Climatic Change manuscript No.
(will be inserted by the editor)

What are robust strategies in the face of uncertain climate threshold responses?

David McInerney · Robert Lempert · Klaus Keller

Submitted to Climatic Change
29 October 2009

Abstract We use an integrated assessment model of climate change to analyze how the implementation of different decision-making criteria affects preferred investments into climate control, the distribution of outcomes, the robustness of the strategies, and the economic value of information. Specifically, we modify the Dynamic Integrated model of Climate and the Economy (DICE-07) to include a simple representation of a climate threshold response, parametric uncertainty, structural uncertainty, and learning. Economic analyses of climate change strategies typically adopt the expected utility maximization (EUM) framework. We compare EUM with two decision criteria adopted from the finance literature, namely Limited Degree of Confidence (LDC) and Safety First (SF). Both criteria increase the relative weight of the performance under the worst-case scenarios compared to EUM. The tradeoff between maximizing expected utility and minimizing the worst case outcomes is shown to be equivalent for the LDC

David McInerney
Department of the Geophysical Sciences
5734 S. Ellis Ave.
HGS 571
Chicago, IL 60637 USA
Fax: (773) 702-9505
E-mail: dmcinern@geosci.uchicago.edu

Robert Lempert
Frederick S.Pardee Center for Longer Range Global Policy and the Future Human Condition
RAND
1776 Main St.
Santa Monica, CA 90401 USA
Phone: (310) 393-0411 x6217
E-mail: lempert@rand.org

Klaus Keller
436 Deike Building
Department of Geosciences
Penn State
University Park, PA, 16802 USA
Phone: (814) 865-6718
Email: klaus@psu.edu

and SF criteria and identical preferred abatement strategies can be obtained using these two criteria.

We show that the LDC and SF criteria provide a computationally feasible foundation for identifying greenhouse gas mitigation strategies that may prove more robust than those identified by the EUM criterion. Robustness is here defined as trading a small decrease in a strategy's expected performance for a significant increase in a strategy's performance in the worst cases. More robust strategies show higher near-term investments in carbon dioxide emissions abatement and a higher economic value of climate information.

Keywords climate change strategy · decision-making under uncertainty · climate thresholds.

1 Introduction

Anthropogenic greenhouse gas emissions are projected to cause climatic changes with non-trivial and likely overall negative impacts on human welfare (Nordhaus 2008). Reducing anthropogenic climate forcing is possible, but requires significant investments. Determining a climate change strategy requires choosing a level of investment that appropriately balances the costs and benefits of reducing emissions. The fact that current estimates of the impacts of climate change and the costs of climate mitigation are deeply uncertain imposes non-trivial methodological challenges (Keller et al 2008; Lempert 2002). Deep uncertainty (sometimes also referred to as ambiguity, Knightian uncertainty, or vagueness) refers to a situation where the decision-makers do not know or cannot agree on a single probability density function (PDF) of the outcomes (Lempert et al 2003; Lempert and Collins 2007). Nonetheless, one prominent analytic approach frames climate change decision-making as an inter-temporal optimization problem where policy makers seek to maximize the discounted expected utility of current and future generations (Nordhaus 2008; Keller et al 2004). This expected utility maximization (EUM) framing poses at least three potentially important problems: (i) EUM poorly describes the actual decision rules people often use under conditions of deep uncertainty, (ii) implementing EUM under conditions of deep uncertainty raises important methodological challenges, for instance how to aggregate differing expert estimates, and (iii) the specification of EUM strategies can be highly sensitive to misspecificaion of low-probability high-impact events. This paper explores alternative decision criteria that may help overcome these challenges.

Most Integrated Assessment Models of climate change (IAMs) adopt the classic decision-making framework of expected utility maximization (EUM) (Ackerman et al 2009; Bretteville Froyn 2005) that identifies an optimum strategy contingent on a single best estimate joint probability distribution over the uncertain input parameters of the model. EUM has at least two important advantages. The approach rests on a solid theoretical foundation built on a small number of intuitive axioms. Computationally the approach is also relatively straightforward to implement (Quiggin 2008; McInerney and Keller 2008; Tol 2003). However, as recently emphasized by the U.S. Climate Change Science Program, for both theoretical and practical reasons, there are limits to the applicability and usefulness of classic decision analysis to climate-related problems (Morgan et al 2009, p.25).

As one challenge, under conditions of deep uncertainty the EUM framework has poor descriptive power (Ellsberg 1961). Empirical studies (Budescu et al 2002; Du and Budescu 2005) show that most decision makers are sensitive to the (im)precision of the decision parameters. Typically, they are vagueness (ambiguity) averse and are willing to pay a premium to avoid it, although is some cases they like small degrees of imprecision. Decision problems involving long-term climate change strategies are arguably better described as deeply uncertain (as opposed to just uncertain) (Keller et al 2008; Lempert 2002; Welsch 1995). For example, risk estimates of potential low-probability high-impact events such as a shutdown of the North Atlantic thermohaline circulation hinges critically on the highly diverging expert prior (Zickfeld et al 2007; Keller et al 2000; Alley et al 2002; Link and Tol 2004).

Analyzing climate change strategies using an expected utility maximization framework also presents methodological challenges, such as the need to aggregate different PDFs, a step typically requiring strong assumptions (Keith 1996). One typical approach to this aggregation step is to neglect the aspects of deep uncertainty and pick just one example PDF (Nordhaus 2008; Keller et al 2004). Neglecting the effects of deep uncertainty can lead to inaccurate and insensitive predictions.

Finally, the expected utility maximization framework can suggest strategies that prove vulnerable to misestimation of deeply uncertain low-probability high-impact events (Peterson et al 2003; Keller et al 2004; Weitzman 2009). This vulnerability can arise, for example, when the EUM strategy is located at a narrow peak of the utility function that is surrounded by a deep valley (a situation sometimes referred to as "dancing on the top of a needle"). In this case, an alternative strategy that is located at a slightly lower but much broader peak of the utility has the advantage of trading a small amount of utility in the expected case for a considerable increase in utility if the estimate of the expected case is slightly wrong.

Previous analyses of decision-making under deep uncertainty have typically addressed these two problems by employing alternative decision-making criterion than EUM (Borsuk and Tomassini 2005; Lempert and Collins 2007; Lange 2003; McInerney and Keller 2008), by adopting alternative representations of uncertainty than a single joint probability distribution (Kriegler and Held 2005; Hall et al 2007; Lempert and Collins 2007), or both. Such analyses have broken important new ground, but they typically analyze highly stylized problems. For example, Lempert and Collins (2007) use a simple model describing a lake that can abruptly turn eutrophic if pollution concentrations exceed some unknown threshold in order to compare the implications of alternative decision making criteria. The study compares EUM and three types of robustness criteria, representing deep uncertainty with a set of alternative probability distributions. The study finds robust strategies preferable to expected utility maximization when deep uncertainty exists about potentially catastrophic impacts and when decision makers have a sufficiently rich menu of decision options to allow them to find one that is robust. Lange (2003) also analyzes the effects of the limited degree of confidence decision criterion – which we will employ below – on the preferred climate change strategies, but does not consider the effects of a potential climate threshold response.

Our study expands on previous work in several ways. It explores alternative decision criteria to EUM adapted from the finance literature, using the particularly stressing case of an IAM with an uncertain, potentially abrupt climate threshold response. It considers both parametric and structural uncertainty, where the former involves uncertainty about input parameters to the model and the latter involves uncertainty about the model itself. The study considers the effects of learning about the threshold response. The study also demonstrates computational methods that can implement these new decision criteria on a climate change IAM with abrupt changes.

With these resources, the study addresses three main questions. First, what are the effects of considering a new set of decision-criteria typically used in the finance literature on the preferred climate change strategy? Second, what is the economic value of information for these decision-criteria? Third, can relatively simple modifications of the decision-criteria be used to identify more robust decisions in a computationally efficient way?

We analyze these questions by modifying the Dynamic Integrated model of Climate and the Economy (DICE) to include the potential for a climatic threshold response. These impacts, representing those that might occur from the shutdown of the North Atlantic Meridional Overturning Circulation (MOC), are considered to be deeply uncertainly (Zickfeld et al 2007; Keller et al 2008). We consider three decision-criteria that treat the uncertainty about the potential for large climate impacts in different ways. As the first decision-criterion we analyze expected utility maximization. Given the potentially severe problems associated with EUM discussed above, we analyze two additional decision-criteria that provide avenues to (i) represent the effects of the deep uncertainty on the decision-making process and (ii) offer the possibility of improving the coherence between the decision-making criterion adopted by real decision-makers and the decision-making criterion used in Integrated Assessments Models of climate change (Budescu et al 2002; Du and Budescu 2005).

