Available online at www.sciencedirect.com
SCIENCE DIRECT

ELSEVIER Journal of Environmental Economics and Management 48 (2004) 723–741

JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT

http://www.elsevier.com/locate/jeem

Uncertain climate thresholds and optimal economic growth

Klaus Keller,[a,*] Benjamin M. Bolker,[b] and David F. Bradford[c,d,e,f]

[a] Department of Geosciences, The Pennsylvania State University, University Park, PA 16802-2714, USA
[b] Department of Zoology, University of Florida, Gainesville, FL 32611-8525, USA
[c] Woodrow Wilson School of Public and International Affairs, Princeton University, Princeton, NJ 08544-1013, USA
[d] School of Law, New York University, New York, NY 10012-1074, USA
[e] Department of Economics, Princeton University, Princeton, NJ 08544-1013, USA
[f] National Bureau of Economic Research, Cambridge, MA 02138-5398, USA

Received 14 August 2000; revised 17 August 2001

Abstract

We explore the combined effects of a climate threshold (a potential ocean thermohaline circulation collapse), parameter uncertainty, and learning in an optimal economic growth model. Our analysis shows that significantly reducing carbon dioxide (CO_2) emissions may be justified to avoid or delay even small (and arguably realistic) damages from an uncertain and irreversible climate change—even when future learning about the system is considered. Parameter uncertainty about the threshold specific damages and the CO_2 level triggering a threshold can act to decrease near-term CO_2 abatements that maximize expected utility.

Keywords: Climate policy; Climate thresholds; Abrupt climate change; North Atlantic thermohaline circulation; Optimal growth model; Learning; Uncertainty

1. Introduction

Anthropogenic emissions of carbon dioxide (CO_2) may cause significant climate change during the next century [1]—a situation prompting divergent policy recommendations. Some climate researchers have argued that CO_2 emissions should be considerably abated to avoid climate thresholds such as a collapse of the ocean thermohaline circulation (e.g., [46,51]). In contrast, many optimal economic growth models suggest that the tradeoff between uncertain future climate damages and certain present costs for controlling CO_2 emissions justify only low levels of

*Corresponding author.
E-mail address: kkeller@geosc.psu.edu (K. Keller).

0095-0696/$ - see front matter
doi:10.1016/j.jeem.2003.10.003

near-term CO_2 abatement (e.g., [39,55]). Here we explore how this discrepancy might be explained by different assumptions and value judgments.

This disagreement about future CO_2 abatements stems partly from a difference in modeling priorities, exacerbated by the difficulty of constructing models that address the concerns of economists and climate scientists simultaneously. On the one hand, economists emphasize the need to find policies for controlling climate change that optimize the trade-off between climate damages and lost opportunities for consumption or economic growth. Economists have typically used simple coupled models of climate–economic systems with smooth relationships among climatic and economic variables [40,55]. On the other hand, climate scientists often focus on the strongly nonlinear character of the global climate system, including thresholds and changes that are irreversible on the time scale of human civilizations (hysteresis) [3]. Economists and climate scientists recognize the importance of uncertainty in parameters, but analyzing the combined effects of uncertainty and climate thresholds in optimal growth models has been difficult since this combination can pose nearly intractable computational problems [22,25,36]. This paper demonstrates a methodological step towards overcoming these computational limitations, finds the optimal policy response to uncertain and sharply nonlinear climate changes, and helps to reconcile the priorities of economists and climate scientists. The results are of more than pure methodological interest; we show that surprisingly small threshold-specific damages (as low as 0.5% of gross world product, roughly two orders of magnitude smaller than the catastrophic events typically assumed in optimal growth models [9,39]) significantly increase optimal CO_2 abatement. Further, in a curious reversal of the perhaps intuitive effect of uncertainty if risk aversion were the dominant effect, parameter uncertainty in this example can act to decrease optimal CO_2 abatement in the near term.

We first modify an archetypical optimal growth model [39] to incorporate parameter uncertainty, learning, and a climate threshold. Specifically, we consider the example of a change in the global ocean circulation caused by anthropogenic climate change. The climatic and economic effects of such an ocean circulation change can be rather abrupt [8,56], and show a considerable hysteresis response [45]. Other climate thresholds (such as a potential collapse of the West Antarctic Ice sheet [42]) are, of course, possible and could be analyzed within a similar framework. We then explore the range of possible outcomes in sensitivity studies with respect to the climate sensitivity and the threshold-specific damages. Finally, the results of the sensitivity study are used to explain the (maybe) counterintuitive results of the decision analysis under uncertainty.

2. The model

2.1. Model structure

We adopt the DICE94 model as described in [39] as a starting point. This model is relatively simple, impressively well documented, and has been used in numerous studies to analyze the effects of varying modeling assumptions on climate policy. The DICE model's simplicity necessarily limits its ability to represent each subsystem with high accuracy. For example, the representation of the carbon cycle in the DICE model has been extensively criticized and revised

[23,40]. One cannot take the DICE model's policy recommendations too literally, but it has repeatedly provided useful insights into the basic phenomena determining optimal policy under climate change [27,39,41,54]. In the following section we briefly outline the salient features of the DICE model for this study. Nordhaus [39] provides a detailed description of the model used in this study.

The DICE model is a dynamic growth model that incorporates a simple feedback mechanism between economic activity and climate change. The objective in the model is to maximize social well-being with a particular set of decisions about investment and CO_2 abatement over time. Well-being is represented in the model by a flow of utility U, defined as the product of the logarithm of per capita consumption per year c, and the exogenously given population L:

$$U(t) = L(t) \ln c(t) \tag{1}$$

(see the appendix for a list of symbols). The objective is to identify a policy that maximizes the discounted sum of utility (U^*):

$$\max U^* = \sum_{t=t_o}^{t^*} U(t)(1+\rho)^{-t}, \tag{2}$$

which is calculated by applying a pure rate of social time preference (ρ) of 3% per year to the flow of utility at time t from some starting point t_o to an appropriate time horizon t^*. The finite time horizon is a necessary numerical approximation to the infinite horizon optimal growth model (e.g., [27,38]). The nonlinearities of the integrated assessment model preclude the use of analytical solutions. We choose a time-horizon for the numerical model implementation of 470 years to minimize numerical artifacts. Using a longer time horizon does not change the results for the analyzed time horizon of 2150. In addition, extrapolations beyond the year 2150 would arguably be quite speculative.

Consumption is derived from production (Q), determined by a standard production function influenced by population, capital, climate damages, and investment into CO_2 abatement. Economic activities result in CO_2 emissions into the atmosphere. A simple carbon cycle model relates CO_2 emissions to atmospheric CO_2 concentrations. In the model, a constant fraction β of carbon emissions is added to the atmospheric carbon stock M (the rest is assumed to be absorbed immediately by carbon sinks). A portion δ_M of the atmospheric carbon in excess of the preindustrial stock of 590 Gt is exported during each time step to the deep ocean so that the atmospheric stock evolves according to

$$M(t) = 590 + \beta\Delta t E(t-\Delta t) + (1-\delta_M)[M(t-\Delta t) - 590], \tag{3}$$

where Δt is the length of the time step in the model and $E(t)$ are the anthropogenic CO_2 emissions in Gigatons of carbon per year. The anthropogenic CO_2 emissions ($E(t)$) are controlled by the policy choice of CO_2 emissions abatement (μ) according to

$$E(t) = BAU(t)(1-\mu(t)), \tag{4}$$

where BAU are the business as usual emissions at no control (i.e., $\mu = 0$), calculated from the gross world product and the carbon intensity of production.

