Therefore, in regions 2 and 3, SGE experimentation on net produces valuable information that leads to more informed decisions.²⁹

A change in SGE also affects the mean of the posterior distribution.³⁰ The first part of Proposition 4 shows that the overall effect on information depends on how SGE affects both the mean and variance of the prior. We can rewrite equation (27) as:³¹

$$\frac{\mu_H(g_2)^2}{\text{var}_H(g_2)} > \frac{\mu_H(g_1)^2}{\text{var}_H(g_1)}. \quad (28)$$

Hence, increasing SGE is more Blackwell informative if doing so increases the signal-to-noise ratio. SGE experimentation can produce a stronger or a weaker signal. Indeed, since μ' depends on the random variable ϵ' , the effect of experimentation on the signal varies across periods. Nonetheless, from Bayes' theorem, $E[\mu'] = \mu$ regardless of X , so in expectation experimentation has no effect on the signal. Thus, the effect of experimentation on the noise shown in part 2 of Proposition 4 is

²⁸See equation (17).

²⁹The information value of SGE experimentation must be weighed against the cost and damage of implementing SGE, and the benefit in terms of less forcing. This is addressed in the next section.

³⁰See equation (15)

³¹From Appendix equations (1.48) and (1.55), the signal for T' is $E[T'] = \hat{F}' - \xi_1 \xi_3 (T_{AT} - T_{LO})(t) + \mu^{tr} X$, of which only the last term matters since only the last term depends on g . Further, Appendix equation (1.56) indicates the noise with respect to T' is var_H .

likely to be quantitatively more important. Finally, the value of X matters and the informativeness of experimentation varies over time.

4 Calibration

We now turn to the computational simulations, beginning with the calibration of the model. Appendix section 4 gives the details of our calibration, except for the calibration of SGE effectiveness and short run temperature response which are discussed here. Notably, we use a novel strategy to calibrate SGE effectiveness and uncertainty using data on aerosol optical depth (AOD) from historical volcanic eruptions.

4.1 SGE Effectiveness

4.1.1 Data

We use two sources of information. For volcanoes, our independent variable, we use VolcanEESM: Global volcanic sulphur dioxide (SO₂) emissions database from 1850 to present - Version 1.0.³² We use the Volcanic Explosivity Index (VEI) of each emission based on the Global Volcanism Program Database and select the volcanoes with $VEI \geq 4$ which have the capacity to reach the stratosphere. We also have the total emission of SO₂ in TgS that we use to approximate the amount in the stratosphere in each period. The dependent variable is the global and hemispheric mean Aerosol optical depth at 550 nm, collected from NASA.³³ The frequency of data is monthly between 1850 and 2012.

Figure 1 plots the data. The top panel shows each volcanic eruption date and SO₂ emissions, while the bottom panel shows monthly AOD. After each eruption there is a clear increase in AOD, which we match to the amount of sulfur associated to each volcano to the whole year following the explosion.

³²This dataset is available at <http://dx.doi.org/10.5285/76ebdc0b-0eed-4f70-b89e-55e606bcd568>.

³³These data are available at <https://data.giss.nasa.gov/modelforce/strataer/>.

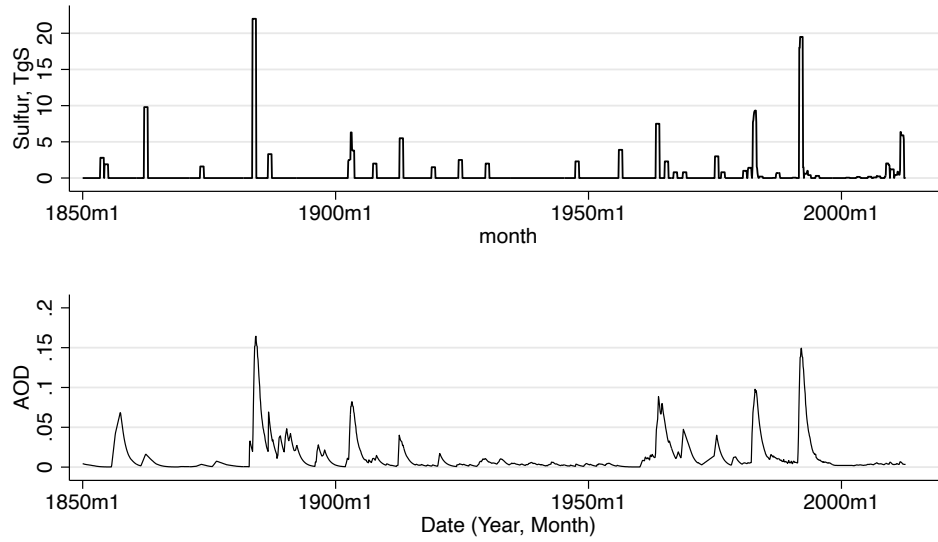


Figure 1: Volcanic Eruptions, SO_2 emissions, and Aerosol Optical Depth

4.1.2 Methodology

The timing of the model is such that sulfur injected into the stratosphere in year t causes temperatures to fall in year $t + 1$, after which sulfur must be re-injected (that is, g_t is a flow variable). Thus, we mechanically assume the effect of the volcano ends one year after the eruption.³⁴ The model does not take a stand on when within the current year the injection takes place, so we are free to assume the sulfur enters the stratosphere $z < 12$ months after the eruption. We are thus calibrating a discrete model to what is in fact a continuous process whereby a volcanic eruption for which stratospheric sulfur levels slowly build up and then decay in a few months to a year or more.

After choosing z , we step-wise interpolate the sulfur data from the eruption in month/year t such that the stratospheric sulfur level equals the recorded emissions in the data in months $t + z$ to $t + 12$. Figure 2 shows the interpolated data for Mt Pinatubo for several values of z . The top left panel shows the interpolation for $z = 0$ so the interpolated data has a sulfur level of approximately 18Tg immediately after the eruption date and remains at this level for one year. The top right panel shows the interpolation assuming the effect starts $z = 3$ months after the eruption; the bottom left shows 6 months after the eruption, and the bottom right 11 months after. All end 12 months after the

³⁴While this assumption is purely for simulation and calibration purposes, the sharp cut-off could be achieved technically using self-levitating particles (Keith, 2010).

eruption, so that the interpolated stratospheric sulfur level is zero for any month that has not seen an eruption in the previous $12 - z$ months. The regression coefficient is the average effect over all months. We can see from this example that lower values of z result in regression coefficients which underestimates the effect of solar geoengineering while starting later can, in principle, overestimate the effect.

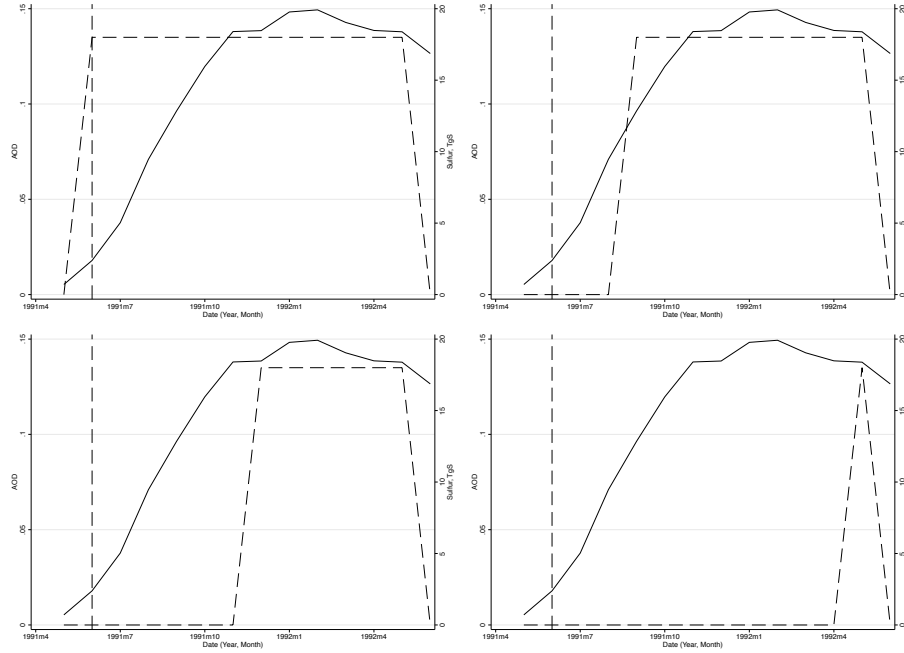


Figure 2: Interpolated data for Mt Pinatubo eruption. The dotted line is the eruption date, the rectangular line is the interpolated data (right scale), and the curve is the AOD (left scale).

For calibrating the prior distribution of uncertainty over SGE effectiveness, ϕ , we first estimate the AOD response to stratospheric sulphur dioxide injections. The estimating equation is

$$AOD_t = \zeta_1 + \zeta_2 \text{sulfur}_t + \epsilon_t, \quad (29)$$

where sulfur_t is the interpolated monthly volcanic eruption data, ζ_1 is a constant, and ζ_2 is the parameter of interest. In addition, ϵ_t is the error term at year/month t . To account for the possibility of autocorrelation in the errors, we use the Newey estimator in Stata that returns the Newey-West standard errors. Calibrating the effectiveness of SGE using volcanic eruptions comes with at least three important caveats (Robock et al., 2013, explains that there are substantial differences between

SGE and volcanic eruption). First, most SGE proposals are designed to have a relatively even pole-to-pole radiative forcing while volcanoes are highly localized producing an asymmetric radiative forcing depending on the latitude of the volcano and the time of the eruption (Jones et al., 2016). Second, the response of the climate to a volcanic eruption is transient and expected to be very different response to the relatively continuous application of aerosols that would be associated with solar geoengineering (Haywood et al., 2013). Both of these caveats suggest we are underestimating the effectiveness of solar geoengineering using volcanoes. Third, we are using interpolated data since actual stratospheric sulfur levels are unobserved, which may mean our estimates understate the uncertainty with respect to SGE effectiveness.

Table 1 shows the estimated effects of sulfur from volcanic eruptions on AOD. The rows represent different values of time after the eruption until the sulfur reaches the stratosphere, z , while the columns represent errors autocorrelated up to a given lag. As expected from the examples above, we see an increase in the estimated effect as the window focuses on the maximum impact within a year. The error terms, in turn, increase with the number of lags, but plateau at around 3-6 months.

z	Estimate	Error Lag				
		0	3	6	9	12
0	0.0041	0.0004	0.0007	0.0008	0.0008	0.0008
3	0.0051	0.0004	0.0006	0.0007	0.0007	0.0007
6	0.0059	0.0003	0.0005	0.0006	0.0006	0.0006
9	0.0063	0.0004	0.0006	0.0007	0.0007	0.0007
11	0.0064	0.0007	0.0008	0.0008	0.0008	0.0008

Table 1: Regression results for different assumptions about the time from eruption until the sulfur reaches the stratosphere and different error lags. Each cell is a different regression with the estimates of ζ_2 given in the second column and the standard errors of ζ_2 given in columns 3-7. The units of ζ_2 are AOD per TgS.

For our calibration, we decided to use a middle of the road assumption that the sulfur enters the stratosphere 6 months after the eruption and an error structure with 3 lags. This assumption also matches the estimated effect of maximum dispersion across the stratosphere, but minimizes the time of achieving that maximum effect. The result is $\zeta_2 = 0.0059 \pm 0.0005$. The results correspond to a two standard deviation confidence interval of approximately plus or minus 10.2% of the mean

estimate. Each eruption generates a very strong effect on AOD relative to the years without an eruption, which allows for a more precise estimate.