The two alternative decision criteria, Safety First (SF) and Limited Degree of Confidence (LDC), both balance the goal of maximizing the expected utility with minimizing the worst-case performance (Aaheim and Bretteville 2001; Bretteville Froyn 2005). The LDC criterion maximizes a weighted average of the expected utility and the utility in the worst-case futures. The SF criterion maximizes the expected utility with the constraint that the utility in the worst cases exceed some prescribed threshold. The weighting in the LDC criteria can be interpreted as a measure of decision makers confidence in the best estimate probability distribution used to calculate the expected value. The constraint used by SF criterion can be interpreted as an ethical requirement that one should maximize utility only after guaranteeing some minimum level of utility (Rawls 1971). SF also has similarities with criterion from the finance literature that combine expected utility with conditional value at risk to help a firm, for example, maximize profits while holding the probability of bankruptcy below some threshold. By placing weight on the performance in the worst-case outcomes, the LDC criterion is similar to decision rules described by Hurwicz (1951) and (Ellsberg 1961). The former's optimism–pessimism criterion prescribes a weighted combination of the best case and worst-case outcome. Ellsberg (1961) suggests that when making decisions under what he calls conditions of ambiguity, people will use as their decision criteria a weighted combination of the expected utility contingent on the best-estimate distribution and of the expected utility of the worst case distribution, where the weighting is a function of decision makers confidence in the available data.

These three criteria (EUM, SF and LDC) can generate a wide range of recommendations for emissions reduction policies depending on the parameters used in the SF and LDC decision-criteria. The Safety First and Limited Degree of Confidence criteria can recommend higher preferred investments in emissions abatement than the EUM criterion, depending, in the former, on the prescribed threshold for worst cases and, in the latter, how much confidence decision makers express in the best estimate distribution. For the SF and LDC criteria we map out the tradeoff between maximizing expected utility and minimizing the effect of the worst cases for a range of parameter values. We find that the tradeoff curves for the SF and LDC criteria are equivalent within numerical accuracy. Moreover, the LDC and SF strategies that are at the same location in this tradeoff curve show the same abatement strategy (again within numerical accuracy).

We then examine the robustness of various strategies. Specifically, we analyze how well strategies that emphasize expected utility maximization perform in the worst cases and how well strategies that emphasize worst-cases performance affect the expected performance. We show how applying the two alternative decision-making criteria help identify strategies decision makers may find more robust given the range of impacts considered. Finally, we evaluate how the choice of decision-criterion affects the choice of climate strategy given the potential for future learning, as well as the expected economic value of climate information. We demonstrate that constraining the worst cases (as implemented in the Safety First decision criterion) can significantly increase the expected economic value of climate information. Finally, we show that the expected value of information can be considerably higher under the SF criterion, compared to the EUM criterion.

2 The integrated assessment model

We use a recent version of the Dynamic Integrated model of Climate and the Economy (DICE-07) (Nordhaus 2007a). DICE is a simple and transparent integrated assessment model that has been widely-used (Popp 2004; Keller et al 2004; Nordhaus 2007b; Tol 1994) and is well-documented (Nordhaus 1992, 1994, 2008). We modify this model to (i) incorporate a representation of a potential climate threshold response, (ii) allow for parametric uncertainty, (iii) consider the potential to learn, and (iv) allow different decision-criteria to be considered. We describe the model and the modifications below.

2.1 The DICE-07 model

The DICE-07 model couples a Ramsey type economic growth model (Ramsey 1928) to simple models of the carbon cycle, the climate system, economic damages due to climate change, and the costs of reducing carbon dioxide (CO_2) emissions. The objective of the decision-maker is to determine optimal trajectories over time for investments in capital stock and investments in reducing CO_2 emissions. Specifically, the goal of the stylized decision-maker is to maximize an objective function. For the standard DICE-07 model, this objective function is defined by the discounted sum of utility:

$$W = \sum_{t=t_0}^{t=t_{N-1}} U\left[c(t_n), L(t_n)\right] R(t_n)\,, \tag{1}$$

where $t_n = t_0 + \Delta t$, $t_0 = 2005$ is the initial time, $t_{N-1} = 2595$ is the considered time horizon, and $\Delta t = 10$ years is the time increment. The notation so far assumes certainty in the values of parameters. We return to the three considered extensions of this objective function to allow uncertainty below (Section 2.4).

The first term in this equation, U, is the stream of utility. U is a function of the per capita consumption, $c(t)$, the exogenously specified population, $L(t)$, and the marginal elasticity of utility, α, according to

$$U[c(t), L(t)] = L(t)\left(\frac{[c(t)]^{1-\alpha} - 1}{1-\alpha}\right). \tag{2}$$

The second term in Equation (1) is the discount factor, $R(t)$, which aggregates the vector of the flows of utility over time into a single number by applying a pure rate of social time preference, ρ, according to:

$$R(t) = (1+\rho)^{t_0 - t_n}. \tag{3}$$

The values for α and ρ are chosen by Nordhaus (2008) to result in monetary discount rates that are roughly consistent with observations. We adopt this descriptive approach here but note the thorny methodological and ethical issues (Nordhaus 2008; Tol 1998).

The relationships between the modules representing the economic system, the carbon cycle, climate system, CO_2 abatement costs, and climate change damages are outlined below. A detailed mathematical description of the governing equations and parameter values can be found in Nordhaus (2008) and McInerney and Keller (2008).

The economic module calculates the per capita consumption from the economic output, net the impacts investment in capital stock, damages due to climate change, cost in CO_2 abatement, and capital depreciation. The gross output is calculated from a constant returns to scale Cobb-Douglas production function in technology, capital and labor. The levels of technology and labour are specified exogenously. Investment increases capital stock, which then depreciates over time. The production function is modified by including factors for emission reduction costs and climate induced damages, discussed below.

Economic output is divided into investment in capital stock and consumption. Per capita consumption is defined as consumption per population. Increasing consumption at a given time raises the present flow of utility, while increasing investment into capital stock enhances future economic output, thereby allowing for greater future utility of consumption.

The economic and climate components of DICE-07 are linked through anthropogenic CO_2 emissions (which affect the climate system and the climate change damages) and the climate change damages (which affect the per capita consumption). Absent CO_2 abatement, emissions and world output are linearly related through the exogenously specified carbon intensity. In addition to industrial CO_2 emissions, DICE-07 specifies emissions due to change in land use. CO_2 emissions to the atmosphere increase atmospheric CO_2 stock. Atmospheric CO_2 in excess of the pre-industrial level is removed from the atmosphere following a linear three-box carbon cycle module, resolving the atmosphere, a fast mixing reservoir consisting of the upper ocean and biosphere, and the deep ocean.

Atmospheric CO_2 concentrations in excess of the pre-industrial level cause an increase in radiative forcing according to a logarithmic relationship, while radiative forcing from additional sources (*e.g.* methane and aerosols) are specified exogenously. The climate in DICE-07 is represented by a two-box model consisting of a combined atmosphere and upper ocean, and the deep ocean (Schneider and Thompson 1981). The temperature of the atmosphere and upper ocean is driven by the net surface radiative forcing, while temperature is transferred between the two boxes at a rate proportional to the difference in these temperatures. The deviation in global mean surface temperature from its pre-industrial level (ΔT) is used as a proxy for economic climate change damages, with damages following a quadratic function in ΔT. These damages determine the reduction in economic output in the modified production function.

The primary mechanism for reducing climate induced damages in DICE-07 is through CO_2 abatement, *i.e.* replacing carbon intensive energy with cleaner alternatives. An abatement level, $\mu(t)$, yields industrial emissions that are $(1-\mu(t))$ times the size of uncontrolled emissions. The cost of CO_2 abatement is expressed as a fraction of world output, and is a nonlinear function of $\mu(t)$.

Without investments in CO_2 abatement, the economic output results in industrial CO_2 emissions according to an exogenously defined ratio called the carbon intensity, $\sigma(t)$. These emissions increase atmospheric CO_2 concentration, thereby raising global temperature due to the greenhouse effect. Changes in global mean temperature are then used as a proxy for climate induced economic damages, and these damages reduce future consumption. The CO_2 emissions hence provide a negative feedback on the objective function. With investments in CO_2 abatement, this negative feedback is weakened, but the investment itself constitutes a decrease in the objective function because CO_2 abatement comes at a cost, thus reducing present consumption.

2.2 Representation of a potential climate threshold response

Keller et al (2000) incorporate an MOC threshold into an earlier version of DICE using a simple statistical parameterization of model results reported by Stocker and Schmittner (1997). Specifically, Keller et al (2000) define a critical equivalent CO_2 concentration as a function of climate sensitivity. If this critical equivalent CO_2 is exceeded an irreversible shutdown of the MOC is triggered. Associated with this shutdown are persistent economic damages, θ_3, expressed as a proportion of economic output. This threshold representation is transparent and has been widely used and investigated (Keller et al 2004; Lempert et al 2006; McInerney and Keller 2008). In addition, its integration into an IAM requires few additional computations, therefore enabling extra dimensions of uncertainty to be examined in this study. The limitations of this MOC threshold representation are duly noted (Zickfeld and Bruckner 2008). First, the MOC is not only sensitive to atmospheric CO_2 concentrations, but to the rate of temperature change (Stocker and Schmittner 1997). Second, this analysis neglects uncertainty in the initial strength of the MOC and the sensitivity of fresh-water forcing in the North Atlantic to changes in temperature.

Here we adopt the basic structure of this simple representation but add one additional factor representing the effects of structural model uncertainty. Structural model uncertainty is important in the assessment of a future MOC threshold response (Zickfeld et al 2007). While the results of Stocker and Schmittner (1997) suggest that the MOC is sensitive to climate change, other models report the presence of stabilizing feedbacks that make the MOC virtually insensitive to global warming (Latif et al 2000). To account for this structural uncertainty, we introduce a Bernoulli variable called the MOC sensitivity (p_{MOC}). The MOC will shutdown when the critical threshold is exceeded and the MOC sensitivity is unity.