Anthropogenic CO_2 emissions increase the atmospheric CO_2 concentrations. Atmospheric CO_2 warms the earth's surface because it absorbs some of the outgoing infrared radiation and reflects parts of this energy back to the earth's surface. Atmospheric CO_2 behaves much like a greenhouse roof which is transparent to the incoming shortwave radiation but opaque to the outgoing infrared radiation—hence it is called a "greenhouse gas". Changes in greenhouse gas concentrations change the earth's energy balance. The magnitude of this change in the energy balance is typically called "radiative forcing" [1, p. 24]. The deviation (T) of the globally averaged temperature from the preindustrial value is determined from the radiative forcing using a simple climate model. An important parameter in this climate model is the climate sensitivity (λ^*), which represents the predicted increase in equilibrium temperature for a hypothetical doubling of atmospheric CO_2.

The economic damages caused by climate change (D) are taken in the standard DICE model as a fraction of gross world product (GWP) and are a function of the temperature deviations only.

$$D(t) = \theta_1 T(t)^{\theta_2}, \tag{5}$$

where θ_1 and θ_2 are model parameters. The cost of CO_2 emissions abatement TC, expressed again as a fraction of GWP, is given by

$$TC(t) = b_1 \mu(t)^{b_2}, \tag{6}$$

where b_1 and b_2 are model parameters. The abatement costs and climate damages are subtracted from the global output.

2.2. Modifications to the DICE model

We depart from the original DICE model structure in two ways: (i) We consider damages caused by an uncertain environmental threshold imposed by an ocean circulation change (technically known as a North Atlantic thermohaline circulation collapse, discussed below); (ii) we account for central uncertainties by formulating the model as a probabilistic optimization problem.

First, we consider a realistic environmental threshold—a potential change in the global ocean circulation [50]. At the present, warm surface currents flow towards the North Atlantic where they cool due to the heat exchange with the cooler atmosphere. The cooling of the surface waters acts to increase the water density, while the net freshwater input into the North Atlantic decreases the salt concentration which acts to decrease the water density [5]. At some point, the surface waters become dense enough that they sink and form deep waters. This deep water formation is part of global flow pattern which is sometimes referred to as the "global conveyor belt" [7]. Because this ocean circulation system is driven by changes in water densities caused by changes in the heat and salinity content of the water, it is also called the "thermohaline circulation" (THC). It is important to note that this ocean circulation system warms North Western Europe considerably by the heat transport from the low latitudes to the high latitudes [47]. Greenhouse gas emissions could cause a THC collapse by two main mechanisms [50]. First, a warmer climate would result in warmer surface waters with lower densities (and hence lower deep water formation rates). Second, warmer climates would likely increase the freshwater input into the North Atlantic, which would reduce the water densities (and again the deep water formation rates).

Ocean modeling studies suggest that the THC could change dramatically as the results of anthropogenic greenhouse gas emissions [31,32,49,51,53]. Many models suggest that the THC may collapse when the equivalent CO_2 concentration ($P_{CO_{2,e}}$, the concentration of CO_2 and all other greenhouse gases expressed as the concentration of CO_2 that leads to the same radiative forcing) rises above a critical value ($P_{CO_{2,e,crit}}$) [31,51]. The $P_{CO_{2,e,crit}}$ derived by Stocker and Schmittner [51] depends on the climate sensitivity and the rate of CO_2 increase in the atmosphere. As a result, $P_{CO_{2,e,crit}}$ depends on various economic assumptions. Here we approximate previous results [24] for optimal CO_2 stabilization levels to avoid a THC collapse under a variety of climate sensitivities. These estimates are derived from the DICE model adopting the vulnerability of the THC as estimated by Stocker and Schmittner [51], and using the preservation of the THC has an absolutely binding constraint (Fig. 1, panel A). It is important to note that some of our

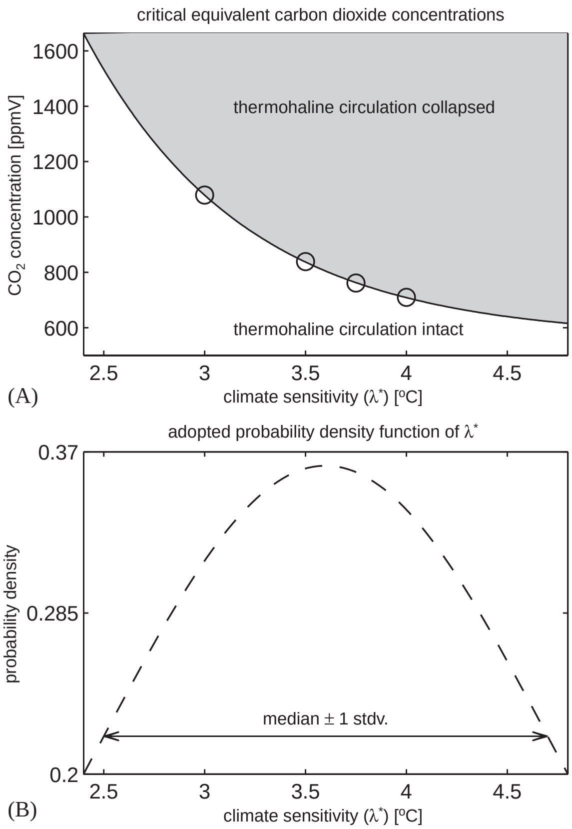

Fig. 1. Critical equivalent carbon dioxide concentrations ($P_{CO_{2,e,crit}}$) as a function of climate sensitivity (panel A) and the parameter uncertainty in climate sensitivity (panel B). The critical equivalent carbon dioxide concentration (line in panel A) separates the region where the ocean thermohaline circulation (THC) in the model is collapsed (shaded region) and the region where the THC is intact. The line in panel A is a fit to stabilization levels (circles) calculated from the ocean model results of Stocker and Schmittner [51] in conjunction with previously estimated optimal stabilization policies in the DICE model [24]. Using an exponential function [$P_{CO_{2,e,crit}} = a_1 + a_2 e^{(a_3 \cdot \lambda^*)}$] to represent the relationship instead of a linear relationship increases the model fit considerably (F-value increase by a factor of 27). The lower panel (B) illustrates the adopted probability density function for the climate sensitivity parameter (over the range shown in panel A). See text for details.

calculations extrapolate beyond the 2–4°C range of climate sensitivity explored by Stocker and Schmittner [51].

A THC collapse would likely affect the global climate system (with the changes likely centered around northwestern Europe) [31,46], and would certainly disrupt the economic system. Potentially affected ecosystem services could be agriculture, fisheries, and the oceanic CO_2 sink [24,56]. We represent the economic damages cause by a THC collapse by imposing a threshold-specific climate damage (θ_3) for all times after the THC has collapsed in the model.