We use the estimated effect of sulfur on AOD to predict the effect of sulfur on radiative forcing. Following Ricke (2011), we assume a change of 0.1 AOD is equivalent to a reduction in radiative forcing of 2.5 Wm^{-2} . This implies based on our results that 1 TgS (= 1 MT sulfur) reduces radiative forcing by $0.15 \pm 0.013 \text{ Wm}^{-2}$. In the estimation above we assume the climate effects of the volcano dissipate after one year. Our calibration suggests a 20TgS (Pinatubo) eruption reduces forcing by 3 Wm^{-2} . This is larger than the usually assigned value of 2 Wm^{-2} , but follows from our assumption of a more effective date the sulfur enters the stratosphere of six months after eruption that captures most of the peak effect of these volcanic eruptions. Using a window that starts at zero months after the eruptions, we would find the effect of a Pinatubo-like explosion to be 2.05 Wm^{-2} , very much aligned with the estimates in the literature. Our calibration is aligned with an engineered approach to SGE that increases the effectiveness of SO_2 relative to a natural explosion.

The calibration of the damage parameter π_g follows Heutel et al. (2018) in that $g = 1$ is normalized such that 5MT sulfur causes a decrease in radiative forcing of $\phi\eta_f$ when carbon concentrations are twice preindustrial. Since $\zeta_2 = 0.15$ has units of forcing per MT of sulfur, we multiply by $5/\eta_f$ to get $\phi = 0.2 \pm 0.0174 \text{ CO}_2$ equivalent forcing per 5 MT sulfur.

Finally, we consider the calibration of the initial covariance, $W_{2,0}$. The covariance comes from the passive learning model, in which $x_{1,t}$, $x_{2,t}$ and H_{t+1} are observed, generating new information to update the parameter estimates (see equations 11 and 21). However, the priors are actually formed from two different sources. The feedback prior comes from GCMs (Roe and Baker, 2007) and the effectiveness prior comes from the volcano/optical depth regressions above. One way to resolve this is to generate a prior for the covariance as if the other priors were generated from a regression using historical data on x_1 and x_2 via equation (11). If so, then:

$$W_0 = \rho \begin{bmatrix} \sum_{t=0}^{2015} x_{1,t}^2 & \sum_{t=0}^{2015} x_{1,t}x_{2,t} \\ \sum_{t=0}^{2015} x_{1,t}x_{2,t} & \sum_{t=0}^{2015} x_{2,t}^2 \end{bmatrix} = \rho \begin{bmatrix} 0.13^2 & W_{2,0} \\ W_{2,0} & 0.0174^2 \end{bmatrix}. \quad (30)$$

Note that $x_{2,t} = 0$ except for the $m = 60$ years with volcanic eruptions noted above and $x_{1,t}$ is the $n = 140$ year temperature record. Under the further simplifying assumption that x_1 and x_2 are

constant when not zero,³⁵ it follows that:

$$W_{2,0} = \sqrt{(m/n) W_{1,0} W_{3,0}} = 0.0385^2. \quad (31)$$

Hence if $W_{1,0}$ and $W_{3,0}$ were generated from historical data consisting of a small number of years with volcanic eruptions and a larger number of years, without an eruption, then $W_{2,0} = 0.0385^2$.

4.2 Initial Feedback Uncertainty and Calibration of Temperature Speed of Adjustment

For the initial uncertainty, Roe and Baker (2007) find that a normal prior distribution of $\tilde{\beta}$ with prior mean $\beta_0 = 0.65$ and variance $W_{1,0} = 0.0169$ implies a climate sensitivity distribution which matches the distribution of climate sensitivities from various global circulation models (GCMs). We adopt these priors as initial conditions. Following the arguments of Weitzman (2009) and Roe and Baker (2007), we assume the net feedback parameter is uncertain but the forcing effect of CO₂ (η) is known with certainty.³⁶ We thus follow DICE-2016 and calibrate $\eta = 3.613$.

It remains to calibrate the parameter ξ_1 , which governs the response of temperature to radiative forcing from CO₂ in equation (7). We would like the temperature change to reflect certain stylized facts. In particular (1): the best estimate of the climate sensitivity is approximately 3°C (Solomon et al., 2007), (2) temperatures respond quickly to changes in CO₂ concentrations (Dietz et al., 2020), and (3) the short run feedbacks generate a temperature increase as in Dietz et al. (2020), and the temperature increases in the medium term as long run feedbacks begin to take hold (Dietz et al., 2020). Typically, GCMs have both short and long term feedbacks, which allows for temperature responses that feature (1)-(4). In contrast, almost all IAMs including this paper have only a single short run feedback effect, β . The calibration thus cannot replicate all of these effects and some compromises must be made (see Conte and Kelly, 2021, for a discussion of short and long run feedbacks and fat tails).

Consider first the climate sensitivity, equation (12). Given the above calibration of η and the prior uncertainty of $\tilde{\beta}$, the calibration of ξ_1 determines the prior beliefs about the climate sensitivity

³⁵This assumption can be relaxed and the result is quantitatively similar in that I_1 will still be well above one and the optimal policy is largely independent of $W_{2,0}$ in any event.

³⁶Still, some debate exists in the literature here. For example, Kelly and Kolstad (1999) assume the forcing effect of CO₂ is unknown.

through equation (12). In particular, a value of $\xi_1 = 0.2852$ results in the best estimate value of $\Delta\bar{T}_{2\times} = 3$. However, this value substantially exceeds the value given in Nordhaus (2016) of $\xi_1 = 0.1005$. In addition, such a large value of ξ_1 is inconsistent with recent temperature fluctuations in the sense that imposing $\xi_1 = 0.2852$ causes the temperature to increase discretely in the first period (see equations 6 and 7). The problem is that the best guess GCM value of the climate sensitivity reflects in large part long run feedback effects, whereas the model has only short run feedback effects. Therefore, to match the long run climate sensitivity requires an unrealistic short run response of temperature to forcing.

Next, consider the short run response of temperature to radiative forcing, illustrated by the short run response of temperature to a 100GT carbon impulse given in Dietz et al. (2020). Given our best guess initial feedback parameter, $\beta_0 = 0.65$, the model fits the relatively fast response of temperature to the 100GT carbon impulse well, getting to within 0.01 of the peak in 16 years, as compared to the peak of the best fit of GCMs which occurs after 14 years. However, to match the magnitude of the impulse response, a maximum temperature difference of 0.225°C , requires imposing $\xi_1 = 0.0742$. This implies a climate sensitivity of only 0.78°C . Because no separate long run feedback mechanism exists in the model, matching exactly the short run feedbacks implies that the climate sensitivity is too low because the long run feedbacks are absent. Conversely, assuming $\xi_1 = 0.2852$ matches the 3°C climate sensitivity but implies an unrealistically large short run temperature response to a 100GT carbon impulse equal to 0.87°C . As a compromise, we follow Kelly and Tan (2015) and impose $\xi_1 = 0.2624$. This implies that if $\tilde{\beta} = \beta_0 = 0.65$ that (1) the climate sensitivity is 2.76°C , close to the IPCC best estimate, (2) the short run temperature response peaks in 16 years, close to the best fit of about 14 years, (3) the magnitude of the short run response is 0.77°C , which exceeds the best fit of GCMs. Although not perfect, our view is that this calibration is an improvement over the DICE 2016 calibration, which has $\xi_1 = 0.1005$ and $\beta_0 = 0.8807$, which has a realistic climate sensitivity of 3.1°C , but is a poor match for the short run impulse response (Dietz et al., 2020), which given discounting is likely more important (Roe and Bauman, 2013).

5 Simulation Results

Propositions 2 and 3 in Section 3 give conditions under which deployment of either SGE or abatement increases the speed of learning over climate sensitivity and/or SGE effectiveness. Given initial

conditions, we can say whether there is information value in experimentation with either policy instrument. The theoretical results, however, do not give the optimal paths for both policies, taking into account their informational values as well as all other costs and benefits. To answer that question, we turn to a set of numerical simulations.

The computational solution method is provided in Appendix section 3. We run 10,000 simulations of the solved model, each simulation lasting 185 years. We choose 185 years since by the end of that period (the year 2200), all policy values have roughly leveled off for all simulations. All of the certain parameter values are identical across each of the 10,000 simulations, and each simulation has the same prior distribution of beliefs about the two uncertain parameters in the initial year. For each simulation, we draw a new true value of $\tilde{\beta}$ and a new true value of $\tilde{\phi}$ from that prior distribution. These true values are unknown to the planner, though the planner learns endogenously over time. We also draw a new weather shock ϵ_{it} for each simulation i in each year t . We report the mean value of the policy variables across the 10,000 simulations. These results are the means conditional on the identical initial (period zero) prior across simulations. These mean values will differ from assuming $\tilde{\beta} = \beta_0$ and $\tilde{\phi} = \phi_0$, which would give only the values conditional on the unlikely event that the initial mean belief was exactly correct.

5.1 Optimal Policy

Figure 3 plots the mean optimal SGE for each period over the 10,000 simulations (line with the plus symbol). We also plot the mean SGE across just the simulations where the planner's current beliefs about the values of the uncertain parameters β_{it} and ϕ_{it} are in the top ten percentiles of the belief distribution (line with circles), and the mean SGE across just the simulations where the planner's current beliefs β_{it} and ϕ_{it} are in the bottom ten percentiles (line with squares).³⁷ The curve with circles thus represents optimal SGE policy conditional on the planner believing that SGE is very effective and climate sensitivity is very high, while the blue curve represents optimal SGE policy conditional on the planner believing that SGE is very ineffective and climate sensitivity is very low.

³⁷We require that both β_{it} and ϕ_{it} are in the top ten percentiles of their respective distributions. By definition, 1,000 of the 10,000 simulations have β_{it} in its top ten percentile, and likewise for ϕ_{it} . Given the prior beliefs imply the true $\tilde{\beta}$ and $\tilde{\phi}$ are positively correlated, in most simulations if β_{it} is in the top 10 percentile then so is ϕ_{it} . A similar result holds for the simulations of the bottom ten percentiles.