2.3 Experimental Design

We consider uncertainty in four parameters that previous studies have identified as key drivers of preferred abatement strategies (Nordhaus 1994; Keller et al 2004; Nordhaus 2008; McInerney and Keller 2008). These four factors are: (i) the climate sensitivity, λ^*; (ii) initial growth rate of the carbon intensity, $g_\sigma(2005)$; (iii) the specific damages from crossing the MOC threshold, θ_3; and (iv) the sensitivity of the MOC to increasing atmospheric carbon dioxide concentrations, p_{MOC}. It is important to stress that these four parameters are just a small subset of the relevant uncertainties (Nordhaus 1994; Keller et al 2008).

We adopt the empirical distribution of Andronova and Schlesinger (2001) for climate sensitivity, and adopt subjective estimates for the distributions of $g_\sigma(2005)$, θ_3, and p_{MOC}. The mean value of $g_\sigma(2005)$ follows its original value in DICE-07 (-0.073 per decade). We adopt a uniform PDF for this parameter and assign bounds of 50% and 150% of this mean value, equivalent to -0.11 and -0.04 ($-0.073 \pm 0.5 \times 0.073$). MOC specific damages are deeply uncertain (Nordhaus 1994; Keller et al 2000; Alley et al 2002; Link and Tol 2004). Following Keller et al (2004), McInerney and Keller (2008) and Tol (2003) we assign a mean value of 1.5% GWP to MOC specific damages. To allow for low-probability high-impact outcomes, we use a Weibull distribution with a large standard deviation (3.7% GWP) about this mean value. Finally, we adopt a binary distribution of the MOC sensitivity, resulting in a mean of 0.5. As a result, it is equally likely that the MOC is sensitive or insensitive to anthropogenic climate forcing.

The properties of these distributions are summarized in Table 1 and the cumulative density functions for λ^*, $g_\sigma(2005)$ and θ_3 are displayed in Figures 1. We draw eleven equally likely samples from these three distributions using the technique described in McInerney and Keller (2008). These samples are marked by the star symbols in Figure 1. After drawing eleven samples from each of these three distributions, and allowing the MOC to be either sensitive or insensitive to anthropogenic warming, there are $N_{SOW} = 11 \times 11 \times 11 \times 2 = 2662$ equally likely states of the world (SOWs).

2.4 Alternative decision criteria

We consider the three decision criteria discussed earlier. As summarized in Table 2, these criteria are referred to in the literature as (i) expected utility maximization, (ii) limited degree of confidence, and (iii) safety first, and represent alternative ways to balance between the expected performance over the entire joint probability distribution and the performance in the tails of that distribution.

2.4.1 Expected utility maximization

We first consider the decision criterion of expected utility maximization (EUM). As discussed in the introduction, this is the typical decision criterion adopted by the majority of Integrated Assessment Models of climate change. The EUM criterion maximizes the product of the probability of each considered State of the World (SOW) and the well-being W_i in this this SOW. The well-being in each SOW is calculated using Equation

(1). Given that our sampling results in N_{SOW} equally likely SOWs, this results in:

$$\max\{E[W]\} = \max\left\{\frac{1}{N_{\mathrm{SOW}}}\sum_{i=1}^{N_{\mathrm{SOW}}} W_i\right\} . \tag{4}$$

Note that this typical implementation of the EUM criterion eliminates the aspects of deep uncertainty. For example, the uncertainty about how different estimates of climate sensitivity should be weighted is neglected and a single best estimate is chosen with an implicit weight of unity. This implies that all other estimates have zero weight.

2.4.2 Limited degree of confidence

The second criterion we consider is the limited degree of confidence (LDC) criterion that is discussed, for example, by Aaheim and Bretteville (2001), Bretteville Froyn (2005) and Lange (2003). The LDC criterion maximizes a weighted average of the worst outcome, *i.e.*, the maximin criterion, and the expected utility. To review, the maximin criterion ranks strategies according to their worst case outcomes, and the strategy that maximizes the minimum utility is chosen:

$$\max\{\min[W]\} .$$

The LDC criterion is a combination of EU maximization and maximin according to:

$$\max\{\beta E[W] + (1-\beta)\min[W]\} . \tag{5}$$

The decision-maker's degree-of-confidence in the underlying probability distributions, β, lies between 0 and 1. For $\beta = 0$, this criteria reduces to maximin, while $\beta = 1$ recovers the decision criterion of EU maximization. This criteria balances between the EUM and maximin criteria, depending on the decision makers' degree-of-confidence in the underlying probability distribution. The LDC criterion is similar to the first robustness definition used in Lempert and Collins (2007), which considers a robust strategy to be one that may give up a small amount of optimum performance for less sensitivity to assumptions. However Lempert and Collins (2007) use regret rather than absolute performance to balance expected and worst case performance. Both these approaches are premised on the notion that the distribution of outcomes may be imprecise.

However, the implementation of the LDC criterion as discussed, for example, in Aaheim and Bretteville (2001) and Lange (2003) poses two problems when applied to the situation at hand that do not apply to previous studies. Lange (2003) analyzes a very simplified analytical model with a fixed upper bound for the environmental damages. This means that the worst case is precisely known. In contrast, we do not know the upper bound for the considered problem of abrupt climate change (Keller et al 2008). In addition, for the numerical approach used in thus study, the worst case depends on the considered number of states of the world (*i.e.*, on the sampling resolution) (Tol 2003). For example, the PDFs of θ_3 and λ^* both contain high values with low probabilities (Figure 1), so that as more samples are taken, the maximum sample from each distribution increases and does not converge to an upper limit. If these samples correspond to the worst-case outcomes, the strategy that maximizes the minimum utility depends heavily on sampling resolution (see Tol 2003, for a similar result). Second, a maximin decision criterion only considers the single worst outcome, *i.e.*, it is blind to other poor outcomes. This may be a problem when the worst outcomes correspond to different parts of the parameter space, where improving the worst-case may not necessarily improve the second-worst outcome. The maximin decision criterion does not include probabilistic information by design.

Rather than considering the minimum value of well-being (as in Equation 5) we use an alternative risk metric: the Conditional Value at Risk (CVaR). The CVaR is the expected value of the worst q-th portion of the utility distribution, which we denote $E[W_q]$. W_q is the portion of outcomes below what is known as the Value-at-Risk for q. The CVaR is also known as the Mean Excess Loss or Mean Shortfall, and has been previously applied to risk analysis in the energy sector (Fortin et al 2007), the crop insurance industry (Liu et al 2008), and financial markets (Andersson et al 2001; Alexander et al 2006; Quaranta and Zaffaroni 2008). Using the CVaR metric addresses our above concerns with using the minimum value, since it is less dependent on sampling resolution than the minimum value; as the number of samples increases, $E[W_q]$ converges. We modify the limited degree of confidence criterion as originally discussed, for example, by Lange (2003) by replacing the minimum value metric with the CVaR metric and implementing a typical value of $q = 0.01$ (Andersson et al 2001; Larsen et al 2002; Krokhmal et al 2002):

$$\max\left\{\beta E[W] + (1-\beta)E[W_{0.01}]\right\} . \tag{6}$$

2.4.3 Safety first

The safety first (SF) criterion is an extension of the EUM criterion that imposes an additional constraint on the lower tail of the strategy's performance. The implementation discussed, for example, by Bretteville Froyn (2005) maximizes EU with the additional constraint that the probability that the well-being is less than a critical value W^* is less than a certain value q:

$$\max\left\{E[W]\right\} \quad \text{such that} \quad \Big\{\Pr[W \le W^*] \le q\Big\} . \tag{7}$$

This criterion retains the best estimate probability distribution used by EUM but adds a term sensitive to size of the impacts in the tails of the distribution. This criterion is similar to the reliability constraint analyzed by McInerney and Keller (2008) where the probability of an undesirable event such as an MOC shutdown is constrained to be less than a certain threshold value.

However, similar to the LDC formulation of Bretteville Froyn (2005), this representation of a safety first approach has shortcomings. For example, Equation (7) does not consider the performance of the very worst case outcomes. When $q = 0.01$, we know that 99% of the outcomes will be greater than W^*, but we do not know how extreme the worst 1% of outcomes may be. To address this concern, we will again use the CVaR metric to constrain the worst case outcomes. By doing this, we take into account the extremity of the worst outcomes, as well as consider outcomes that are nearly as bad as the worst case outcomes. We hence modify the safety first criteria to utilize the CVaR metric with $q = 0.01$:

$$\max\left\{E[W]\right\} \quad \text{such that} \quad E[W_{0.01}] \le W^* . \tag{8}$$

2.5 Implementing the Decision Criteria

Incorporating a climate threshold in an economic optimal growth model can result in a non-convex objective function (Keller et al 2004). Gradient based optimization methods, as implemented by popular optimization software programs such as GAMS and MINOS, face the potentially serious problem of converging to a local maxima rather than the global maximum (Moles et al 2004). To address this issue, we use an evolutionary algorithm (Storn and Price 1997). We assess the convergence of the optimal strategy by comparing two independent solutions of the optimization (Keller et al 2007c; McInerney and Keller 2008). Solutions that differ by less than 1% in the optimal abatement strategy are considered practically identical, and are deemed to have converged within the relevant precision.