As a second change to the original DICE model [38], we explore the effects of parameter uncertainty on a policy that maximizes the expected value of the objective function. To this end, we formulate the model as a probabilistic optimization problem. Because we use a numerical solution method, we have to approximate the continuous probability density functions by discrete subsamples. We refer to the discrete subsamples as "States of the World". The technical details of the subsampling procedure are discussed below.

We maximize the expected value of the total discounted utility U^* over all states of the world, weighted by their probability. Expected utility maximization is used as a decision criterion to be consistent with previous studies [27,39,41]. Other decision criteria (such as minimizing the consumption loss of the worst-case scenario analogous to Loulou and Kandia [30]) are possible and could be investigated with our model.

We focus on uncertainty in the threshold damages (θ_3) and the climate sensitivity (λ^*). Tol [56] estimates the potential damages associated with a THC collapse in Western Europe between 0% and 3%. The probability distribution of the threshold-specific damages seems at this time unknown. We explore the effects of the parameter uncertainty for this value by assuming (rather arbitrarily) a uniform distribution between 0 and 3% of GWP, which results in an expected value of 1.5% of GWP. The extrapolation to the global scale might be justified as Tol's [56] damage estimate for Western Europe excludes some potential global damages caused by a THC collapse (e.g., a decrease in the natural CO_2 sinks and fishery yields [24]). Note that even our upper estimate of the THC specific damages is much smaller than the typically assumed 30–60% of GWP for "catastrophic" climate damages [9,39].

For the parameter value of the climate sensitivity (λ^*), we rely on the analysis of Tol and deVos [57], who estimate a mean and standard deviation of 3.6°C and 1.1°C. We furthermore approximate the probability density distribution estimated by Tol and deVos [57] as normal (Fig. 1, panel B). The expected value of 3.6°C is within the range of 1.5−4.5°C range, reported by the IPCC [1]. In fact, the estimate of Tol and deVos [57] is obtained by a Bayesian data analysis using the 1.5–4.5°C range of the IPCC as prior. The range of the climate sensitivity estimated by Tol and deVos [57] exceeds the IPCC range—a fact consistent with numerous other studies. For example, Wigley et al. [62] conclude from their data analysis that the climate sensitivity "could lie anywhere between 1.7°C and 6.4°C". Climate sensitivities estimated from model simulations range from 0.16°C to 8.7°C [21].

2.3. Solution method

Because analytical solutions for the considered type of dynamic optimization problem are not available, they are typically formulated in specialized programming languages like GAMS and solved with commercially available optimization software such as MINOS [39]. MINOS uses a quasi-Newton local search algorithm that relies on symbolic partial derivatives. However, the representation of climate thresholds and hysteresis effects introduces nonsmooth gradients and local optima (discussed below), which severely limit the modified Newton algorithm and other methods based on explicit or implicit derivative information (e.g., [44], pp. 34,35). The numerical problems introduced by the local optima are typically addressed by the choice of a global optimization method such as simulated annealing or a genetic algorithm [2,59].

To overcome the computational problems caused by the consideration of the climate threshold, we apply a global optimization algorithm, which does not rely on gradient information. Specifically, we use the differential evolution algorithm by Storn and Price [52], augmented with a line search method. Here we give a brief summary of the method. Detailed analyses and descriptions of the numerical solution method are given elsewhere [29,36]. The basic idea behind the differential evolution algorithm is to find a global optimum by simulating mutation and selection processes on trial vectors. The method uses an initial guess for the optimal policy and produces a number of slightly different (mutated) policies. These mutated policies are then tested with the objective function. Some of the fittest policies are then selected to yield the next generation of mutated trial vectors. The iteration of the sequential selection and mutation process can be quite successful in identifying the global optimum. In fact, our algorithm converges to basically the same solution from many different initial conditions, in contrast to many other global optimization algorithms [36]. Furthermore, our algorithm reproduces the GAMS/MINOS solution for the original DICE model (i.e., for no threshold-specific damages) very closely [29]. The solution algorithm is available from the authors upon request.

The long time horizon of 470 years used to avoid time-horizon end effects results in 94 optimization variables (47 for abatement and 47 for investment at a time step of 10 years). Identifying the global optimum in the 94-dimensional space is a nontrivial task, and we cannot rigorously prove that the reported results are global optima. However, a large number of positive tests for global optimality suggest that we have indeed identified the global optimum or at least a very close approximation to it. Performed tests include: (i) using different initial conditions [36], (ii) increasing the number of iterations, (iii) deceasing the step-length in the line search method, (iv) increasing the mutation rate in the evolutionary algorithm, and (v) graphically inspecting the utility function along selected sections of the solution space.

Note that local optima or nonsmooth gradients can arise whenever abrupt changes or path dependency are included in an optimal growth model. For example, considering effects such as technological learning curves [34], hysteresis effects in damages [43], catastrophic damages [11], or path-dependent multiple steady states [10] would likely lead to optimization problems similar to those discussed here. As illustrated below, our algorithm can handle this class of problems.

3. Results and discussion

3.1. Sensitivity analyses without parameter uncertainty

We use our model to analyze the response of the optimal abatement policy to a range of threshold specific damages and climate sensitivities. These sensitivity analyses assume perfect information (i.e., the parameters vary between the scenarios but are otherwise assumed to be perfectly known).

3.2. Abatement as a function of threshold-specific damages

When threshold-specific damages are low ($\theta_3 < 0.3$), optimal abatement levels are very close to the original DICE model [39], ranging from roughly 10% in the near future to plateau around 20% in about a century (Fig. 2), and the THC collapses. At our upper bound for the estimated threshold-specific damages ($\theta_3 = 3\%$) near-term abatement is still around 10%, but the model suggests a sharp increase in abatement in the next few decades, up to a plateau of roughly 65% approximately after 2100. For intermediate θ_3 values ($\theta_3 \approx 0.5-1.5\%$), the near-term abatement follows the trajectory calculated for the high-damage case but *decreases* in the long run to a low level, similar to the optimal long-term level for the low-damage case. Such a policy delays a THC collapse for some time but fails to preserve the THC in the distant future. The optimal policy allows a thermohaline collapse in this case because investment in preserving the THC increases the total discounted utility less than consumption and/or investing in the economy. However, delaying a THC collapse as long as abatement levels (and costs) are low is worthwhile. Note that the future optimal abatement increases significantly only for medium to high damages. For very small threshold damages (e.g., θ_3 below 0.3%), the near-term policy is virtually unaffected. We will return below to this important asymmetry in the optimal abatement levels with respect to threshold-specific damages.