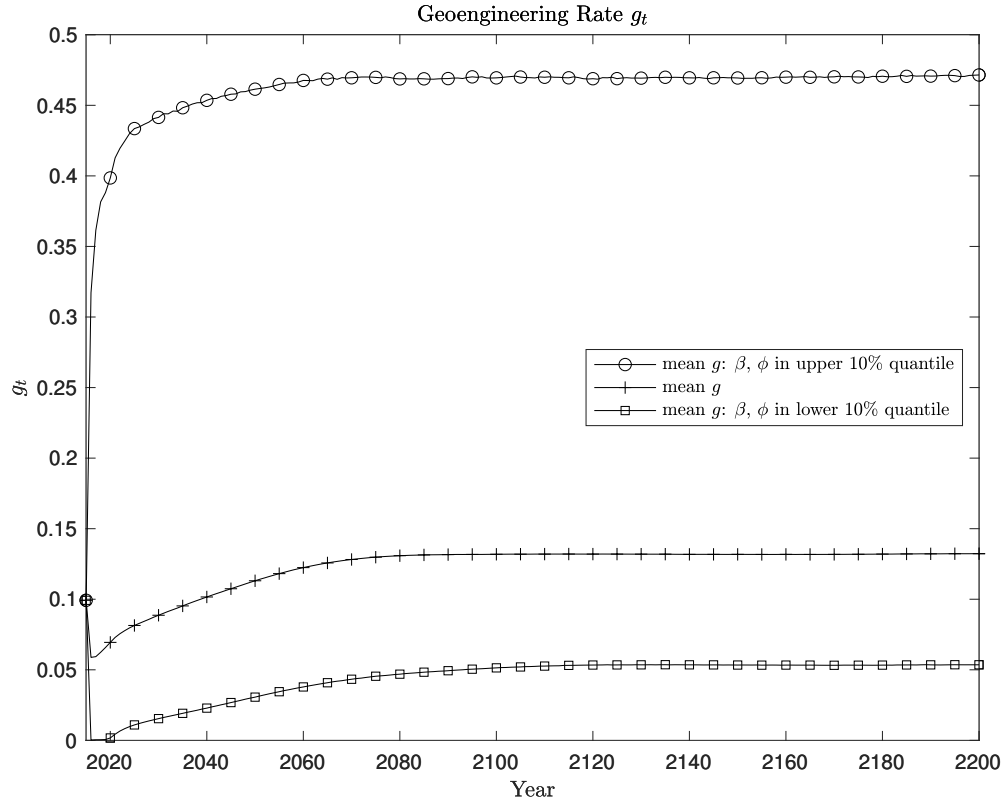


Figure 3: SGE policy, g_t , 10,000 simulations. Markers are every 5 years for clarity, but the data are annual in all figures.

For the overall mean, in 2015 the geoengineering rate g is 0.10 (offsetting 10% of radiative forcing) and then falls over the next nine years. We will show that initially the planner is concerned about the possibility that the climate sensitivity might be much higher than the initial prior. In most simulations, the planner is able to rule out these worst-case scenarios relatively quickly, and so SGE declines. Hence, initially, a moderate amount of SGE is optimal for precautionary reasons.

The upper and lower quantiles also begin in 2015 at about $g = 0.10$, since all simulations start with the same initial beliefs and thus the same SGE policy. However, when beliefs are that climate sensitivity and SGE effectiveness are high, optimal SGE rises, peaking at around $g = 0.46$ (offsetting 46% of radiative forcing). Conversely, when beliefs indicate SGE effectiveness and the climate sensitivity are likely to be low, the optimal SGE plummets to less than 0.01. Optimal SGE is thus highly sensitive to beliefs about the SGE effectiveness and climate sensitivity. SGE is largely a policy used quite sparingly, and only under extraordinary beliefs is SGE deployed at anything

more than about $g = 0.10$. However, for most beliefs in the first few years there is some SGE for precautionary purposes.

Over time, the mean SGE increases and then stabilizes around its peak as rising carbon concentrations and wealth makes SGE more attractive. The mean optimal SGE is closer to the 10% quantile than the 90% quantile. The climate sensitivity distribution is bounded below by zero but is unbounded above, and the prior distribution is asymmetric and fat-tailed (Kelly and Tan, 2015).

We next compare our results on optimal SGE policy with those of Heutel et al. (2018) (see their Figure 4f). In our model, the initial beliefs (prior mean) of the net feedback parameter and SGE effectiveness are $\beta_0 = 0.65$ and $\phi_0 = 0.201$, respectively. In their model, the mean value of the net feedback parameter is $\beta_0 = 0.87$ and SGE effectiveness is fixed at $\phi_0 = 0.5$. As a result, although in their model optimal geoengineering has a similar shape as in our model, SGE rises to about $g = 0.32$ before declining. In our model in contrast, SGE rises to only $g = 0.13$ and then declines more modestly. SGE is used less here relative to the prior literature for three reasons. First, our calibration has an initial belief that effectiveness is only 0.201. Second, SGE effectiveness is uncertain in our model. Third, SGE at these levels slows learning, leading to less informed decision-making. Even if we solve our model assuming the mean of the prior beliefs of SGE effectiveness is 0.5, SGE peaks at $g = 0.23$. Therefore, uncertain effectiveness and slower learning are quantitatively significant deterrents for SGE deployment, reducing the optimal SGE by about 28% relative to Heutel et al. (2018).

Figure 4 plots the mean optimal abatement policy α , as well as the mean abatement policy when beliefs about the net feedback parameter and SGE effectiveness are in the 10th and 90th quantiles.

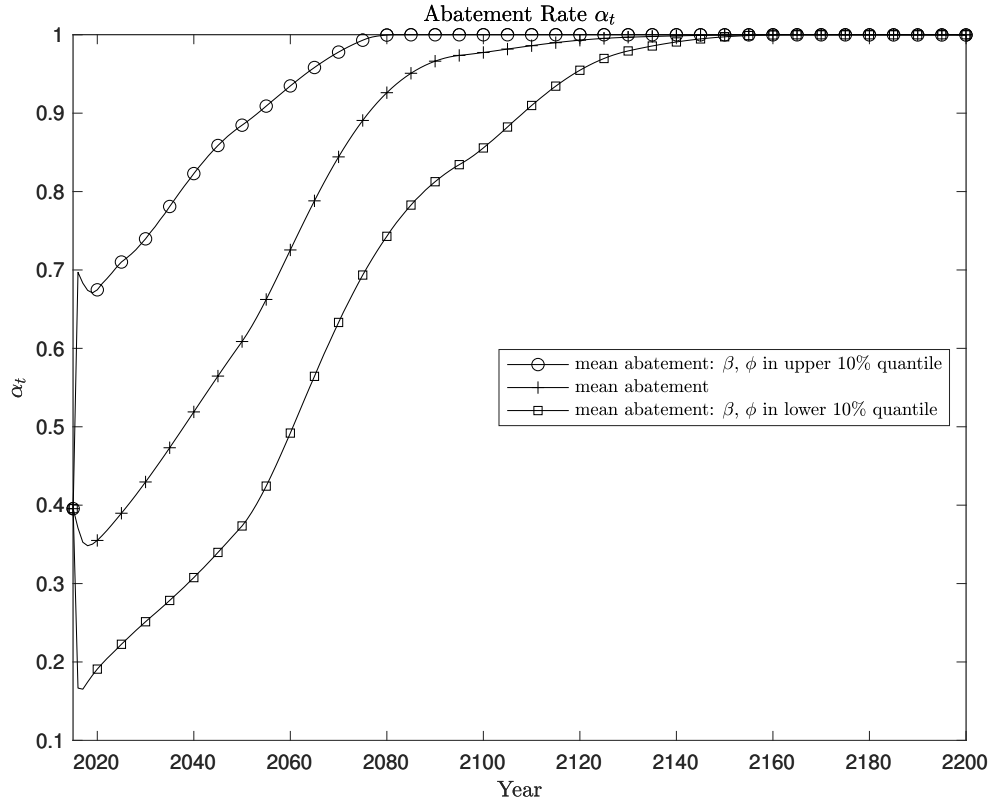


Figure 4: Abatement policy α_t , 10,000 simulations.

The initial abatement rate is about 40% in 2015. Like SGE, mean abatement is initially high for precautionary reasons. Mean optimal abatement subsequently falls to about 35% as learning rules out some of the worst climate sensitivity cases. However, after a few years mean abatement begins to rise again, as the cost of abatement falls to near zero and growth in capital, population, and productivity quickly raise emissions gross of abatement. Eventually abatement is so inexpensive that 100% abatement becomes optimal. This result is consistent with the prior DICE literature.

When beliefs about the climate sensitivity and SGE effectiveness are in the 90th percentile, abatement rises more quickly. Because beliefs are that the climate is very sensitive to carbon, the planner expects that emissions cause higher temperatures and more damage, so higher abatement is optimal. Thus, when beliefs about the climate sensitivity rise, the planner uses more of both SGE and abatement. SGE has a number of advantages, including (1) a linear effect on forcing versus the logarithmic effect of abatement, (2) an immediate, flow effect on forcing whereas abatement can only slowly lower the CO_2 stock, and (3) an initial cost advantage. Despite the advantages

of SGE in terms of cooling temperatures more quickly and cheaply than abatement, the efficient policy portfolio still includes a substantial role for abatement. This is because (1) SGE does not directly reduce damages from CO₂, (2) SGE creates damages itself, (3) SGE slows learning, and (5) the cost of abatement falls over time. Further, the advantages and disadvantages change over time. However, regardless of the current magnitude of each, the cost of both SGE and abatement are convex. Hence, it is less expensive to use some of each rather than one or the other exclusively. SGE does not eliminate the need for abatement, a concern expressed in some of the prior literature.³⁸

5.2 Effect of Learning on Optimal Policy

We explore how uncertainty and learning affect policy by comparing three sets of simulations. First, in the *learning* case, we present the optimal policy simulations under our full model where there is both uncertainty and endogenous learning over two parameters (these are the results presented earlier in Figures 3 and 4). Second, in the *certainty* case, we present analogous optimal policy simulations without uncertainty. That is, we again simulate optimal policy for 10,000 simulations, each of which lasting 185 years, where at the start of each simulation we draw a true value of each uncertain parameter β and ϕ from the prior distribution. But, the prior mean belief is set equal to this true value and the variance of the prior is set equal to zero, so that there is certainty over these values at the start and thus no learning.³⁹ Third, in the *no-learning* case, we present analogous optimal policy simulations under uncertainty but without learning. To do this, we solve the model such that the planner knows no learning will take place in the future (W is no longer a state, and beliefs over β and ϕ are fixed at the means of the prior distribution). We again simulate the optimal policy for 10,000 simulations, where at the start of each simulation we draw a true value for β and ϕ , unobserved by the planner. We average 10,000 such paths without learning, to get the expected path conditional on the prior information and no learning.

Figures 5 and 8 present the simulation results for optimal SGE policy. Figure 5 contrasts the certainty and learning simulations, and Figure 8 contrasts the learning and no-learning simulations. In each figure, we present the mean optimal policy, and the mean optimal policy when the feedback and SGE effectiveness beliefs are in the 10th and 90th percentile, both for the simulations with

³⁸In Appendix section 5, we show how the forecast errors over the uncertain variables evolve, which illuminates how the speed of learning differs across the variables.

³⁹However, a random weather shock is drawn in each period of each simulation, so that even in the certainty case future weather shocks are unknown.

learning (replicating Figure 3) and for the simulations under either certainty or no-learning.