Solving the optimal control problem for the considered long time horizon of several centuries and two control variables at each time-step (relative investment in capital stock as well as relative abatement of CO_2 emissions) can impose considerable computational challenges. We overcome these challenges by (i) using a sequential approach (Keller et al 2007c), (ii) using stretched discretization of the time variable, and (iii) implementing the model as a FORTRAN program on a high performance parallel computer cluster. The sequential solution method first optimizes the relative investment in capital stock for the business-as-usual (BAU) case (*i.e.*, where no GHG abatement is employed). We then fix the relative investment to these values and optimize the abatement strategy. The variable time discretization uses 60 decision variables for the investment between 2005 and 2595 and 31 decision variables for the optimization of abatement. The first 17 variables for the abatement strategy represent decadal abatement between 2015 and 2165, while the remainder represent abatement between 2185 and 2595 with lower temporal resolution. Linear interpolation is used to determine the decadal abatement levels required for DICE-07. Abatement in 2005 is fixed at 0.02 (2%), while all other abatement variables are constrained between 0 and 1.

3 Results and discussion

We first use this model to analyze the effects of the climate threshold and uncertainty on optimal abatement in the EUM case. We then characterize the effects of choosing different decision criteria on (i) preferred abatement, (ii) the distribution of outcomes, (iii) the tradeoffs between the expected utility and the utility in the lower tail, (iv) the robustness of the strategies with respect to uncertainty about key parameters, and (v) the economic value of information.

3.1 Effects of the climate threshold and uncertainty in the expected utility maximization framework

In the absence of both parametric uncertainty and the climate threshold, optimal abatement rises from 15% in 2015 to around 60% in 2150 (Figure 2, circles). Considering the climate threshold in DICE-07, but still neglecting parametric uncertainty, optimal abatement remains the same within the numerical accuracy of our optimization (Figure 2, crosses). In this case, including the climate threshold is inconsequential since the threshold is not crossed during the considered time horizon. This result differs from previous studies (*e.g.* Keller et al (2004), Figure 5) where the introduction of the identical MOC representation increases optimal abatement in an attempt to avoid or delay the threshold crossing. The key reason for this difference is the reduction in baseline CO_2 concentrations between the DICE-94 and DICE-07 models due to a complex mixture of changes in the utility function, the pure rate of social time preference, and the carbon intensity (Nordhaus 2008).

The effects of including parametric uncertainty and the climate threshold is also very small for the EUM decision criterion (Figure 2, squares). Considering parametric uncertainty and the climate threshold increases optimal abatement by a few percentage points by 2100. From an EU perspective, the considered climate threshold of an MOC shutdown has little effect on the optimal near-term abatement strategy in DICE-07. As we will show below, the considered climate threshold does have considerable effects on optimal near term abatement as well as the economic value of information once decision criteria other than EUM are considered.

3.2 Effects of the confidence in the PDF estimate on preferred abatement and the distribution of utilities

The weight that the decision-maker assigns to the worst case scenarios for the LDC criterion (parameter β in Equation 6) has a considerable effect on the near-term preferred abatement (Figure 3). When the decision-maker has no confidence in the underlying parametric distributions (*i.e.*, $\beta = 0$) and thus maximizes the expected utility of the worst 1% of cases (*i.e.* $E[W_{0.01}]$) the preferred abatement increases considerably compared the EUM case (Figure 3 squares vs. circles). In this “zero confidence” case (squares), abatement increases steadily from more than 20% in 2015 to around 80% in 2100, and reaches 100% around 2135. By comparison, the EU maximizing abatement rises from around 15% in 2015 to 40% in 2100 (circles).

The key difference between the EUM and zero confidence criteria is how much importance they assign to the expected utility of the worst cases, as seen in the PDFs of their utilities in Figure 4 (b) and (d). Since the units for discounted sum of utility are not particularly meaningful, we have linearly scaled these values to the range of BAU utility (Figure 4 a). A rescaled utility of 0 corresponds to the minimum value of BAU utility, while 100 scales to the maximum value of the BAU utility. In each case, the distribution is left-skewed due to the low-probability, high-impact events (panels a and c in Figure 1).

The improvement obtained by maximizing EU instead of following the BAU path is clear. The expected value of the distribution (solid vertical line) improves, and the variance in the distribution is also reduced. Importantly, the worst cases are greatly improved; the expected value of the worst 1% of outcomes ($E[W_{0.01}]$), marked by the dotted vertical line, increases significantly. In the “zero confidence” case, we see that the distribution is reduced even further, with a noticeable improvement in the worst case outcomes and hence the expected value of the worst 1% of outcomes. However, the decision maker must sacrifice some of the expected utility of the entire distribution for this improvement of the worst cases.

3.3 Tradeoff between expected utility and worst cases performance

Implementing the LDC and SF criteria with a range of parameter values allows us to examine the tradeoff between expected utility and utility averaged over the worst case outcomes (Figure 5). For the LDC criterion, we sample values of the confidence parameter β in Equation (6) of $0, 0.1, \ldots, 1$, and additional values 0.95, 0.97 and 0.99 to increase the resolution of the trade-off curve (Figure 5, crosses). As we move from left-to-right in this figure, we move from $\beta = 1$ (EU maximization) to $\beta = 0$ (zero-confidence). Similarly, the tradeoff curve for SF can also be mapped out by considering a range of values for the threshold W^* in Equation (8) based on values from the EU maximization and zero-confidence strategies (Figure 5, circles). Interestingly, the LDC and SF points appear to lie on the same trade-off curve. Abatement strategies for LDC and SF points that lie in close proximity to each other on these tradeoff curves are practically identical. For example, abatement strategies for the right-most SF point (with $W^* = 63$) and the closest LDC point (with $\beta = 0.8$) are displayed in Figure 6 (circles vs crosses). Krokhmal et al (2002)[Theorem 3] implies that if $E[W]$ and $E[W_{0.01}]$ are convex functions, the tradeoff curves for SF and LDC will in fact be equivalent. Despite the potential for non-convexity in $E[W]$ and $E[W_{0.01}]$ due to the climate threshold response, for practical purposes these criteria produce equivalent strategies. This allows the two sets of points to be combined to trace out this curve with higher resolution (solid line in Figure 5).

For the LDC criterion, we may interpret the points on the tradeoff curve as follows. The preferred solution with $\beta = 1$ (*i.e.*, expected utility maximization) corresponds to the point on the curve which has the maximum value in the y-direction, *i.e.* the maximum value in the $[0, 1]$ direction in Cartesian coordinates. For $\beta = 0$ (*i.e.*, zero-confidence) the preferred solution corresponds to the point with maximum value in the x-direction, *i.e.*, in the direction $[1, 0]$. Similarly, values of β between 0 and 1 correspond to the point on the curve with the maximum value in the direction $[1 - \beta, \beta]$. For example, when $\beta = 0.5$, the preferred solution corresponds to the maximum point in the $[0.5, 0.5]$ direction, *i.e.* along the axis that runs at 45° to the positive x-axis. For the SF criterion, the interpretation is more intuitive: the parameter W^* in Equation (8) corresponds to the value of $E[W_{0.01}]$ on the curve.

The abatement strategy that corresponds to the SF point in the middle of this trade-off curve that is found by constraining $W^* = 55$ is displayed in Figure 3. The abatement for this strategy lies between expected utility maximization and zero-confidence, and we will refer to this as our "intermediate-confidence" strategy. This strategy greatly improves the worst case outcomes compared with EUM, but maintains a high expected utility (Figure 4 c).

The tradeoff curve is also useful for identifying "sweet spots" in the tradeoff between expected utility and minimizing the worst cases. At the right-hand end of the curve we see a sudden drop off in expected utility. In this region, we must make a large sacrifice in expected utility to obtain a small improvement in the worst cases. A decision-maker might prefer abatement strategies at the top of this "cliff" compared to stragies that are at the bottom.

3.4 Robustness of the preferred strategies

The LDC and SF decision-making criteria can be used to derive strategies that may prove more robust with respect to deeply uncertain parameters than the strategy that maximizes expected utility (Figure 7). An increase in robustness with respect to a deeply uncertain parameter is here interpreted as a decreased slope of the expected performance of the strategy over the considered parameter range. Figure 7 demonstrates this effect for (a) climate sensitivity, λ^*, and (b) threshold specific damages, θ_3. Each panel shows a sensitivity study where the considered parameter on the x-axis is certain and all other parameters remain uncertain. The two panels are plotted such that an increase of the parameters implies an increase in climate change damages. The fact that the high climate sensitivity values coincide with the lowest values of the utility function for each strategy (*c.f.* Figure 4) shows that the current fat high tail of the climate sensitivity (Knutti and Hegerl 2008; Urban and Keller 2009) is the dominating factor explaining the worst case scenarios in our analysis. Over the considered ranges of the parameter values, the EU maximizing strategy has the steepest overall slope, followed by the intermediate and no-confidence strategies (the same strategies as analyzed in Figures 3 and 4). To some extent, this ranking of slopes is expected given the tradeoff between the expected utility and the expected utility for the worst case scenarios (which are located at the far right of these two panels). What is perhaps surprising is that the curves cross over. This results in a reversal of the strategies ranking with respect to the expected utility for the best case and the worst case scenarios. For example, the expected utility maximizing strategy has the highest utility for the best case scenarios (left hand side of panel a) but has the worst expected utility for the worst case scenarios (right hand side of panel a). This result also shows that the SF criterion can automatically identify strategies in a high dimensional decision-space that may increase robustness relative to EUM.