We can understand the decision to let the THC collapse, and explore some of the difficulties of optimizing models that consider thresholds, by examining how the objective function (the total

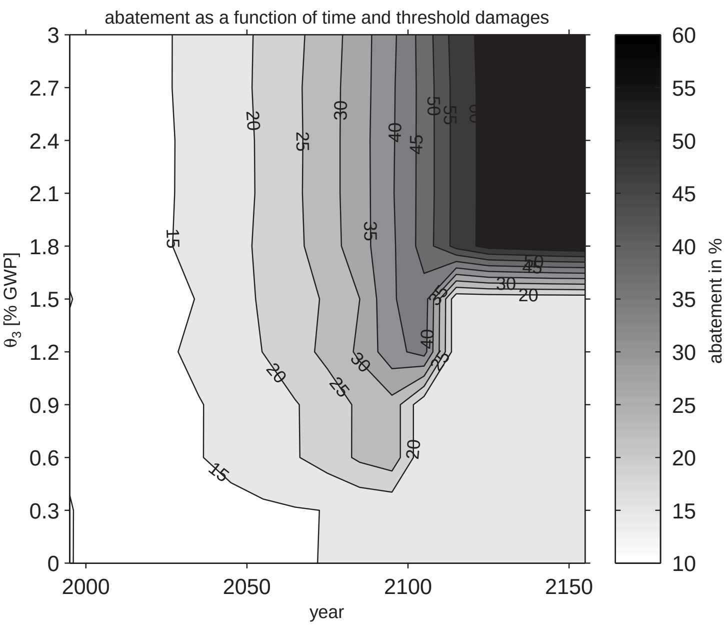

Fig. 2. Optimal abatement policies as a function of threshold-specific damages (θ_3) under parameter certainty. The optimal policy for the central estimate (i.e., $\theta_3 = 1.5\%$ of GWP, and $\lambda^* = 3.6°C$) can also be seen in Fig. 5).

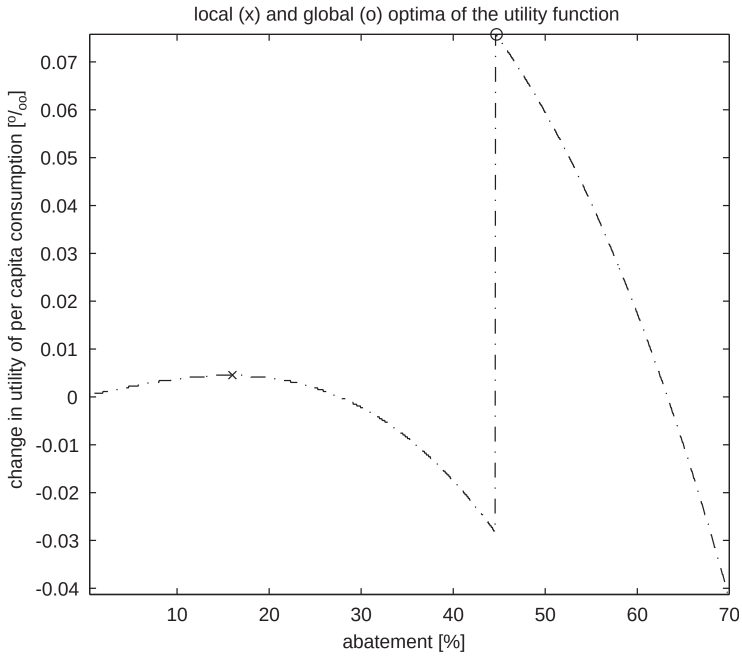

Fig. 3. Relative changes in the objective function with increasing CO_2 abatement in the year 2095. Shown are the changes in the objective function (Eq. (2)) relative to zero abatement. The circle denotes the identified global optimum. The cross depicts a local optimum. The specific model parameters are given in the text.

discounted utility, Eq. (2) varies as a function of abatement in a specific year (Fig. 3). In this example, which perturbs an optimal policy for a climate sensitivity of 3.5°C and a threshold-specific damage of 2%, the objective function increases as abatement in the year 2095 rises from zero to approximately 15% (cross) to yield a first local optimum. This local optimum is close to the solution without a threshold damage (i.e., $\theta_3 = 0\%$). A small increase in abatement from the local optimum decreases the objective function since the marginal costs of abatement exceed the marginal benefits of reducing climate damages and the small increase in abatement does not affect the THC over the relevant time-horizon. However, increasing abatement to a point that stabilizes $P_{CO_{2,e}}$ below the critical value (i.e., abatement levels at roughly 45%, circle in Fig. 3) avoids the damages of a THC collapse in the next time-step. Increasing abatement from roughly 15% to roughly 45% to avoid a THC collapse in the next time step recovers the global optimum in this case and maximizes the total discounted utility.

3.3. *Abatement as a function of climate sensitivity*

Adopting a different climate sensitivity changes $P_{CO_{2,e,crit}}$ (Fig. 1) and, in turn, the optimal abatement levels. For very low climate sensitivities (e.g., $\lambda^* = 1.4$°C $P_{CO_{2,e,crit}}$ is not reached in the next 150 years with or without threshold-specific damages. The increase in CO_2 causes only a small temperature increase and small climate damages, leading to relatively small optimal abatement levels (Fig. 4, which shows results for a threshold-specific damage of 1.5%.).

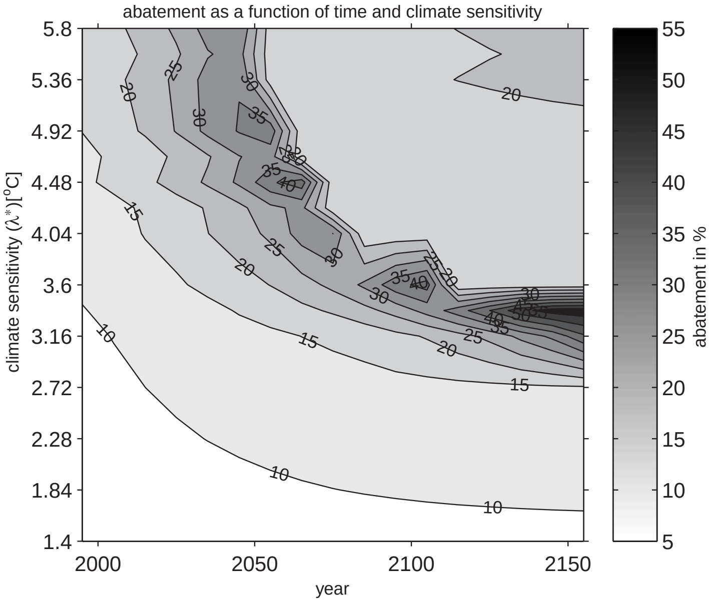

Fig. 4. Optimal abatement policies as a function of climate sensitivity (λ^*) under parameter certainty. The optimal policy for the central estimate (i.e., $\theta_3 = 1.5\%$ of GWP, and $\lambda^* = 3.6°C$) can also be seen in Fig. 5).

Adopting a median climate sensitivity ($\lambda^* = 3.6°C$) changes the situation considerably. The $P_{CO_{2,e,crit}}$ of roughly 805 ppmV (Fig. 1) is reached earlier and the threshold-specific damages of 1.5% justify abatement levels that preserve the THC at least for the next 100 years. For high climate sensitivities $P_{CO_{2,e,crit}}$ is even lower. For example, for a climate sensitivity of 5.6°C, the necessary $P_{CO_{2,e}}$ stabilization level is roughly 580 ppmV (Fig. 1). Because a lower $P_{CO_{2,e}}$ threshold is reached earlier, the near-term abatement levels increase. However, to keep $P_{CO_{2,e}}$ permanently at this lower level would require more expensive abatement measures. The median THC-specific damage of 1.5% does not justify the preservation of the THC in the considered optimum growth model and abatement decreases after roughly 60 years.