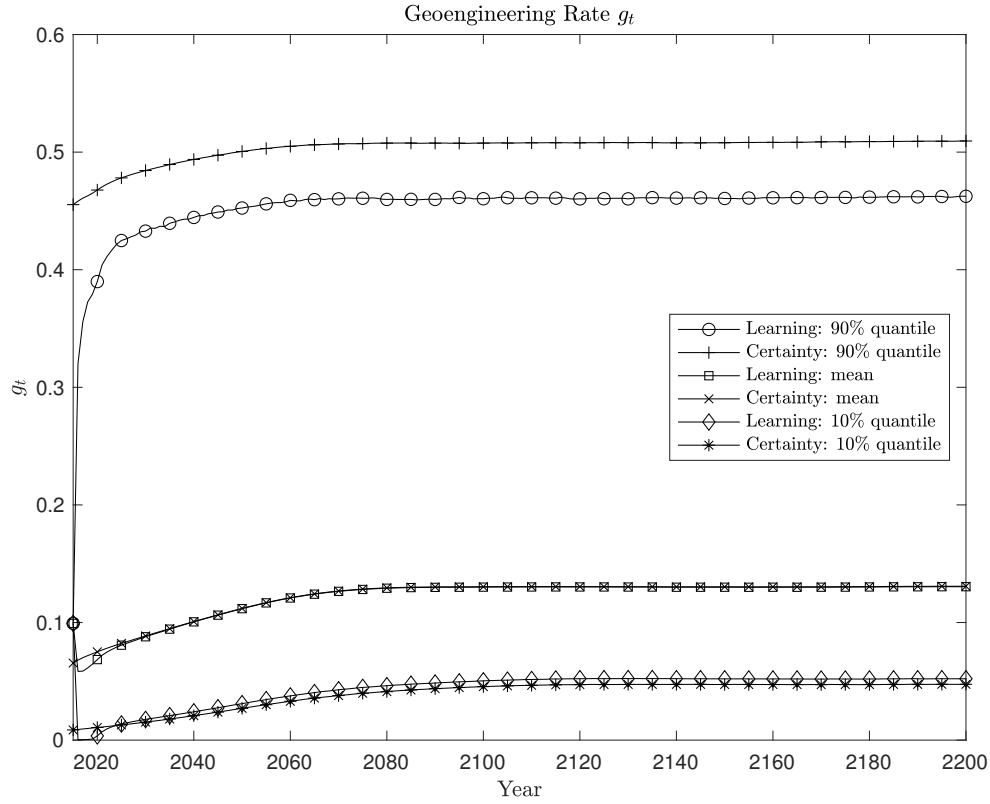


Figure 5: SGE policy, g_t , 10,000 simulations, learning and certainty cases.

Figure 5 shows that uncertainty matters initially. In 2015 the mean SGE policy is $g = 0.10$, yet if the planner knew the climate sensitivity and SGE effectiveness were in the 90% quantile, mean SGE policy would be $g = 0.46$ and if the planner knew the climate sensitivity and SGE effectiveness were in the 10% quantile, mean optimal SGE would be less than one percent ($g = 0.0087$). Since the prior distribution is normal, with Bayesian learning forecast errors are unbiased (see Figures A1-A3 in Appendix section 5). In the 90% quantile in Figure 5, the true values are higher than the prior, and so $\beta_t < \tilde{\beta}$ and $\phi_t < \tilde{\phi}$ and the planner is revising beliefs upward. Similarly, in the lower 10% cases the planner is revising estimates downward. However, SGE policy is much more sensitive to beliefs when beliefs are that the climate sensitivity is very high. In this case, small differences in feedbacks magnify over time into very large temperature changes, and so the planner makes a much bigger policy adjustment. Hence the difference between learning and certainty in

the 90% quantile in Figure 5 represent more a difference in policy sensitivity than a difference in forecast errors. Appendix Figures A1-A3 also show that after the first couple of decades further learning about SGE effectiveness and the climate sensitivity proceed slow enough so that there is little further convergence in the next 100 years.

Since in the 90% quantile beliefs increase from below the true value, the optimal policy given learning uses less SGE than is optimal given certainty, as the planner believes SGE is not as effective as it actually is. Conversely, when SGE is not very effective (10% quantile), beliefs are slow to decline, and the learning policy uses more SGE than is optimal under certainty as the planner believes SGE is more effective than it actually is.

5.2.1 Sensitivity of Optimal Policy

To get an idea of the relative importance of SGE effectiveness versus climate feedbacks for optimal SGE policy, Figure 6 graphs the results under certainty for various values of the true feedback parameter holding fixed the true SGE effectiveness parameter. Figure 7 graphs the results under certainty for various values of the true SGE effectiveness holding fixed the true feedback parameter. The optimal SGE policy is convex in the feedback parameter, becoming increasingly sensitive as the feedback parameter increases. Feedbacks have a compound effect on the temperature (generating the fat tail in the climate sensitivity). Therefore, SGE usage responds more when feedback parameter is large, to prevent the large temperature change that would result from the feedback effects compounding over time in the absence of a policy intervention.

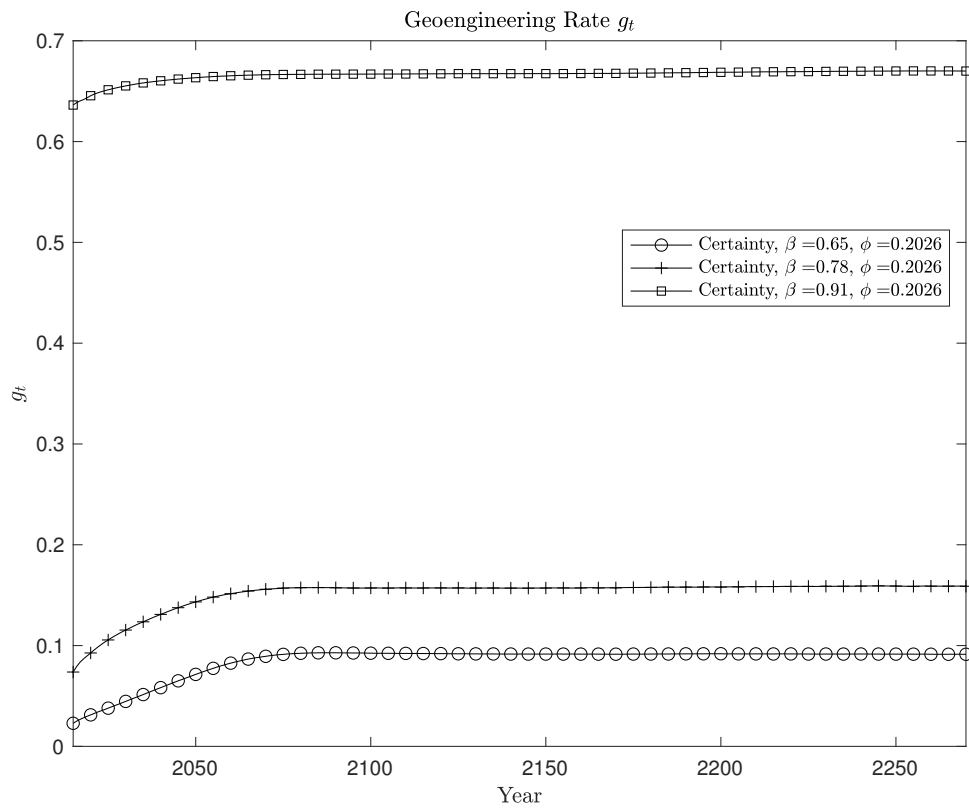


Figure 6: SGE policy, g_t under certainty, with the true value of ϕ equal to the mean of the prior beliefs and the true value of β equals the mean of the prior beliefs and one and two standard deviations above the mean of the prior beliefs. Each line is the mean of 1000 simulations to smooth the effect of the weather shocks.

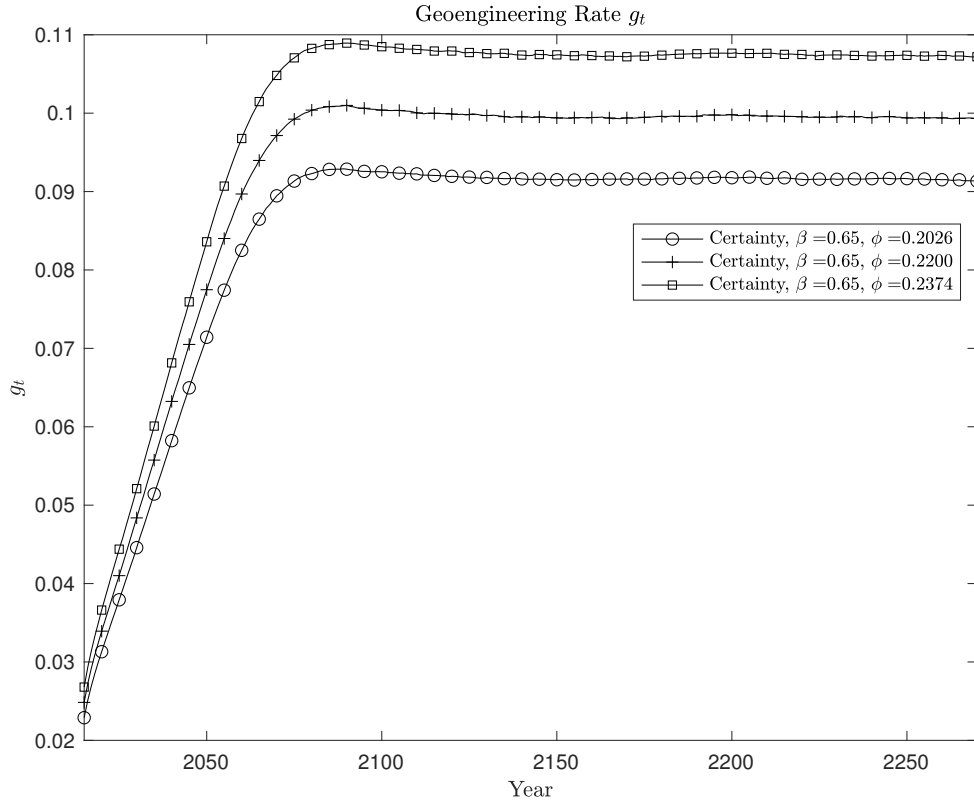


Figure 7: SGE policy, g_t under certainty, with the true value of β equal to the mean of the prior beliefs and the true value of ϕ equals the mean of the prior beliefs and one and two standard deviations above the mean of the prior beliefs. Each line is the mean of 1000 simulations to smooth the effect of the weather shocks.

The value of perfect information is important for policy makers considering investment in R&D that might reduce uncertainty. The value of perfect information relative to uncertainty with learning, vpi , is:

$$vpi = \int_{\tilde{\beta}, \tilde{\phi}} v(k_0, T_0, m_0, t_0, \tilde{\beta}, \tilde{\phi}, 0, 0, 0) \Phi([\beta, \phi], [\beta_0, \phi_0], W_0) d\tilde{\beta} d\tilde{\phi} - v(k_0, T_0, m_0, t_0, \beta_0, \phi_0, W_0). \quad (32)$$

Equation (32) reflects that if the planner had access to an experiment which revealed the true values with certainty today, then ex ante in expectation the experiment is equivalent to a draw from the prior distribution, the result of which is the true value the planner then knows with certainty. Thus, perfect information allows the planner to begin with correct priors regardless of what $\tilde{\beta}$ and $\tilde{\phi}$ turn

out to be, whereas the value with learning reflects that the planner must start with priors of β_0 and ϕ_0 regardless of the true values. Performing the integration numerically on the value function results in a value of perfect information of approximately 1%. Given the risk aversion coefficient of 1.45, this implies households would be indifferent between perfect information and the learning model with a 2.15% increase in consumption per person per year in perpetuity. Although not large in percentage terms, this represents a world-wide increase in consumption of over one trillion dollars per year.⁴⁰

Figure 5 also shows that the mean SGE policy under learning is initially much greater than the certainty policy, as SGE is used initially as a precaution against the possibility that the climate sensitivity is very high. Since learning quickly rules out this case in most simulations, SGE policy becomes similar under the two simulations.⁴¹

5.2.2 Learning Case vs. No-learning Case

Figure 8 contrasts optimal SGE policy under learning with optimal policy under uncertainty but no learning.

⁴⁰Note however, that small computational errors in the percentage gain magnify when converting to dollar values.

⁴¹Kelly and Tan (2015) find that, for abatement, the mean certainty and uncertainty policies are similar after the first decade.