3.5 Effects of learning

Thus far, we have neglected the possibility of future learning on preferred abatement strategies. Uncertainties in the natural and human systems will often decrease over time as observations are made (*cf.* Keller and McInerney (2008) and Ricciuto et al (2009), but see also Oppenheimer et al (2008) for counter examples). To explore the importance of learning we consider a scenario where we instantaneously learn about the MOC sensitivity in the year 2075. Prior to the time of learning a single abatement path is followed, while after the time of learning one strategy will be implemented if we learn that the MOC is sensitive to climate change, while a different strategy will be used if we find otherwise. Obviously, this representation of learning is somewhat unrealistic; the learning process will likely happen gradually as observations are made (Keller and McInerney 2008; Keller et al 2007a). However, a more realistic representation that would allow for gradual learning is beyond the scope of this paper.

Preferred abatement strategies for the learning scenario outlined above are shown in Figure 8; the three panels correspond to (a) EUM, (b) intermediate-confidence, and (c) zero-confidence. Again, the intermediate confidence strategy is determined using safety first criteria with $E[W_{0.01}] = 55$. Learning about the MOC in 2075 has only a small effect on optimal policy for EUM and zero-confidence (panels a and c). The EU maximizer is not particularly concerned about the MOC threshold (*c.f.* Figure

2), so they will not drastically alter their abatement at the time of learning. On the other hand, the decision-maker with zero-confidence in the underlying parameter distributions is trying to avoid the worst case outcomes at all cost. Even if the MOC is not sensitive to warming, the worst case outcomes will still occur when the climate sensitivity is largest and high abatement levels will help improve these outcomes. So regardless of what information this decision-maker happens to learn, abatement levels will be high.

In contrast, the abatement strategies for the decision-maker with intermediate confidence (Figure 8 b) is noticeably different after the time of learning. If they learn that the MOC is sensitive to climate change, they will increase abatement to avoid or delay an MOC shutdown. If they learn that the MOC is insensitive, abatement levels will decrease shortly after the time of learning since they no longer need to worry about the economic damages of an MOC shutdown.

The tradeoff between $E[W]$ and $E[W_{0.01}]$ for the cases where do not learn about the MOC in the considered time-horizon and when we learn in 2075 is mapped out in Figure 9. The end points of these curves, corresponding to EU maximization (left end) and zero-confidence (right end) are similar, which is expected since the abatement strategies are almost the same. However, for the intermediate-confidence strategies, with $E[W_{0.01}] = 55$, we see a marked improvement in expected utility due to learning. This improvement is even greater for values between 55 and 62.

The difference in expected utility between the two tradeoff curves can be translated into an economic value of information (VOI). We approximate the VOI in the optimal growth model by the difference in 2005 consumption that would equalize the objective function for the case with and without the additional information. The VOI for various values of $E[W_{0.01}]$ is expressed as percentage points of 2005 Gross World Product (GWP) in Figure 10. For reference, 1% of 2005 GWP in DICE-07 is approximately 600 billion US dollars. For the intermediate-confidence strategies, this value is around 20% of GWP, and is larger for values of $E[W_{0.01}]$ between 55 and 62. If learning occurs at an earlier time, this value will likely be greater (Keller et al 2007b).

4 Caveats

We adopt a highly stylized decision-analytical framework. This simplicity improves the transparancy, but neglects several potentially important aspects. Many of the general limitations of Integrated Assessment Models of Climate change are discussed in great detail in the excellent overviews provided, for example, by Ackerman et al (2009), Weyant (2009), or Nordhaus (2008). Given this careful discussion of the general issues, we focus here on key issues that are specific to the considered problem of deeply uncertain climate threshold responses.

First, we analyze only one specific example of potential climate threshold responses. Our conclusion may well change for other climate thresholds such as a disintegration of the Greenland Ice sheet that might be triggered much earlier than the considered example of a shutdown of the North Atlantic Meridional Overturning Circulation (Keller et al 2007a; Schneider et al 2007). Second, we consider only parametric uncertainty about four parameters and only one type of structural uncertainty, a small subset of the full range of such uncertainties that influence any actual climate policy. For instance, the analysis does not account for back-stop technologies and induced technological change (Keller et al 2008). Incorporation of CO_2 sequestration as a backstop technology would increase both $E[W]$ and $E[W_{0.01}]$; however, the cost and global acceptance of this technology is deeply uncertain. Finally, the study considers only a narrow range of policy options – the degree of emissions abatement – that may significantly restrict any ability to identify robust strategies.

5 Conclusion

We modify a relatively parsimonious integrated assessment model of climate change strategies to include simple representations of a potential climate threshold response, structural uncertainty, and learning. We use this model to analyze the effects of different decision-criteria on preferred strategies, the distribution of outcomes, the robustness of the strategies, and the economic value of information. The simplicity of the model enables an arguably transparent analysis, but it also imposes considerable caveats. Subject to these caveats, we draw four main conclusions.

First, choosing decision criteria that put greater weights on the performance under the worst-case scenarios compared to expected utility maximization acts to increase the preferred investments in abating anthropogenic climate forcing. This increase in near term preferred abatement occurs for scenarios with and without learning. Second, increasing the relative importance of the worst-case scenarios is a promising conceptual and computational technique for identifying strategies that may prove more robust, in the sense that they trade off a small decrease in the expected performance for a sizable increase in the performance under the worst-case scenarios. In particular, we have shown that it is possible to identify such strategies in a numerically efficient way for relatively high dimension dynamic and non-convex decision problems such as the ones posed by anthropogenic climate change. Third, increasing the importance of a strategy's performance under worst-case scenarios compared to expected utility maximization can considerably increase the economic value of information. Finally, a key policy relevant conclusion that follows from our analysis is that increasing near term investment in reducing anthropogenic climate forcing may be a promising avenue for increasing the robustness of climate strategies.

Acknowledgements We thank David Budescu, Nathan Urban, Marlos Goes, and Brian Tuttle for invaluable discussions. Funding from the U.S. National Science Foundation (Grant SES-0345925) is gratefully acknowledged.

Appendix: List of symbols

symbol	definition
α	marginal elasticity of utility
BAU	business-as-usual
β	degree of confidence in the underlying parameter distributions
CO_2	carbon dioxide
c	per capita consumption per year
Δt	time increment
ΔT	deviation in global mean temperature form pre-industrial level
EU	expected utility
g_σ	growth of carbon intensity
L	population
LDC	Limited Degree of Confidence
$\lambda^\star$	climate sensitivity
MOC	Meridional overturning circulation
μ	CO_2 abatement relative to the BAU scenario
N	number of time increments
N_{SOW}	number of SOWs
PDF	Probability Density Function
Pr	Probability
p_{MOC}	probability that MOC is sensitive to climate change
q	portion of considered worst-case outcomes
R	discount factor
ρ	pure rate of social time preference
SF	Safety First
SOW	state of the world
σ	carbon intensity
t	time
t_0	initial time
t_n	time at increment n
t_{N-1}	end of considered time horizon
θ_3	threshold specific fractional economic damage
U	flow of utility
W	discounted sum of utility
W_i	for integer i, this is the discounted sum of utility for the i-th SOW
W_q	for probability q, this is distribution of the lowest q portion of the W distribution
W^*	threshold parameter for the SF criterion

References

Aaheim H, Bretteville C (2001) Decision-making frameworks for climate policy under uncertainty. Working paper 2001:02, CICERO, Oslo, Norway

Ackerman F, DeCanio SJ, Howarth RB, Sheeran K (2009) Limitations of integrated assessment models of climate change. Climatic Change 95(3-4):297–315, DOI:10.1007/s10584-009-9570-x

Alexander S, Coleman T, Li Y (2006) Minimizing CVaR and VaR for a portfolio of derivatives. Journal of Banking and Finance 30:583–605, DOI:10.1016/j.jbankfin.2005.04.012

Alley RB, Marotzke J, Nordhaus W, Overpeck J, Pielke R, Pierrehumbert R, Rhines P, Stocker T, Talley L, Wallace JM (2002) Abrupt climate change: Inevitable surprises. National Research Council

Andersson F, Mausser H, Rosen D, Uryasev S (2001) Credit risk optimization with conditional value-at-risk criterion. Mathematical Programming 89:273–291, DOI:10.1007/PL00011399

Andronova NG, Schlesinger ME (2001) Objective estimate of the probability density function for climate sensitivity. Journal of Geophysical Research 106:22,605–22,612, DOI:10.1029/2000JD000259

Borsuk ME, Tomassini L (2005) Uncertainty, imprecision and the precautionary principle in climate change assessment. Water Science and Technology 52(6):213–225

Bretteville Froyn C (2005) Decision criteria, scientific uncertainty, and the global warming controversy. Mitigation and Adaptation Strategies for Global Change 10:183–211, DOI:10.1007/s11027-005-3782-9

Budescu DV, Kuhn KM, Johnson T (2002) Modeling certainty equivalents for imprecise gambles. Organizational Behavior and Human Decision Processes 88:748–768, DOI: 10.1016/S0749-5978(02)00014-6

Du N, Budescu DV (2005) The effects of imprecise probabilities and outcomes in evaluating investment options. Management Science 51:1791–1803, DOI: 10.1287/mnsc.1050.0428

Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. Quarterly Journal of Economics 75(4):643–669

Fortin I, Fuss S, Hlouskova J, Khabarov N, Obersteiner M, Szolgayova J (2007) An integrated CVaR and real options approach to investments in the energy sector. Economic Series 209, Institute for Advanced Studies, Vienna

Hall J, Fu G, Lawry J (2007) Imprecise probabilities of climate change: aggregation of fuzzy scenarios and model uncertainties. Climatic Change 81(3-4):265–281, DOI:10.1007/s10584-006-9175-6

Hurwicz L (1951) Optimality criteria for decision-making under ignorance. Cowles Commission Discussion Paper, Statistics, No. 370

Keith DW (1996) When is it appropriate to combine expert judgments? Climatic Change 33(2):139–143, DOI:10.1007/BF00140244

Keller K, McInerney D (2008) The dynamics of learning about a climate threshold. Climate Dynamics 30(2-3):321–332

Keller K, Tan K, Morel F, Bradford D (2000) Preserving the ocean circulation: Implications for climate. Climatic Change 47:17–43, DOI:10.1023/A:1005624909182

Keller K, Bolker BM, Bradford DF (2004) Uncertain climate thresholds and optimal economic growth. Journal of Environmental Economics and Management 48:723–741, DOI:10.1016/j.jeem.2003.10.003

Keller K, Deutsch C, Hall MG, Bradford DF (2007a) Early detection of changes in the North Atlantic meridional overturning circulation: Implications for the design of ocean observation systems. Journal of Climate 20:145–157, DOI:10.1175/JCLI3993.1

Keller K, Kim SR, Baehr J, Bradford D, Oppenheimer M (2007b) Human-Induced Climate Change, ed. M.E Schlesinger, Cambridge University Press, chap What is the economic value of information about climate thresholds?, pp 343–354

Keller K, Robinson A, Bradford D, Oppenheimer M (2007c) The regrets of procrastination in climate policy. Environmental Research Letters 2(2):024,004 (4pp), DOI:10.1088/1748-9326/2/2/024004

Keller K, McInerney D, Bradford DF (2008) Carbon dioxide sequestration: How much and when? Climatic Change 88:267–291, DOI:10.1007/s10584-008-941

Knutti R, Hegerl G (2008) The equilibrium sensitivity of the earth's temperature to radiation changes. Nature Geosciences 1:735–743, DOI:10.1038/ngeo337

Kriegler E, Held H (2005) Utilizing belief functions for the estimation of future climate change. International Journal of Approximate Reasoning 39(2-3):185–209, DOI:10.1016/j.ijar.2004.10.005

Krokhmal P, Palmquist J, Uryasev S (2002) Portfolio optimization with conditional value-at-risk objective and constraints. Journal of Risk 4(2):11–27

Lange A (2003) Climate change and the irreversibility effect combining expected utility and maximin. Environmental & Resource Economics 25(4):417–434, DOI:10.1023/A:1025054716419

Larsen N, Mausser H, Uryasev S (2002) Algorithms for optimization of value-at-risk. In: Pardalos P, Tsitsiringos V (eds) Financial engineering, E-commerce, and supply chain, Kluwer Academic Publisher, pp 129–157

Latif M, Roeckner E, Mikolajewski U, R V (2000) Tropical stabilization of the thermohaline circulation in a greenhouse warming simulation. Journal of Climate 13:1809–1813, DOI:10.1175/1520-0442(2000)013<1809:L>2.0.CO;2

Lempert R, Sanstad A, Schlesinger M (2006) Multiple equilibria in a stochastic implementation of DICE with abrupt climate change. Energy Economics 28:677–689, DOI:10.1016/j.eneco.2006.05.013

Lempert RJ (2002) A new decision sciences for complex systems. Proceedings of the National Academy of Sciences of the United States of America 99:7309–7313

Lempert RJ, Collins MT (2007) Managing the risk of uncertain threshold responses: Comparison of robust, optimum, and precautionary approaches. Risk Analysis 27(4):1009–1026, DOI:10.1111/j.1539-6924.2007.00940.x

Lempert RJ, Popper SW, Bankes SC (2003) Shaping the next one hundred years: New methods for quantitative, long-term policy analysis. Tech. rep., RAND MR-1626-RPC

Link PM, Tol RSJ (2004) Possible economic impacts of a shutdown of the thermohaline circulation: an application of fund. Portuguese Economic Journal 3:99–114, DOI:10.1007/s10258-004-0033-z

Liu J, Men C, Cabrera V, Uryasev S, Fraisse C (2008) Optimizing crop insurance under climate variability. Journal of Applied Meteorology and Climatology 47:2572–2580, DOI:10.1175/2007JAMC1490.1

McInerney D, Keller K (2008) Economically optimal risk reduction strategies in the face of uncertain climate thresholds. Climatic Change 91:29–41, DOI:10.1007/s10584–006–9137-z

Moles CG, Banga JR, Keller K (2004) Solving nonconvex climate control problems: pitfalls and algorithm performances. Applied Soft Computing 5(1):35–44, DOI:10.1016/j.asoc.2004.03.011

Morgan G, Dowlatabadi H, Henrion M, Keith D, Lempert R, McBrid S, Small M, Wilbanks T (2009) U.S. CCSP Synthesis and Assessment Product 5.2, Best practice approaches for characterizing, communicating, and incorporating scientific uncertainty in decisionmaking. Tech. rep., National Oceanic and Atmospheric Administration

Nordhaus W (1992) An optimal transition path for controlling greenhouse gases. Science 258:1315–1319

Nordhaus W (2007a) The challenge of global warming: Economic models and environmental policy, http://www.econ.yale.edu/~nordhaus/DICEGAMS/DICE2007.htm, accessed May 2, 2007, model version: DICE-2007.delta.v7

Nordhaus W (2007b) A review of the Stern review on the economics of climate change. Journal of Economic Literature 45(3):686–702, DOI:10.1257/jel.45.3.686

Nordhaus WD (1994) Managing the Global Commons. The MIT press, Cambridge, Massachusetts

Nordhaus WD (2008) A Question of Balance: Economic Modeling of Global Warming. Yale Press

Oppenheimer M, O'Neill BC, Webster M (2008) Negative learning. Climatic Change 89(1-2):155–172, DOI:10.1007/s10584-008-9405-1

Peterson GD, Carpenter SR, Brock WA (2003) Uncertainty and the management of multistate ecosystems: An apparently rational route to collapse. Ecology 84(6):1403–1411

Popp D (2004) Entice: Endogenous technological change in the DICE model of global warming. Journal of Environmental Economics and Management 48:742–768, DOI:10.1016/j.jeem.2003.09.002

Quaranta A, Zaffaroni A (2008) Robust optimization of conditional value at risk and portfolio selection. Journal of Banking and Finance 32:2046–2056, DOI:10.1016/j.jbankfin.2007.12.025

Quiggin J (2008) Uncertainty and risk: Multidisciplinary perspectives. In: Uncertainty and Risk: Multidisciplinary perspectives, Earthscan, pp 195–204

Ramsey F (1928) A mathematical theory of saving. The Economic Journal 38(152):543–559

Rawls J (1971) A Theory of Justice. Harward University Press

Ricciuto D, Tonkonojenkov R, Urban NM, Wilkinson R, Matthews D, Davis K, Keller K (2009) Assimilation of oceanic, atmospheric, and ice-core observations into an earth system model of intermediate complexity. Global Biogeochemical Cycles in revision:available at: http://www.geosc.psu.edu/ kkeller/publications.html

Schneider S, Thompson S (1981) Atmospheric CO_2 and climate - importance of the transient-response. Journal of Geophysical Research-Oceans and Atmospheres 86:3135–3147

Schneider SH, Semenov S, Patwardhan A, Burton I, Magadza C, Oppenheimer M, Pittock A, Rahman A, Smith J, Suarez A, Yamin F, Corfee-Morlot J, Finkel A, Fssel HM, Keller K, MacMynowski D, Mastrandrea MD, Todorov A, Sukumar R, Ypersele JPv, Zillman J (2007) Assessing key vulnerabilities and the risk from climate change, Cambridge University Press, Cambridge, UK, pp 779–810

Stocker T, Schmittner A (1997) Influence of CO_2 emission rates on the stability of the thermohaline circulation. Nature 388:862–865

Storn R, Price K (1997) Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces. Journal of Global Optimization 11:341–359, DOI:10.1023/A:1008202821328

Tol R (1994) The damage costs of climate change: a note on tangibles and intangibles, applied to DICE. Energy Policy 22(5):436–438

Tol R (2003) Is the uncertainty about climate change too large for expected cost-benefit analysis? Climatic Change 56(3):265–289, DOI:10.1023/A:1021753906949

Tol RSJ (1998) On the difference in impact of two almost identical climate change scenarios. Energy Policy 26(1):13–20

Urban NM, Keller K (2009) Complementary observational constraints on climate sensitivity. Geophysical Research Letters DOI:10.1029/2008GL036457

Weitzman ML (2009) On modeling and interpreting the economics of catastrophic climate change. The Review of Economics and Statistics 91(1):1–19

Welsch H (1995) Greenhouse-gas abatement under ambiguity. Energy Economics 17(2):91–100, DOI:10.1016/0140-9883(95)00010-R

Weyant JP (2009) A perspective on integrated assessment. Climatic Change 95(3-4):317–323, DOI:10.1007/s10584-009-9612-4

Zickfeld K, Bruckner T (2008) Reducing the risk of atlantic thermohaline circulation collapse: Sensitivity analysis of emissions corridors. Climatic Change 91(3-4):291–315, DOI:10.1007/s10584-008-9467-0

Zickfeld K, Levermann A, Morgan M, Kuhlbrodt T, Rahmstorf S, Keith D (2007) Expert judgements on the response of the Atlantic meridional overturning circulation to climate change. Climatic Change 82:235–265, DOI:10.1007/s10584-007-9246-3

Table 1 Summary of the considered uncertain parameters. MOC refers to Meridional Overturning Circulation.