3.4. Decision analysis for uncertain parameter values

The sensitivity studies so far assume perfect information for each scenario. These sensitivity studies, while informative, give only limited guidance on which policy would be optimal over the range of possible outcomes. We explore this question by analyzing the effects of parameter uncertainty in θ_3 and λ^* with a Monte Carlo simulation. Expected utility is optimized with one policy, which applies to 100 random samples of the parameter space (states of the world). Technically, we pick 100 random samples of the climate sensitivity and the threshold-specific damage assuming independent probability distributions. The optimization problem is then to identify a single policy which optimizes the sum of utilities of each state of the world. We limit our analysis to 100 samples as larger samples resulted in prohibitive computational demands. Note that this approach neglects that we are likely to learn in the future. We shall return to this important caveat below.

The optimal policy with parameter uncertainty is then contrasted with the optimal policy without parameter uncertainty (i.e., one sample with the parameters at their expected values). In this example, abatement considering parameter uncertainty (Fig. 5, solid line) is *lower* than abatement without parameter uncertainty (Fig. 5, dashed line). For example, optimal abatement without parameter uncertainty is roughly 38% in 2095. Considering parameter uncertainty reduces this value to roughly 29%. The decrease of optimal abatement with increasing parameter uncertainty is in contrast to the change one would perhaps expect intuitively, if risk aversion were the dominant effect. A consumer is risk averse if a certain outcome is preferred over an uncertain outcome with the same expected value. The observed effect of risk aversion is represented in our model by the logarithmic utility function (Eq. (1)), which is one of the many nonlinearities relating climate policy to the objective function. Increased uncertainty would increase optimal abatement if (i) risk aversion were the dominant nonlinearity in the model, and (ii) increased abatement would decrease the variance of the per capita consumptions across the different states of the world. The fact that increased parameter uncertainty does not cause an increase in optimal abatements indicates that these two conditions are not met in our simple model. For a more detailed discussion of this aspect, see [27,28].

The decrease of optimal abatement with increasing parameter uncertainty can be explained intuitively by the threshold properties and the benefit-cost reasoning in the model. As seen in the sensitivity analyses shown in Figs. 2 and 4, optimal abatement is high for the expected parameter values but falls off for all but one tail of the distributions. Specifically, low thermohaline circulation specific damages and high *and* low climate sensitivities result in lower abatement levels

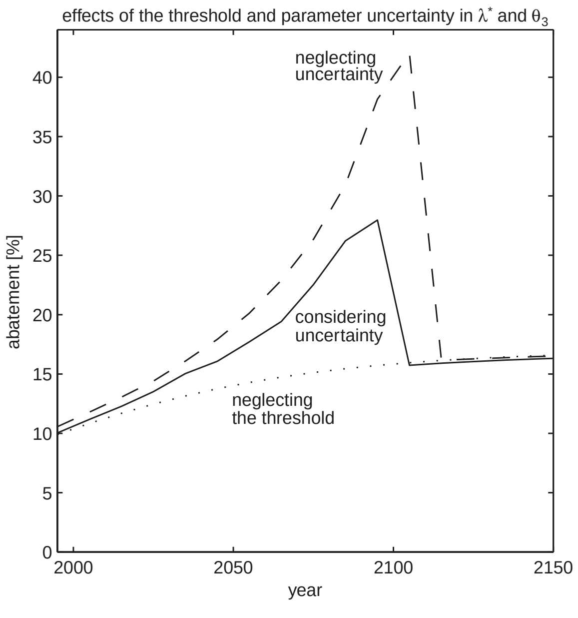

Fig. 5. Decision analysis under parameter uncertainty in θ_3 and λ^*. The solid line shows the optimal policy considering parameter uncertainty. The dashed line shows the optimal policy without parameter uncertainty (i.e., θ_3 and λ^* at their expected values). The dotted line shows the optimal policy neglecting the THC threshold entirely (i.e., $\theta_3 = 0\%$ and λ^* at the expected value).

compared to a policy based on the expected parameter values. Consideration of parameter uncertainty gives more weight to the tails of the parameter distributions and results in this example in lower abatement levels. Note that although the probability weighted average of optimal abatement under certainty (Figs. 2 and 4) gives in this case an intuitive explanation for the observed effects of parameter uncertainty, it is not equivalent to the policy derived from an optimization of expected utility.

The fact that parameter uncertainty reduces the optimal abatement should not be confused with a recommendation for no action. As shown in Fig. 5, the policy that optimizes expected utility considering the threshold (solid line) recommends considerably higher CO_2 abatement than the policy neglecting the threshold (dotted line).

3.5. *Effects of future learning*

The uncertainties about the natural system will likely decline in the long run because we can observe future climates [26]. We explore the effects of uncertainty and learning by analyzing three scenarios: (i) a certainty scenario, (ii) an uncertainty scenario with learning, and (iii) an uncertainty scenario without learning. In the certainty scenario, the climate sensitivity is perfectly known and equal to 3.6°C. In the uncertainty scenario with learning, the climate sensitivity is only imperfectly known in the near term but perfect knowledge of the true climate sensitivity is gained in 2085. The long period of uncertainty reflects the long-time scale Bayesian learning about the climate sensitivity takes in the model simulations of Kelly and Kolstad [26]. The parameter uncertainty about the climate sensitivity is represented by a rather coarse subsample of the estimated probability density function (Fig. 1). We use a Latin Hypercube subsampling procedure [35] with 3 samples resulting in a low, a medium, and a high climate sensitivity of 2.4°C, 3.6°C, and 4.8°C with 25%, 50%, and 25% probability, respectively. Finally, the uncertainty scenario without learning assumes that learning does not occur within the analyzed time horizon of 150 years. (We assume a very late learning in 2185. To assume no learning at all from the observational record would arguably be unrealistic [26].) All scenarios assume a relatively high threshold-specific damage of 3% to simplify the discussion. Note that the learning about the climate sensitivity implies also some learning about the future climate damages. This is because the climate sensitivity affects the temperature changes, which in turn affect the climate damages (Eq. (5)).

These simplified scenarios indicate that uncertainty and learning about the climate sensitivity (and in turn, about climate thresholds) might affect the near-term policy considerably (Fig. 6). In this example, the effect of parameter uncertainty with future learning is to decrease the near-term optimal abatement relative to the situation with certainty (similar to the case of uncertainty without learning shown in Fig. 5). As long as the true value of the climate sensitivity is unknown (i.e., before the year 2085, indicated by the arrow in Fig. 6) the abatement levels with uncertainty are lower than the abatement levels under certainty and the climate sensitivity at the expected value. The optimal near-term abatement for the case of uncertainty without learning in the analyzed time-horizon is also below the policy with certainty (Fig. 6) (cf. [12]).