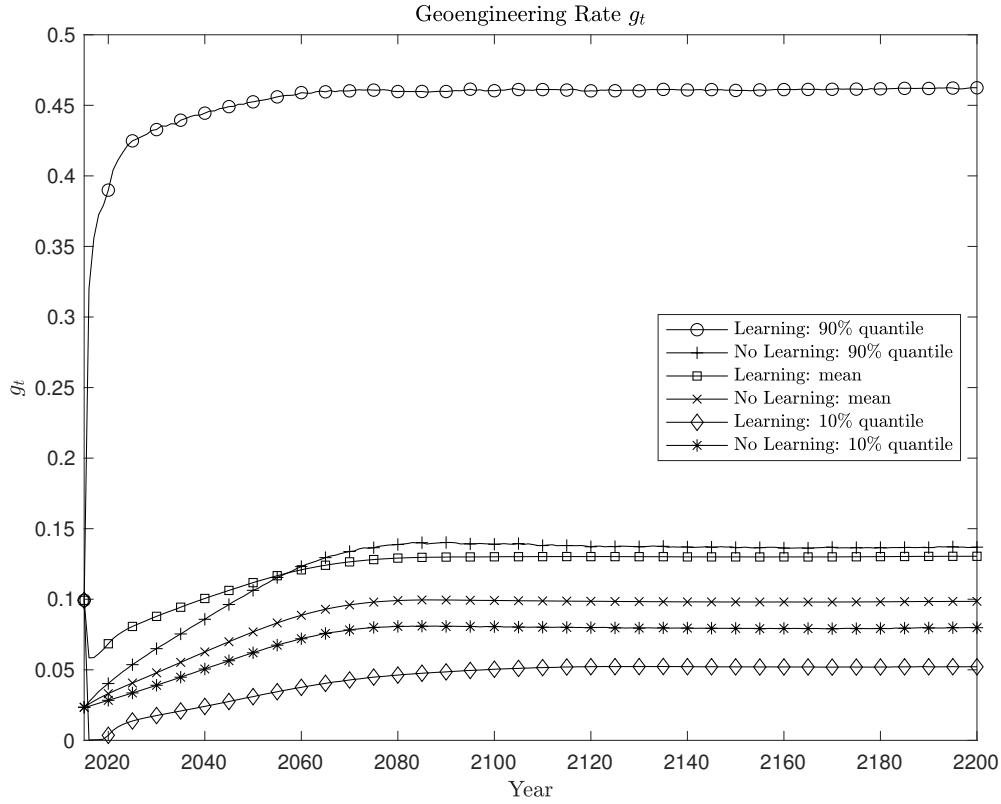


Figure 8: SGE policy, g_t , 10,000 simulations, learning and no learning. Quantiles for the learning case are over beliefs β_t and ϕ_t whereas quantiles for no learning are over true values $\tilde{\beta}$ and $\tilde{\phi}$.

Because the planner knows that no learning will take place, the initial-period policy given no learning differs from the initial policy under learning even though initial beliefs are identical.

Initially, SGE under no learning is about 0.03. This is less than the average of the certainty initial values of about 0.08. Since SGE policy is convex in the feedback parameter under certainty, the average of the starting values is greater than the starting value under certainty when the true value is the prior mean. This effect is offset to some degree as with no learning the policy at the prior mean is greater than the policy under certainty of the prior mean for precautionary reasons. However, this effect is relatively small because Figure 8 shows that even under no learning, some tailoring of the SGE policy is possible.

As the simulations progress, learning allows the planner to tailor SGE policy, substantially increasing SGE when beliefs indicate climate change is severe and SGE is effective. Even without learning, there is some difference across the policies despite no difference in beliefs. This is because

temperature differs across the simulations due to the different true parameter values, and policies respond to that difference. If the temperature rises unexpectedly (due to the true value of the climate sensitivity being higher than the believed value), and no learning takes place, the planner attributes this only to higher weather shocks while anticipating lower future temperatures based on the initial beliefs. Therefore, the planner increases SGE by a relatively small amount. In contrast, under learning, when the temperature rises beliefs also rise, signalling to the planner that future temperatures will also rise. Therefore, the planner increases SGE by a larger amount with learning. The planner under learning bases SGE policy largely on beliefs because beliefs indicate likely paths of future temperatures. A similar argument holds for SGE effectiveness. When beliefs about SGE effectiveness rise, the planner increases SGE. However, without learning, even if SGE is effective the planner keeps the prior belief and does not increase SGE usage. Conversely, when the climate feedbacks are small and the temperature remains low, the planner without learning uses more SGE, anticipating future temperatures will evolve according to the prior distribution with a higher mean.

The mean SGE policies differ between learning and no learning for two reasons. First, after period 0, uncertainty is lower under learning. Therefore, SGE has a more certain effect on temperature. All other things equal, this induces the planner to use more SGE under learning. Second, under learning the use of SGE slows learning. Under no learning, the use of SGE does not affect the speed of information gain, which is fixed at zero. All other things equal, this induces the planner to use less SGE under learning. Figure 8 shows that the mean SGE policy is generally greater under learning, reflecting that the uncertainty effect is generally stronger. Finally, the above characterization is somewhat simplified since all state variables are changing in different ways after the initial period, each of which affects SGE policy.

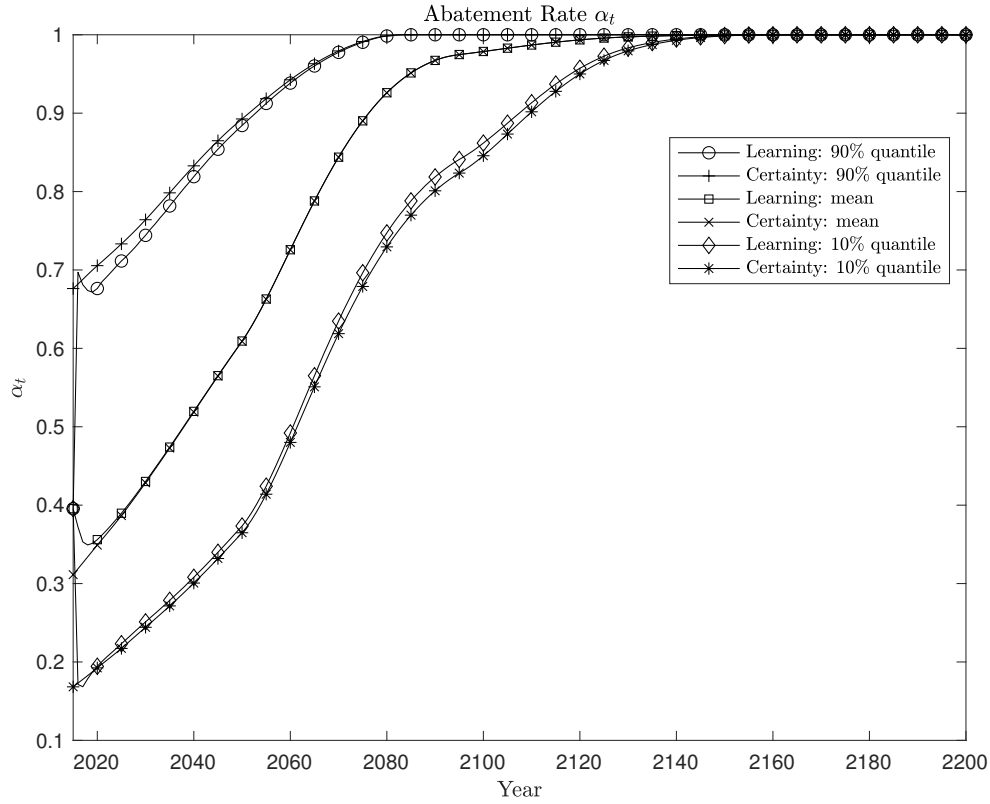


Figure 9: Abatement policy, α_t , 10,000 simulations, learning and certainty.

Figure 9 compares optimal abatement policy under learning and under certainty. Like SGE, the initial mean abatement policy is higher under learning, as a precaution against an potentially high climate sensitivity. Optimal abatement under learning is initially below the level that would occur if the climate sensitivity was known to be high for certain, however. The optimal policy under learning quickly converges to the certainty policy under all cases, since the residual uncertainty about SGE effectiveness is not as important for abatement policy as it is for SGE policy.

Figure 10 contrasts optimal abatement policy under learning with optimal abatement policy under uncertainty but no learning.

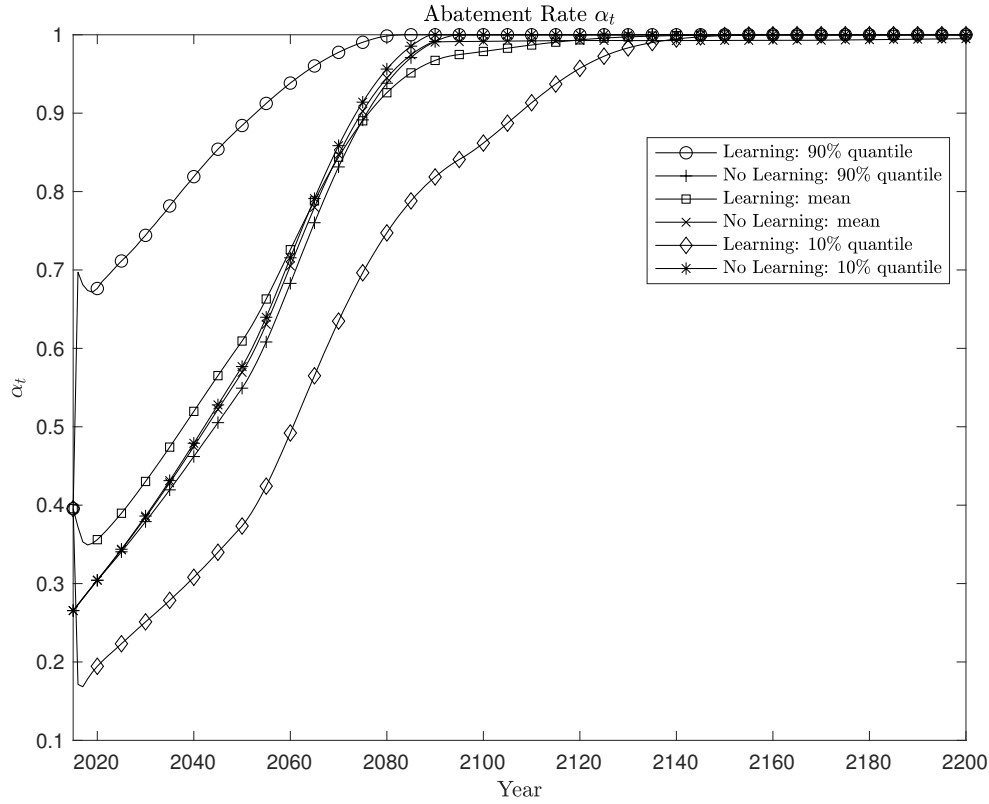


Figure 10: Abatement policy, α_t , 10,000 simulations, learning and no learning.

In the case without learning, the planner's beliefs about the net feedback parameter and SGE effectiveness are unchanged through time and therefore, the policy response in both cases of high and low climate sensitivity hovers around the case with mean climate sensitivity. In contrast, under learning the planner increases abatement when learning indicates the climate sensitivity is relatively high and the reverse.