Parameter	Symbol	Units	Distribution	Mean	Standard deviation	90% confidence interval
Climate sensitivity	λ^*	^{o}C	Empirical	3.4	3.3	[1.0,9.4]
Initial carbon intensity growth	$g_\sigma(2005)$	Per decade	Uniform	-0.073	0.021	[-0.11,-0.04]
MOC specific damages	θ_3	% Economic output	Weibull	1.5	3.7	[0.001,6.9]
MOC sensitivity	p_{MOC}	Dimension-less	Bernoulli	0.5	0.5	[0,1]

Table 2 Summary of the decision criteria considered in this study.

Criteria	Description	Free parameters
Expected Utility Maximization (EUM)	Maximize expected utility contingent on best estimate joint probability distribution	None
Limited Degree of Confidence (LDC)	Maximize weighted average of expected utility and worst-case outcomes, as measured by the Conditional Value at Risk. The CVaR is the expected value of the worst q-th portion of the distribution of utility	1. Weight $(1-\beta)$ on worst-case outcome 2. Percentile q that defines CVaR
Safety First (SF)	Maximize expected utility with the constraint that the CVaR is less than some threshold value.	1. Threshold W^* that constrains CVaR 2. Percentile q that defines CVaR

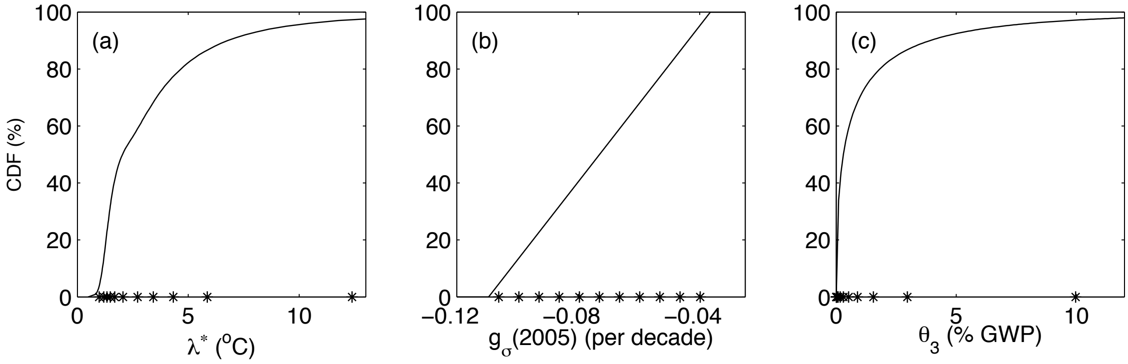

Fig. 1 Cumulative density functions (CDFs) for the considered uncertain parameters. Shown are: (a) the climate sensitivity (λ^*), (b) the initial growth in carbon intensity ($g_\sigma(2005)$), and (c) the damage associated with an MOC shutdown (θ_3). The stars on the horizontal-axis denote the locations of the sampled parameters. See text for details on the sampling procedure.

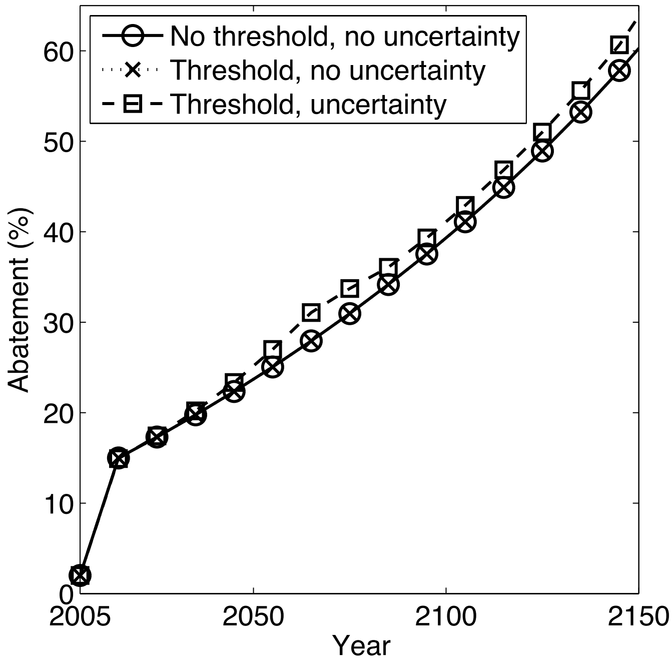

Fig. 2 Preferred abatement strategies without parametric uncertainty and without the climate threshold (circles), with the climate threshold but without parametric uncertainty (crosses), and considering both the climate threshold and parametric uncertainty and using expected utility maximization (squares).

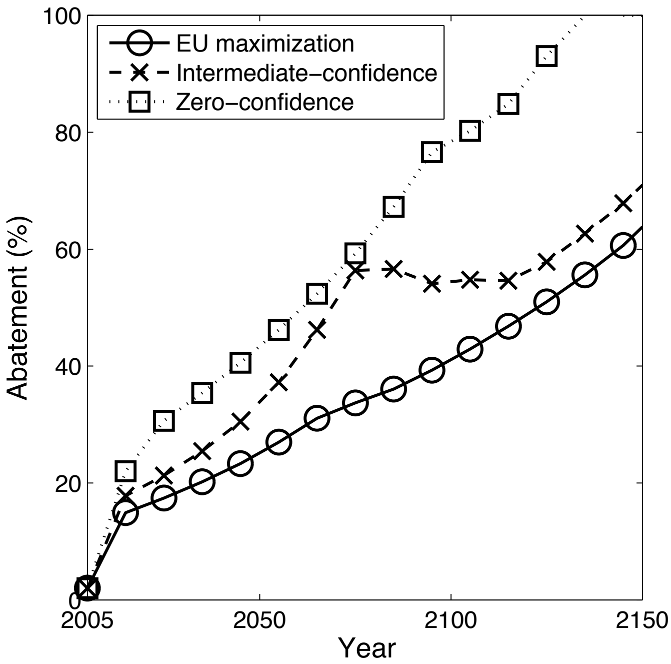

Fig. 3 Preferred abatement levels for expected utility maximization (cirles), intermediate-confidence (crosses) and zero-confidence (squares) strategies. This result considers both the climate threshold and parametric uncertainty, but neglects future learning.

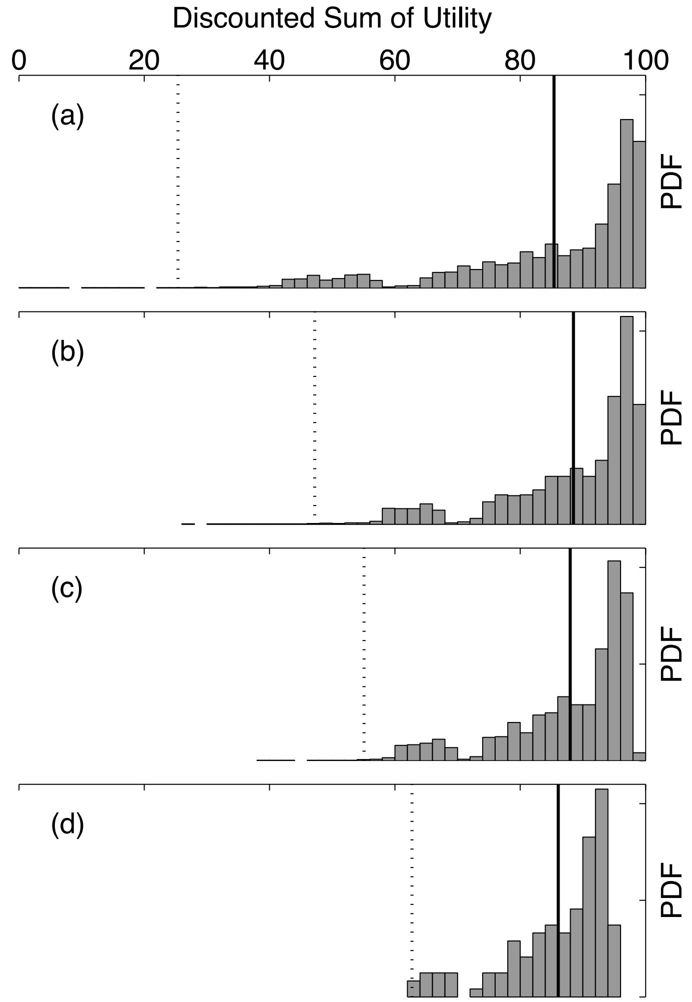

Fig. 4 Probability density functions (PDFs) of the discounted sum of utility for (a) business-as-usual, (b) expected utility maximization, (c) intermediate confidence, and (d) zero-confidence strategies in the absence of future learning. The solid vertical lines indicate the expected value of these distributions ($E[W]$) and dotted lines show the expected value of the worst 1% of outcomes ($E[W_{0.01}]$). Utilities have been rescaled to 100% of the business-as-usual (BAU) range.