The situation changes considerably, once the true states of the world are revealed in the scenario of uncertainty with learning. For a median climate sensitivity, the abatement levels increase sharply above the certainty scenario. This increase in CO_2 abatement compensates for the

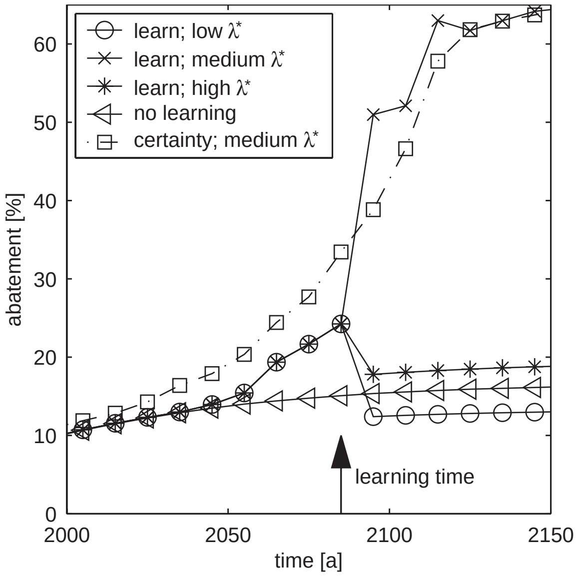

Fig. 6. Decision analysis under parameter uncertainty in climate sensitivity with a rudimentary representation of learning. Shown are the optimal abatement levels as a function of time for a climate sensitivity of 2.4 (circles), 3.6 (crosses), and 4.8 (stars)°C (with 25%, 50%, and 25% probability, respectively). Learning about the climate system is represented as a collapse of uncertainty in the year 2085. Shown for comparison are also the policies for the certainty scenario (squares) and for uncertainty without learning scenario. The shown scenarios are for a threshold-specific damage of 3%. See text for details.

lower abatement levels in the past and the THC is preserved. For a low climate sensitivity, abatement levels decrease relative to the certainty case. This is because the high critical $P_{CO_{2,e}}$ levels (Fig. 1) are not reached in the considered time horizon and the THC is not in danger. For a high climate sensitivity, the abatement levels decrease relative to the certainty case. The explanation for this perhaps counterintuitive results is that the 3% threshold-specific damage is too small to justify the relatively low $P_{CO_{2,e}}$ stabilization levels in the underlying utilitarian framework (similar to the situation depicted in Fig. 4, upper half) and the THC collapses.

We discuss the effects of uncertainty and learning on optimal CO_2 abatement in terms of the expected economic value of information (VOI) and the quasi-option value. The VOI is the increase in expected utility as information becomes available (and uncertainty decreases) [20,33,41]. As shown in Fig. 6, the interplay of the irreversible climate damage of a THC collapse and future learning can affect the optimal strategy. Actions to delay this climate irreversibility beyond the time when learning occurs can increase the VOI and can carry a quasi-option value [4,18,19]. As discussed in [18, p. 34], the quasi-option value is the "expected value of future information conditional on preservation". Preservation, in the example of the considered environmental irreversibility, would be the preservation of the THC. Delaying or avoiding a THC collapse maintains flexibility as it keeps the option open to reevaluate climate policy after the information becomes available. Excellent discussions of the concepts of the value of information and the quasi-option value can be found in [15,16,18].

There are two main irreversibilities in the model that affect optimal CO_2 abatement in opposite directions: (i) the climate irreversibility of a THC collapse (discussed above) and (ii) consumption forgone by unrecoverable investment in CO_2 abatement. Consider the quasi-option value associated with the environmental irreversibility. Increasing CO_2 abatement reduces the risk of the irreversible damages caused by a THC collapse and increases the quasi-option value associated with the environmental irreversibility. The environmental quasi-option value hence favors higher CO_2 abatement. Irrecoverable investments constitute a second irreversibility and give rise to a quasi-option value associated with lower CO_2 abatement [17,27,60]. This is because reducing CO_2 abatement reduces the risk of unrecoverable investments into CO_2 abatement and increases the quasi-option value associated with the investment irreversibility.

The question whether uncertainty with learning acts to increase or decreases optimal CO_2 abatement in the near term relative to the case without uncertainty is a function of the relative importance of the environmental and investment irreversibilities in the specific situation. Previous analyses indicate that uncertainty and learning can either increase or decrease the optimal CO_2 abatement in the near-term [16,17,60]. In the specific example shown in Fig. 6, uncertainty and learning act to decrease the optimal CO_2 abatement in the near term relative to the case without uncertainty. This result suggests that the effect of the investment irreversibility dominates the effect of the environmental irreversibility in our specific example.

4. Comparison to previous studies

Our study differs from previous studies of optimal CO_2 abatement mostly in the representation of the climate threshold and the effects of parameter uncertainty on optimal CO_2 abatement. Studies using an absolutely binding constraint to represent climate thresholds [14,24,58] may arrive at different policy recommendations than our study, which considers finite threshold-specific damages. Studies using a binding constraint do not consider the option of just delaying the damage (i.e., the entire lower half of Fig. 2 would be identical and equal to the case of $\theta_3 = 0\%$). Further, using an absolutely binding constraint implicitly adopts an infinite threshold specific damage—arguably a questionable assumption.

We show that uncertainty about the climate sensitivity and threshold-specific climate damages acts to decrease the near term optimal abatement (and in turn carbon taxes) compared to the certainty scenario (Fig. 5). Our results are similar to the findings of Kolstad [27], where uncertainty and learning about smooth climate damages has the same effect. In contrast, Roughgarden and Schneider [48] conclude that uncertainty about climate damages acts to increase optimal carbon taxes. We hypothesize that this discrepancy is the result of a different decision criterion adopted by Roughgarden and Schneider [48]. Our analysis is based on expected utility maximization. In contrast, Roughgarden and Schneider [48] analyze optimal policies in various scenarios, where each scenario assumes perfect information. Optimal policy is then characterized by the probability distribution of the optimal policies for any given scenario (characterized, for example, by the mean and the median of the carbon taxes). This alternative approach has the advantage of being numerically efficient but it is unclear whether their estimated mean or the median carbon tax should be implemented (a discrepancy as large as a factor of 1.8).

5. Caveats

Integrated assessment models of climate change are nothing more than thinking tools to analyze the interactions between the economic and the environmental system in a highly simplified, but transparent, way. As a result, the resulting policy recommendations should be taken with a grain of salt and are subject to numerous caveats. For example, our results are conditional on the chosen decision criterion (cf. [13]), the assumed rate of technological improvement (cf. [33]), the (so far very uncertain) probability density function of the threshold-specific damages, the choice of the social rate of time preference, and the representation of technological inertia (cf. [14]).