5.2.3 Optimal policy and the Speed of Learning

Section 3 showed that experimentation with SGE slowed learning about both SGE effectiveness and climate sensitivity, given the 2015 initial conditions. However, the thresholds I_1 and I_2 , which determine whether or not SGE speeds learning, evolve over time as functions of the state variables T , m , W_2 , and W_3 . A natural question is whether or not optimal SGE ever exceeds I_1 or I_2 . If so, for these time periods SGE would speed learning about SGE effectiveness and the climate sensitivity, respectively. Optimal SGE would then have an information benefit from experimentation, rather

than an information cost. Figure 11 investigates this issue by comparing the mean of I_1 over time (the mean of I_2 remains far above one for all t) with the mean of optimal SGE. Figure 11 shows that I_1 declines over time.

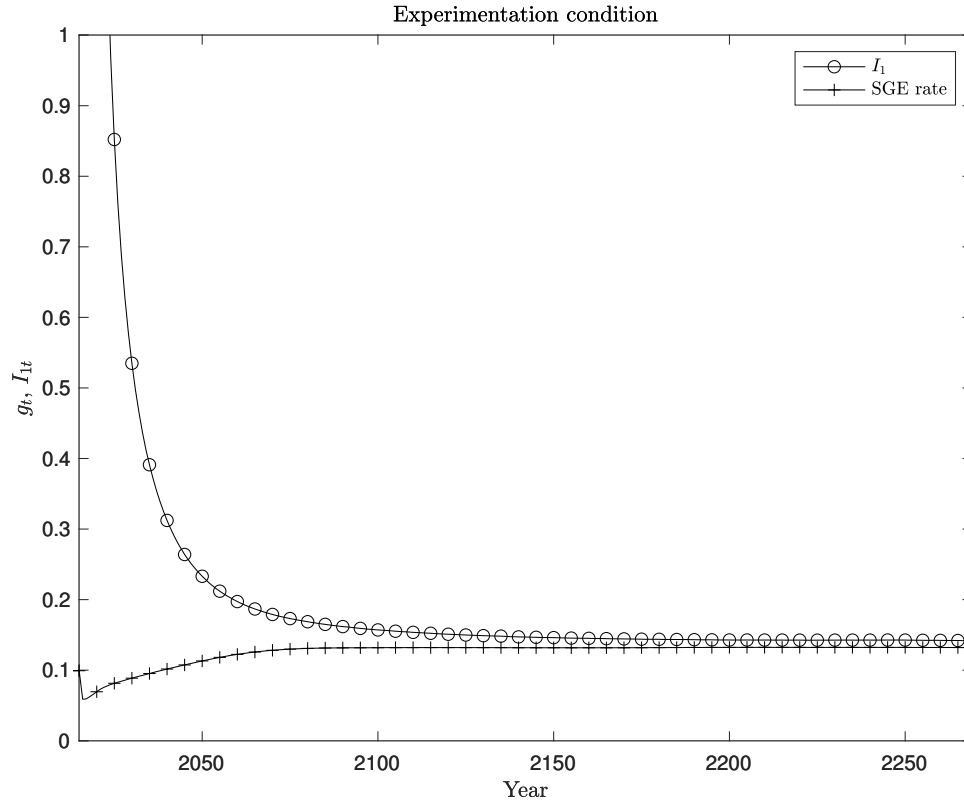


Figure 11: Threshold for SGE to speed learning about SGE effectiveness, I_{1t} , mean 10,000 simulations.

Recall, $I_1 = \frac{T}{F'} \cdot \frac{W_2}{W_3}$ and the first fraction increases modestly as both temperature and forcing through higher CO_2 concentrations increase, but the overall fraction increases because the numerator is linear. In the second fraction, both variances are decreasing as learning proceeds, but the numerator decreases faster (refer to equation 21). Figure 11 shows that the mean of I_1 is greater than the mean SGE deployment g for all time periods (however, $g_{it} > I_{1,it}$ for about 1.45% of all simulations/time periods). The slowing of learning caused by SGE creates a disincentive for SGE use throughout the simulations, not just at the initial conditions.

6 Conclusion

We present a model of optimal climate policy that includes the possibility of using solar geoengineering (SGE) to reduce temperatures. The model allows for uncertainty and endogenous learning over both climate sensitivity (the effect of carbon on temperature) and the effectiveness of SGE (the effect of SGE deployment on temperature). It allows for two climate policy variables – abatement and SGE – both of which affect the rate of learning about those uncertain parameters, creating a potential motivation to experiment with either policy to reduce uncertainty. However, we find that any reasonable level of SGE implementation actually slows learning. Since SGE has an uncertain effect on the climate, use of SGE essentially adds noise to the climate system. Given current initial conditions, no information justification exists for experimentation with SGE. Further, over time, starting at current calibrated initial conditions and continuing along the optimal path, the mean SGE level is below the mean of minimum level required to gain information from SGE experimentation in each period. Thus, the information loss of using SGE continues over time, which acts as a disincentive for SGE usage.

Our model contains several simplifying assumptions that can be relaxed in extensions. We use the simplified version of DICE from Traeger (2014), and DICE itself is a relatively parsimonious integrated assessment model. Optimal policy is defined solely based on net discounted utility for the aggregate representative economy; our model does not incorporate any heterogeneity (e.g. by income, region) and does not consider equity criteria. Our model has endogenous learning over two uncertainties, but other parameters could also be modeled as uncertain, for example, SGE damages.

Theoretically, we decompose the effect of SGE on learning into a signal strength effect and a noise amplification effect. The signal strength effect shows that small-scale SGE implementation reduces the magnitude of the temperature change and thus makes the temperature signal harder to observe amidst the noisy weather shocks. However, with large-scale SGE implementation, the planner expects to see a large change in temperature. Since a weather shock is unlikely to cause a large change in temperature, more information will be gained. Furthermore, any use of an uncertain SGE technology adds noise to the climate system, making learning over climate sensitivity more difficult; we call this the noise amplification effect.

Our numerical results show that SGE deployment is highly sensitive to beliefs about climate sensitivity and SGE effectiveness. Although SGE usage is generally low, it is substantially higher

when beliefs indicate that climate change is likely to be severe and that SGE is effective at reducing temperature. Because the costs of both abatement and SGE are convex, some small amount of SGE is always optimal, as it is less costly than using only abatement.

References

- Abatayo, Anna Lou, Valentina Bosetti, Marco Casari, Riccardo Ghidoni, and Massimo Tavoni. 2020. Solar geoengineering may lead to excessive cooling and high strategic uncertainty. *Proceedings of the National Academy of Sciences* 117 (24): 13393–13398.
- Barrett, Scott, et al. 2014. Climate engineering reconsidered. *Nature Climate Change* 4 (7): 527–529.
- Bickel, J Eric, and Shubham Agrawal. 2013. Reexamining the economics of aerosol geoengineering. *Climatic Change* 119 (3-4): 993–1006.
- Bickel, J Eric, and Lee Lane. 2009. An analysis of climate engineering as a response to climate change. Copenhagen Consensus Center Report, 40.
- Brown, Calum, Peter Alexander, Almut Arneth, Ian Holman, and Mark Rounsevell. 2019. Achievement of paris climate goals unlikely due to time lags in the land system. *Nature Climate Change* 9 (3): 203–208.
- Conte, Marc, and David L. Kelly. 2021. Understanding the improbable: A survey of fat tails in environmental economics. *Annual Review of Resource Economics* 13: 289–310.
- de Oliveira, Henrique. 2018. Blackwell's informativeness theorem using diagrams. *Games and Economic Behavior* 109: 126–131.
- Dietz, Simon, Rick van der Ploeg, Armon Rezai, and Frank Venmans. 2020. Are economists getting climate dynamics right and does it matter? CESifo Working Paper number 8122.
- Emmerling, Johannes, and Massimo Tavoni. 2018. Climate engineering and abatement: A 'flat' relationship under uncertainty. *Environmental and Resource Economics* 69 (2): 395–415.
- Fitzpatrick, Luke, and David L. Kelly. 2017. Probabilistic stabilization targets. *Journal of the Association of Environmental and Resource Economists* 4 (2): 611–57.
- Flegal, Jane A, Anna-Maria Hubert, David R Morrow, and Juan B Moreno-Cruz. 2019. Solar geoengineering: Social science, legal, ethical, and economic frameworks. *Annual Review of Environment and Resources* 44: 399–423.

-
- Forster, Piers, Trude Storelvmo, Kyle Armour, William Collins, Jean-Louis Dufresne, David Frame, Daniel J. Lunt, et al. 2021. The earth's energy budget, climate feedbacks, and climate sensitivity. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, eds. V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Páan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekci, R. Yu, and B. Zhou, chapter 7. New York: Cambridge University Press, 923–1054.
- Goes, Marlos, Nancy Tuana, and Klauss Keller. 2011. The economics (or lack thereof) of aerosol geoengineering. *Climatic Change* 109 (3-4): 719–744.
- Gramstad, Kjetil, and Sigve Tjøtta. 2010. Climate engineering: Cost benefit and beyond. University of Bergen, Department of Economics, No 05/10.
- Hänsel, Martin C, Moritz A Drupp, Daniel JA Johansson, Frikk Nesje, Christian Azar, Mark C Freeman, Ben Groom, et al. 2020. Climate economics support for the un climate targets. *Nature Climate Change* 10 (8): 781–789.
- Haywood, Jim M, Andy Jones, Nicolas Bellouin, and David Stephenson. 2013. Asymmetric forcing from stratospheric aerosols impacts sahelian rainfall. *Nature Climate Change* 3 (7): 660–665.
- Heckendorn, Patricia, Debra Weisenstein, Stephen Fueglistaler, Beiping P. Luo, Eugene Rozanov, M. Schraner, Larry W. Thomason, et al. 2009. The impact of geoengineering aerosols on stratospheric temperature and ozone. *Environmental Research Letters* 4 (4): 045108.
- Helwegen, Koen G, Claudia E Wieners, Jason E Frank, and Henk A Dijkstra. 2019. Complementing co2 emission reduction by solar radiation management might strongly enhance future welfare. *Earth System Dynamics* 10 (3): 453–472.
- Heutel, Garth, Juan Moreno-Cruz, and Katharine Ricke. 2016a. Climate engineering economics. *Annual Review of Resource Economics* 8: 99–118.
- Heutel, Garth, Juan Moreno-Cruz, and Soheil Shayegh. 2016b. Climate tipping points and solar geoengineering. *Journal of Economic Behavior and Organization* 132: 19–45.
- . 2018. Solar geoengineering, uncertainty, and the price of carbon. *Journal of Environmental Economics and Management* 87: 24–41.

- hr/>
- Hwang, In Chang, Frédéric Reynès, and Richard SJ Tol. 2017. The effect of learning on climate policy under fat-tailed risk. *Resource and Energy Economics* 48: 1–18.
- Hwang, In Chang, Richard SJ Tol, and Marjan W Hofkes. 2019. Active learning and optimal climate policy. *Environmental and Resource Economics* 73 (4): 1237–1264.
- Jensen, Sverre, and C.P. Traeger. 2013. Optimally climate sensitive policy under uncertainty and learning. Department of Agricultural & Resource Economics, UC Berkeley.
- Jones, Anthony C, James M Haywood, Andy Jones, and Valentina Aquila. 2016. Sensitivity of volcanic aerosol dispersion to meteorological conditions: A pinatubo case study. *Journal of Geophysical Research: Atmospheres* 121 (12): 6892–6908.
- Keith, David W. 2010. Photophoretic levitation of engineered aerosols for geoengineering. *Proceedings of the National Academy of Sciences* 107 (38): 16428–16431.
- Keith, David W, Riley Duren, and Douglas G. MacMartin. 2014. Field experiments on solar geoengineering: Report of a workshop exploring a representative research portfolio. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372 (2031): 20140175.
- Keller, David P, Ellias Y Feng, and Andreas Oschlies. 2014. Potential climate engineering effectiveness and side effects during a high carbon dioxide-emission scenario. *Nature Communications* 5: 3304.
- Kelly, David L., and Charles D. Kolstad. 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23 (4): 491–518.
- Kelly, David L., and Zhuo Tan. 2015. Learning and climate feedbacks: Optimal climate insurance and fat tails. *Journal of Environmental Economics and Management* 72: 98–122.
- Klepper, Gernot, and Wilfried Rickels. 2014. Climate engineering: Economic considerations and research challenges. *Review of Environmental Economics and Policy* 8 (2): 270–289.
- Kravitz, Ben, and Douglas G MacMartin. 2020. Uncertainty and the basis for confidence in solar geoengineering research. *Nature Reviews Earth & Environment* 1 (1): 64–75.

- hr/>
- LaRiviere, Jacob, David Kling, James N. Sanchirico, Charles Sims, and Michael Springborn. 2018. The treatment of uncertainty and learning in the economics of natural resource and environmental management. *Review of Environmental Economics and Policy* 12 (1): 92–112.
- Leach, Andrew J. 2007. The climate change learning curve. *Journal of Economic Dynamics and Control* 31 (5): 1728–1752.
- Lemoine, Derek, and Ivan Rudik. 2017. Managing climate change under uncertainty: Recursive integrated assessment at an inflection point. *Annual Review of Resource Economics* 9: 117–142.
- Lemoine, Derek, and Christian Traeger. 2014. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy* 6 (1): 137–166.
- Moreno-Cruz, Juan, and David Keith. 2013. Climate policy under uncertainty: a case for solar geoengineering. *Climatic Change* 121: 431–444.
- National Research Council, et. al. 2015. *Climate Intervention: Reflecting Sunlight to Cool Earth*. National Academies Press.
- Nordhaus, William D. 2016. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences* 114 (7): 1518–1523.
- Pastor, AV, DCS Vieira, FH Soudijn, and OY Edelenbosch. 2020. How uncertainties are tackled in multi-disciplinary science? a review of integrated assessments under global change. *Catena* 186: 104305.
- Proctor, Jonathan, Solomon Hsian, Jennifer Burney, Marshall Burke, and Wolfram Schlenker. 2018. Estimating global agricultural effects of geoengineering using volcanic eruptions. *Nature* 560 (7719): 480–483.
- Quaas, Martin F, Johannes Quaas, Wilfried Rickels, and Olivier Boucher. 2017. Are there reasons against open-ended research into solar radiation management? a model of intergenerational decision-making under uncertainty. *Journal of Environmental Economics and Management* 84: 1–17.
- Raftery, Adrian E, Alec Zimmer, Dargan MW Frierson, Richard Startz, and Peiran Liu. 2017. Less than 2 c warming by 2100 unlikely. *Nature Climate Change* 7 (9): 637.

-
- Ramaswamy, Venkatachalam, W Collins, J Haywood, J Lean, N Mahowald, G Myhre, V Naik, et al. 2018. Radiative forcing of climate: the historical evolution of the radiative forcing concept, the forcing agents and their quantification, and applications. *Meteorological Monographs* 59: 14.1–14.101.
- Ricke, Katharine L. 2011. Characterizing impacts and implications of proposals for solar radiation management, a form of climate engineering. Ph.D. thesis, Carnegie Mellon University.
- Ricke, Katharine L, Daniel J Rowlands, William J Ingram, David W Keith, and M Granger Morgan. 2012. Effectiveness of stratospheric solar-radiation management as a function of climate sensitivity. *Nature Climate Change* 2 (2): 92–96.
- Robock, Alan, Douglas G MacMartin, Riley Duren, and Matthew W Christensen. 2013. Studying geoengineering with natural and anthropogenic analogs. *Climatic Change* 121 (3): 445–458.
- Roe, Gerard H., and Marcia B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318 (5850): 629–32.
- Roe, Gerard H, and Yoram Bauman. 2013. Climate sensitivity: Should the climate tail wag the policy dog? *Climatic Change* 117 (4): 647–62.
- Shayegh, Soheil, and Valerie M Thomas. 2015. Adaptive stochastic integrated assessment modeling of optimal greenhouse gas emission reductions. *Climatic Change* 128 (1-2): 1–15.
- Solomon, Susan, D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M.Tignor, et al. 2007. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. New York: Cambridge University Press.
- Spafford, Lynsay, and Andrew H MacDougall. 2020. Quantifying the probability distribution function of the transient climate response to cumulative co2 emissions. *Environmental Research Letters* 15 (3): 034044.
- Traeger, Christian. 2014. A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics* 59: 1–37.
- Wagner, Gernot, and Martin L Weitzman. 2018. Potentially large equilibrium climate sensitivity tail uncertainty. *Economics Letters* 168: 144–146.

Weitzman, Martin L. 2009. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91: 1–19.

Figure A1

This is the author's accepted manuscript without copyediting, formatting, or final corrections. It will be published in its final form in an upcoming issue of JAERE, published by The University of Chicago Press on behalf of the Association of Environmental and Resource Economists. Include the DOI when citing or quoting. <https://doi.org/10.1016/j.jaere.2024.100000>