Besides these so far unexplored sensitivities, our study could certainly be improved by more refined representations of nonmarket values [6,39], uncertainty and learning, or long-term discounting. For example, our assumption that the uncertainty about the THC is resolved once the climate sensitivity is perfectly known neglects the uncertain response of precipitation patterns which additionally affects the THC [50]. A second area for improvement is the simple treatment of discounting. We adopted a constant rate of social time preference (ρ) of 3% per year from the DICE94 model [39] (which yields a monetary discount rate roughly consistent with present observations). The application of a constant discount rate to the far-distant future is problematic, as it would suggest that anything beyond a time horizon of several centuries is insignificant. Weitzman [61], for example, suggests that the far-distant future should be discounted at the lowest possible rate, basically because the lowest possible discount rate dominates the cost-benefit analysis in an expected value sense. Newell and Pizer [37] explore this issues and estimate the effective discount rate implied by the historic record. The long-term discount rates estimated by Newell and Pizer [37] are much lower than the discount rates in our model. One interesting question is how strongly our conclusions would change if we were to adopt the more realistic discount rates suggested by Newell and Pizer [37]. One might hypothesize that the general effect of decreasing the discount rate would be to favor higher CO_2 abatement levels and the preservation of the ocean circulation. This hypothesis is based on the properties of the underlying cost-benefit analysis that compares present costs of CO_2 abatement with future benefits of decreased climate change. Lower discount rates give more weight to future events and hence justify higher investments in the present.

6. Conclusions

We explore the effects of uncertainty and learning about a climate threshold (an uncertain ocean thermohaline circulation collapse) in an economic optimal growth model. Our study demonstrates three important points. First, uncertain thresholds can be realistically incorporated into optimal growth models and can change the policy recommendations considerably. Second, uncertainty about the climate sensitivity and the economic damages caused by climate thresholds can act to decrease optimal CO_2 abatement in the near-term. This effect of uncertainty can occur with or without future learning. Third, and most importantly, surprisingly small and uncertain damages associated with a realistic climate threshold might well render significant reductions of future CO_2 emissions a utility-maximizing policy choice.

Acknowledgments

We thank the participants of the Yale/NBER/IIASA workshop on potential catastrophic impacts of climate change for helpful comments on an earlier version of this paper. We are especially grateful to H. Dowlatabadi, H.D. Jacoby, E. Naevdal, W.D. Nordhaus, R. Mendelsohn, F.M.M. Morel, C.D. Kolstad, and R.S.J. Tol for helpful discussions. The comments of Dr. Deacon, as well as two anonymous reviewers improved the presentation. Financial support from the Princeton Environmental Institute, the Carbon Mitigation Initiative at Princeton University, and the Environmental Consortium at the Pennsylvania State University is gratefully acknowledged. Special thanks go to Dr. Herriges for excellent editorial handling. Any remaining mistakes are solely the responsibility of the authors.

Appendix. List of symbols

BAU	business as usual emissions
β	atmospheric airborne fraction for carbon dioxide
b_1	model parameter in the cost function
b_2	model parameter in the cost function
CO_2	carbon dioxide
c	per capita consumption per year
D	fractional economic damages caused by climate change
δ_M	rate constant for carbon dioxide sink
E	anthropogenic carbon dioxide emissions
L	population size
μ	CO_2 abatement in percent relative to the uncontrolled scenario
M	atmospheric carbon stock
λ^*	climate sensitivity
$P_{CO_{2,e}}$	equivalent CO_2 concentration
$P_{CO_{2,e,crit}}$	critical equivalent CO_2 concentration
Q	production
ρ	pure rate of social time preference
T	change in globally averaged surface temperature
t	time
t_o	start of considered time horizon
t^*	end of considered time horizon
θ_1	model parameter in the damage function
θ_2	model parameter in the damage function
θ_3	threshold specific fractional economic damage
TC	fractional CO_2 emissions abatement cost
U	flow of utility
μ	abatement
U^*	discounted sum of utility

References

[1] D.L. Albritton, L.G. Meira Filho (and 56 co-authors), Climate change 2000: the scientific basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change (IPCC), Technical Summary Chapter, Cambridge University Press, Cambridge, 2001.

[2] N.M. Alemdar, S. Ozyildirim, A genetic game of trade, growth and externalities, J. Econ. Dynam. Control 22 (1998) 811–832.

[3] R.B. Alley, J. Marotzke, W.D. Nordhaus, J.T. Overpeck, D.M. Peteet, R.A. Pielke, R.T. Pierrehumbert, P.B. Rhines, T.F. Stocker, L.D. Talley, J.M. Wallace, Abrupt climate change, Science 299 (2003) 2005–2010.

[4] R.K. Arrow, A.C. Fisher, Environmental preservation, uncertainty, and irreversiblity, Quart. J. Econ. 88 (1974) 312–319.

[5] A. Baumgartner, E. Reichel, The World Water Balance, Elsevier, Amsterdam, 1975.

[6] D.F. Bradford, On the uses of benefit-cost reasoning in choosing policy toward global climate change, in: P.R. Portney, J.P. Weyant (Eds.), Discounting and Intergenerational Equity, Resources for the Future, Washington, DC 20036-1400, USA, 1999, pp. 37–44.

[7] W.S. Broecker, The great ocean conveyor, Oceanography 4 (1991) 79–89.

[8] W.S. Broecker, Thermohaline circulation, the Achilles heel of our climate system: will man made CO_2 upset the current balance?, Science 278 (1997) 1582–1588.

[9] H.-P. Chao, Managing the risk of global climate catastrophe: an uncertainty analysis, Risk Anal. 15 (1995) 69–78.

[10] M. Claussen, On multiple solutions of the atmosphere–vegetation system in present-day climate, Global Change Biol. 4 (1998) 549–559.

[11] M.L. Crooper, Regulating activities with catastrophic effects, J. Environ. Econ. Manag. 3 (1976) 1–15.

[12] C. Deutsch, M.H. Hall, D.F. Bradford, K. Keller, Detecting a potential collapse of the north atlantic thermohaline circulation before a point of no return: implications for the design of an ocean observation system, Paper presented at the EMF Summer Workshop "Climate Change Impacts and Integrated Assessment VIII", Snowmass, Colorado, July 27–August 7, 2002. (available at: http://www.geosc.psu.edu/˜kkeller).

[13] H. Dowlatabadi, M.G. Morgan, A model framework for integrated studies of the climate problem, Energy Pol. 21 (1993) 209–221.

[14] M.H. Duong, M. Grubb, J.-C. Hourcade, Influence of socioeconomic inertia and uncertainty on optimal CO_2-emission abatement, Nature 390 (1997) 270–273.

[15] A.C. Fisher, W.M. Hahnemann, Quasi-option value: some misconceptions dispelled, J. Environ. Econ. Manage. 14 (1987) 183–190.

[16] A.C. Fisher, U. Narain, Global warming, endogenous risks, and irreversibility, Department of Agricultural and Ressource Economics Working Papers, Vol. 908, 2002.

[17] M. Ha-Duong, Quasi-option value and climate policy choices, Energy Econ. 20 (1998) 599–620.

[18] W.M. Hahnemann, Information and the concept of option value, J. Environ. Econ. Manage. 16 (1999) 23–37.

[19] C. Henry, Investment decisions under uncertainty: the "irreversibility effect", Amer. Econ. Rev. 64 (1974) 1006–1012.

[20] S.R. Johnson, M.T. Holt, The value of weather information, in: R.W. Katz, A.H. Murphey (Eds.), Economic Value of Weather and Climate Forecasts, Cambridge University Press, Cambridge, 1997, pp. 75–107.

[21] K. Kacholia, R.A. Reck, Comparison of global climate change simulations for $2 \times CO_2$-induced warming, Climatic Change 35 (1997) 53–69.