Expected Percent Absolute Forecast Errors, Conditional on 2015 Priors

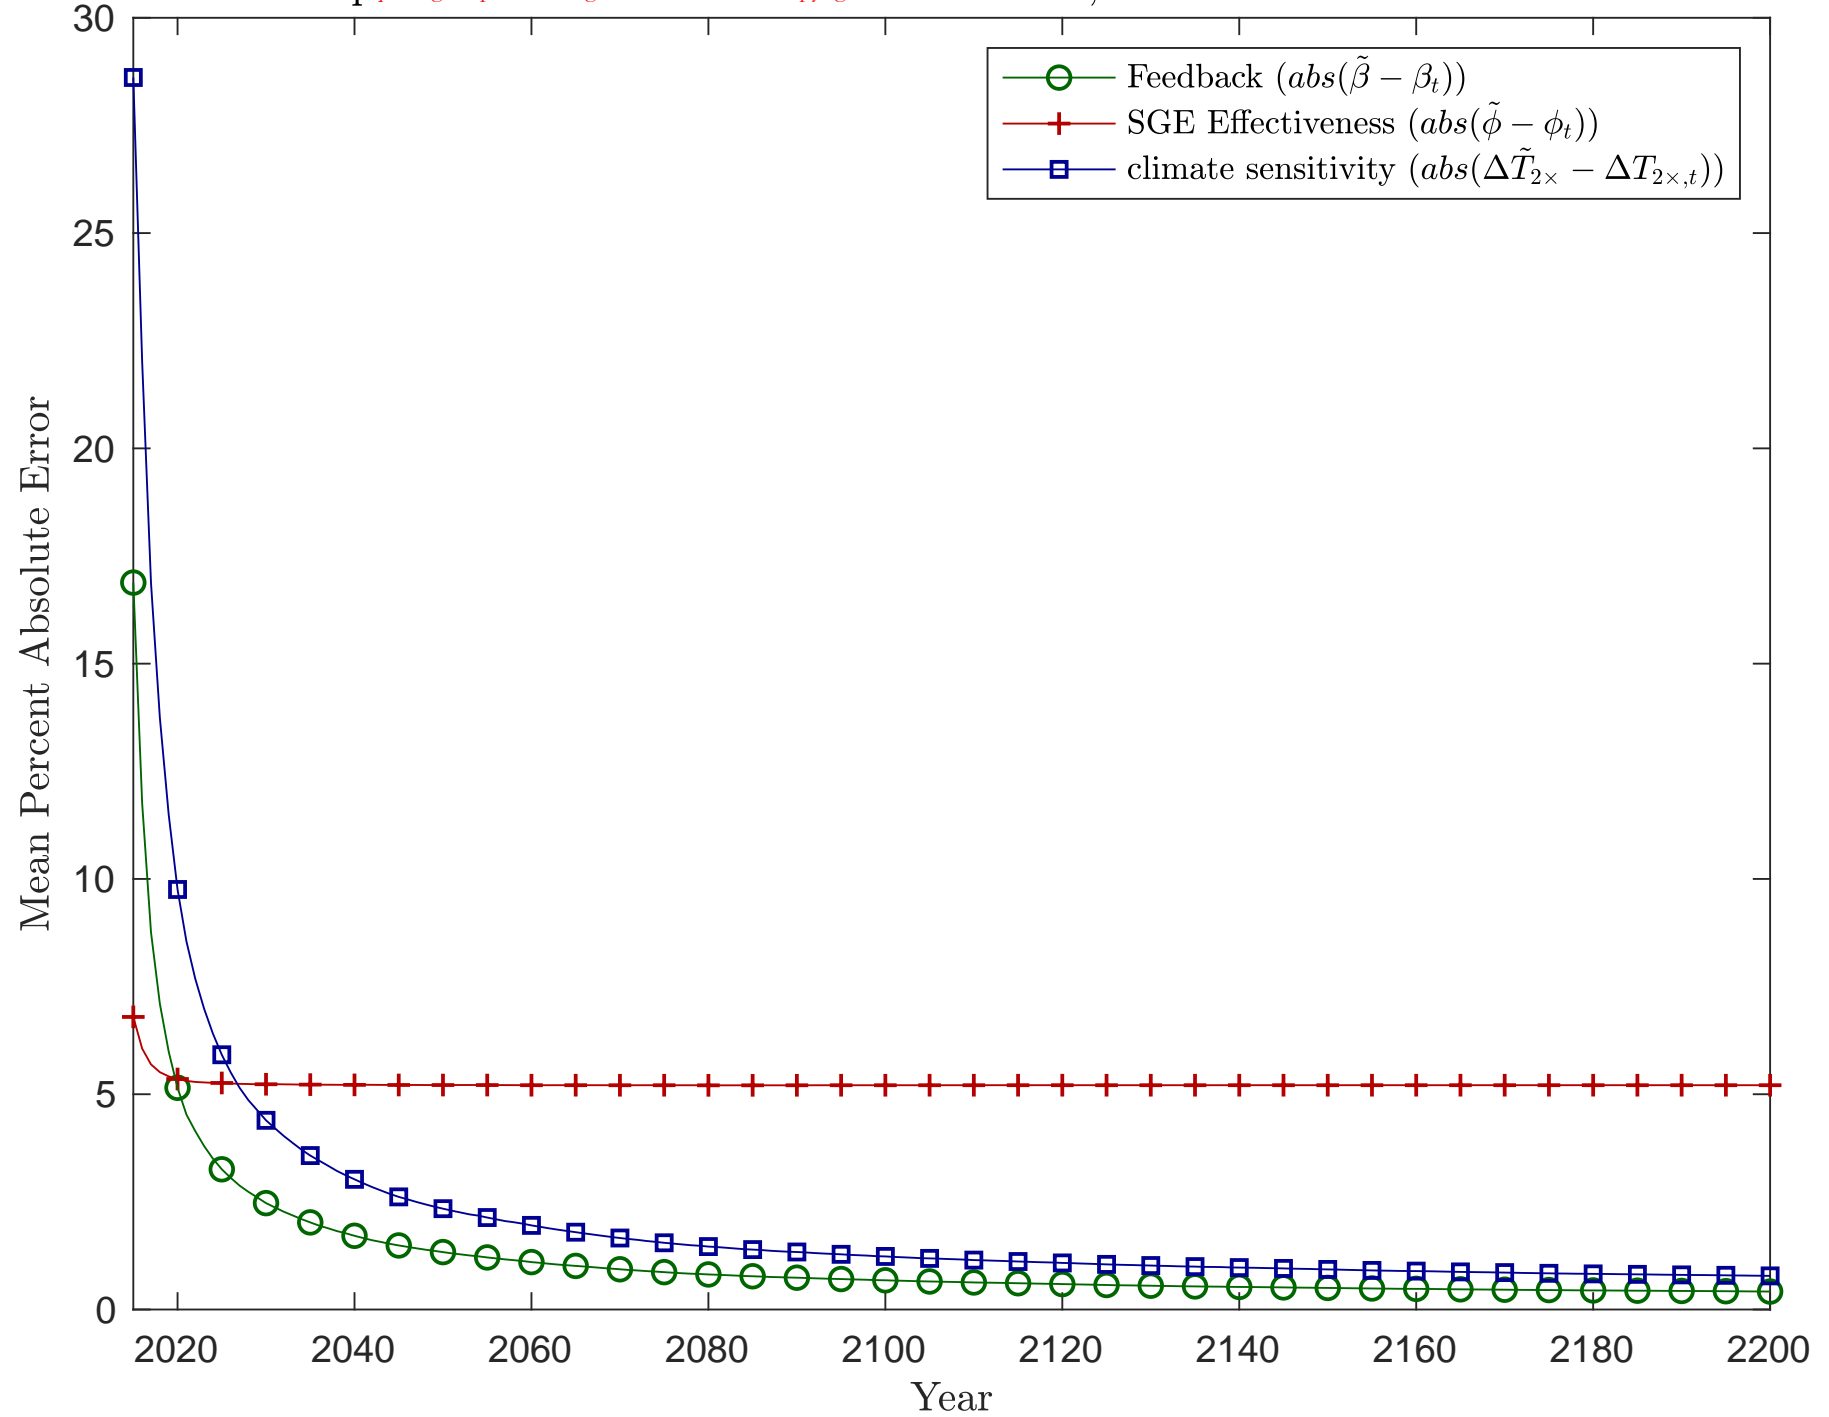


Figure A2

This is the author's accepted manuscript without copyediting, formatting, or final corrections. It will be published in its final form in an upcoming issue of JAERE, published by The University of Chicago Press on behalf of the Association of Environmental and Resource Economists. Include the DOI when citing or quoting: <https://doi.org/10.1086/72908>. Copyright 2024. Association of Environmental and Resource Economists.

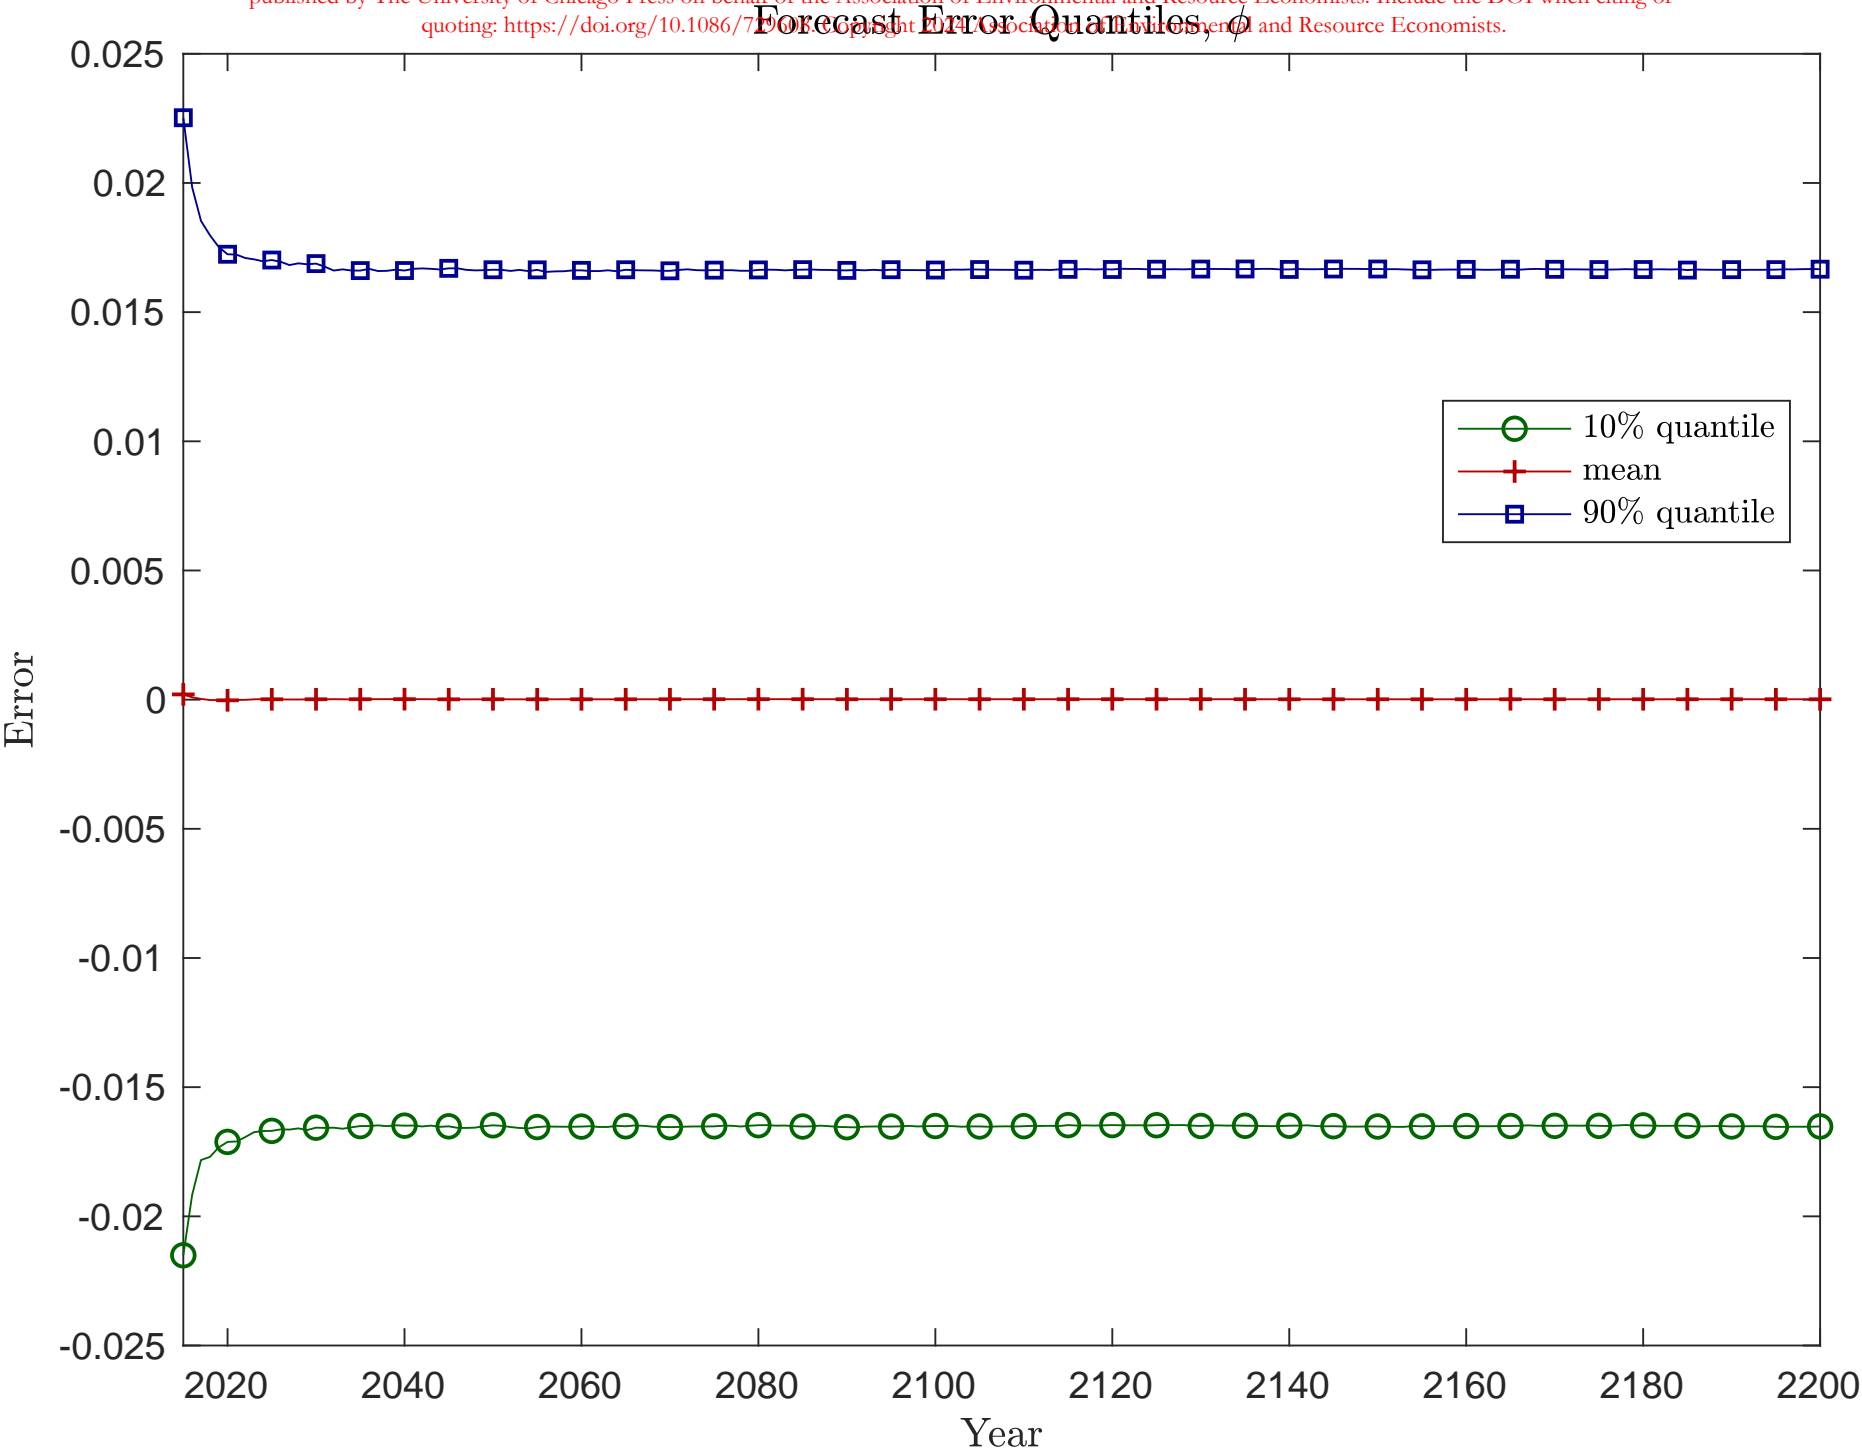


Figure A3

This is the author's accepted manuscript without copyediting, formatting, or final corrections. It will be published in its final form in an upcoming issue of JAERE, published by The University of Chicago Press on behalf of the Association of Environmental and Resource Economists. Include the DOI when citing or quoting: <https://doi.org/10.1086/aeer.2020.11.1.11111111>