[22] A. Kann, J.P. Weyant, Approaches for performing uncertainty analysis in large scale energy/economic policy models, Environ. Model. Assess. 5 (2000) 29–46.

[23] R.K. Kaufmann, Assessing the DICE model: uncertainty associated with the emissions and retention of greenhouse gases, Climatic Change 35 (1997) 435–448.

[24] K. Keller, K. Tan, F.M.M. Morel, D.F. Bradford, Preserving the ocean circulation: implications for climate policy, Climatic Change 47 (2000) 17–43.

[25] D.L. Kelly, C.D. Kolstad, Solving growth models with an environmental sector, Mimeo, University of Miami, 1998.

[26] D.L. Kelly, C.D. Kolstad, Bayesian learning, growth and pollution, J. Econ. Dynam. Control 23 (1999) 491–518.

[27] C.D. Kolstad, Learning and stock effects in envionmental regulation: the case of greenhouse gas emissions, J. Environ. Econ. Manage. 31 (1996) 1–18.
[28] J.-J. Laffont, The Economics of Uncertainty and Information, MIT Press, Cambridge, MA, 1999 (Translated by J.P. Bonin, H. Bonin).
[29] A.S. Lieber, What is the economic value of information about uncertain climate thresholds? Undergraduate Thesis, Department of Economics, Princeton University, Princeton, NJ, 2001.
[30] R. Loulou, A. Kandia, Minimax regret strategies for greenhouse gas abatement: methodology and application, GERAD discussion paper G-97-32, 1997.
[31] S. Manabe, R.J. Stouffer, Multiple-century response of a coupled ocean–atmosphere model to an increase of atmospheric carbon dioxide, J. Climate 7 (1994) 5–23.
[32] S. Manabe, R.J. Stouffer, M.J. Spelman, K. Bryan, Transient responses of a coupled ocean–atmosphere model to gradual changes of atmospheric CO_2. Part I: annual mean response, J. Climate 4 (1991) 785–818.
[33] A.S. Manne, R. Richels, Buying greenhouse insurance, Energy Pol. 19 (1991) 543–552.
[34] N. Mattsson, C.-L. Wene, Assessing new energy technologies using an energy system model with endogenized experience curves, Int. J. Energy Res. 21 (1997) 385–393.
[35] M.D. McKay, R.J. Beckman, W.J. Conover, Comparison of 3 methods for selecting values of input variables in the analysis of output from a computer code, Technometrics 21 (1979) 239–245.
[36] C.M. Moles, A.S. Lieber, J.R. Banga, K. Keller, Efficient global optimization of nonconvex and nondifferentiable climate control problems. in: J.M. Benitez, O. Cordon, F. Hoffmann, R. Roy (Eds.), Advances in Soft Computing: Engineering Design and Manufacturing, Springer, Heidelberg, 2003, pp. 331–342.
[37] R.G. Newell, W.A. Pizer, Uncertain discount rates in climate policy analysis, Energy Policy 32 (4) (2004) 519–529.
[38] W.D. Nordhaus, An optimal transition path for controlling greenhouse gases, Science 258 (1992) 1315–1319.
[39] W.D. Nordhaus, Managing the Global Commons: The Economics of Climate Change, MIT Press, Cambridge, Massachusetts, 1994.
[40] W.D. Nordhaus, J. Boyer, Warming the World: Economic Models of global Warming, MIT Press, Cambridge, MA, 2000.
[41] W.D. Nordhaus, D. Popp, What is the value of scientific knowledge? An application to global warming using the PRICE model, Energy J. 18 (1997) 1–45.
[42] M. Oppenheimer, Global warming and the stability of the West Antarctic ice sheet, Nature 393 (1998) 322–325.
[43] G. Petschel-Held, H.J. Schnellnhuber, T. Bruckner, F L. Toth, The tolerable windows approach: theoretical and methodological foundations, Climatic Change 41 (1999) 303–331.
[44] D.A. Pierre, Optimization Theory with Applications, Dover, New York, 1986.
[45] S. Rahmstorf, Rapid transitions of the thermohaline ocean circulation: a modelling perspective. in: F. Abrantes, A. Mix, (Eds.), Reconstructing Ocean History: A Window into the Future, Kluwer Academic Publishers, Dordrecht 1999.
[46] S. Rahmstorf, Shifting seas in the greenhouse?, Nature 399 (1999) 523–524.
[47] S. Rahmstorf, A. Ganopolski, Long-term global warming scenarios computed with an efficient coupled climate model, Climatic Change 43 (1999) 353–367.
[48] T. Roughgarden, S.H. Schneider, Climate change policy: quantifying uncertainties for damages and optimal carbon taxes, Energy Pol. 27 (1999) 415–429.
[49] A. Schiller, U. Mikolajewicz, R. Voss, The stability of the North Atlantic thermohaline circulation in a coupled ocean–atmosphere general circulation model, Climate Dynam. 13 (1997) 325–347.
[50] A. Schmittner, T.F. Stocker, The stability of the thermohaline circulation in global warming experiments, J. Climate 12 (1999) 1117–1133.
[51] T.F. Stocker, A. Schmittner, Influence of CO_2 emission rates on the stability of the thermohaline circulation, Nature 388 (1997) 862–865.
[52] R. Storn, K. Prince, Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces, J. Global Optimiz. 11 (1997) 341–359.
[53] R.J. Stouffer, S. Manabe, Response of a coupled ocean–atmosphere model to increasing atmospheric carbon dioxide: sensitivity to the rate of increase, J. Climate 12 (1999) 2224–2237.

[54] R.S.J. Tol, The damage costs of climate change—a note on tangibles and intangibles applied to DICE, Energy Pol. 22 (1994) 436–438.
[55] R.S.J. Tol, On the optimal control of carbon dioxide emissions: an application of FUND, Environ Model. Assess. 2 (1997) 151–163.
[56] R.S.J. Tol, Potential slowdown of the thermohaline circulation and climate policy, Discussion Paper DS98/06, Institute for Environmental Studies Vrije Universiteit Amsterdam, 1998.
[57] R.S.J. Tol, A.F. de Vos, A Bayesian statistical analysis of the enhanced greenhouse effect, Climatic Change 38 (1998) 87–112.
[58] F.L. Toth, T. Bruckner, H.-F. Fussel, M. Leimbach, P.-H. Gerhard, H. Schnellnhuber, The tolerable windows approach to integrated assessments, in: Proceedings of the IPCC Asia-Pacific Workshop on integrated assessment models [IAMS], Paris, France, 1997.
[59] M.P. Tucci, A note on global optimization in adaptive control, econometrics and macroeconomics, J. Econ. Dynam. Control 26 (2002) 1739–1764.
[60] M. Webster, The curious role of "learning" in climate policy: should we wait for more data?, Energy J. 23 (2002) 97–119.
[61] M.L. Weitzman, Why the far-distant future should be discounted at its lowest possible rate, J. Environ. Econ. Manage. 36 (1998) 201–208.
[62] T.M.L. Wigley, P.D. Jones, S.C.B. Raper, The observed global warming record: what does it tell us?, Proc. Nat. Acad. Sci. 94 (1997) 8314–8320.