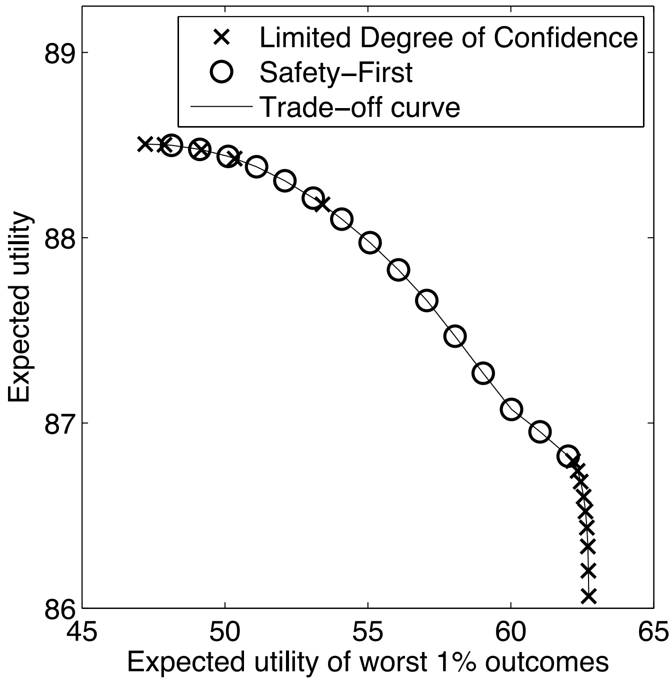

Fig. 5 The tradeoff between the expected utility ($E[W]$) and the expected value of the lowest 1% of utility ($E[W_{0.01}]$). Preferred strategies determined using the limited degree of confidence criterion are marked by the crosses, while those obtained using the safety first criteria are marked by circles. This result considers both the climate threshold and parametric uncertainty, but neglects future learning. Utilities have been rescaled to 100% of the business-as-usual (BAU) range.

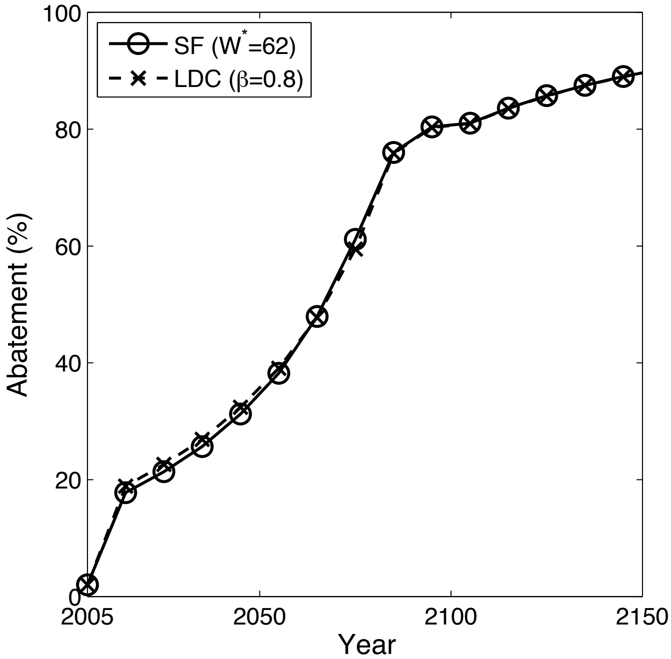

Fig. 6 Preferred abatement levels obtained using the safety first criterion with $W^* = 62$ (circles), and the limited degree of confidence criterion with $\beta = 0.8$. See the main text and Equations (8) and (6) for a description of these variables. This result considers both the climate threshold and parametric uncertainty, but neglects future learning.

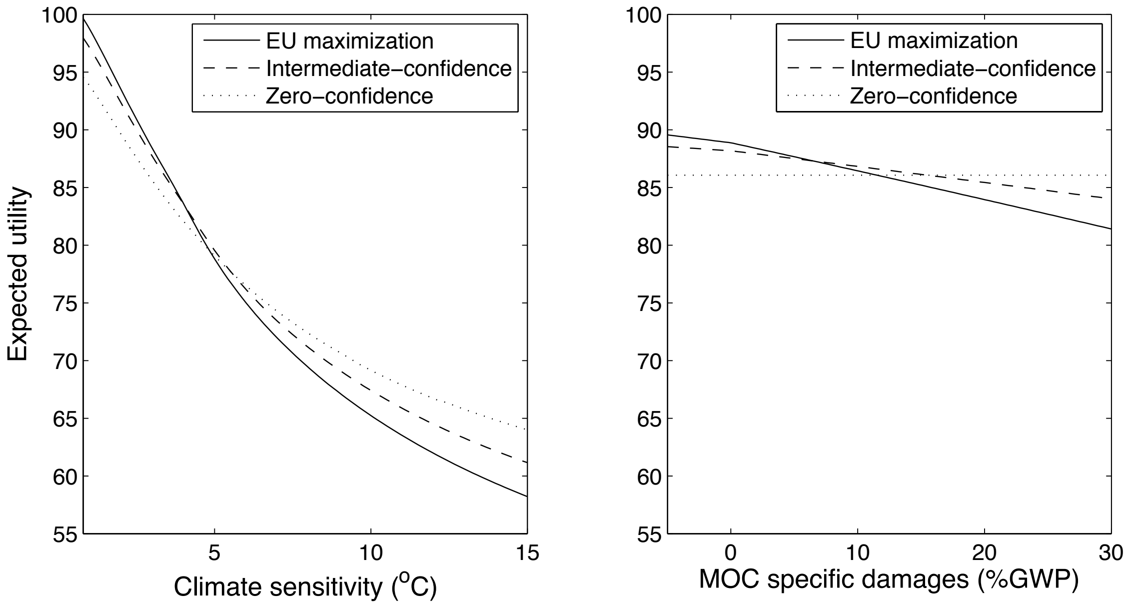

Fig. 7 Expected utility for the three considered strategies when (a) climate sensitivity (λ^*), and (b) MOC specific damages (θ_3), are known precisely. The other uncertain parameters are still sampled over their range (see Table 1 and Figure 1).

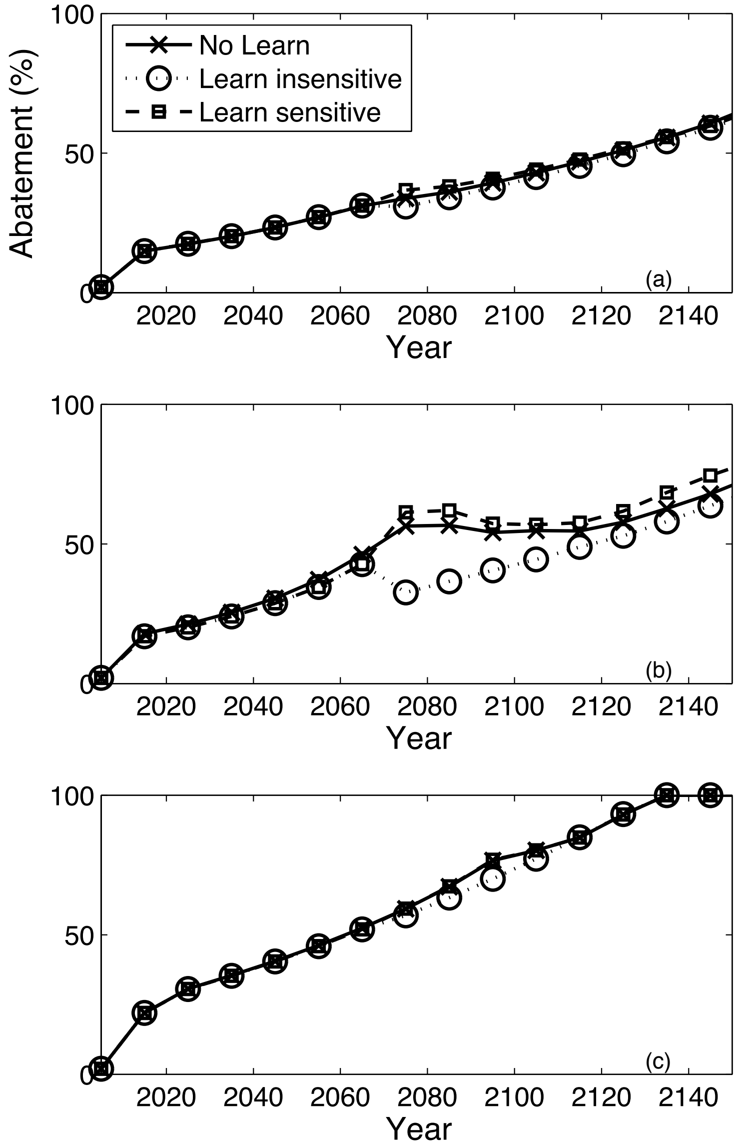

Fig. 8 Preferred abatement levels under uncertainty with learning about the MOC sensitivity in 2075. Decisions are based on (a) expected utility maximization, (b) intermediate-confidence, and (c) zero-confidence criteria.

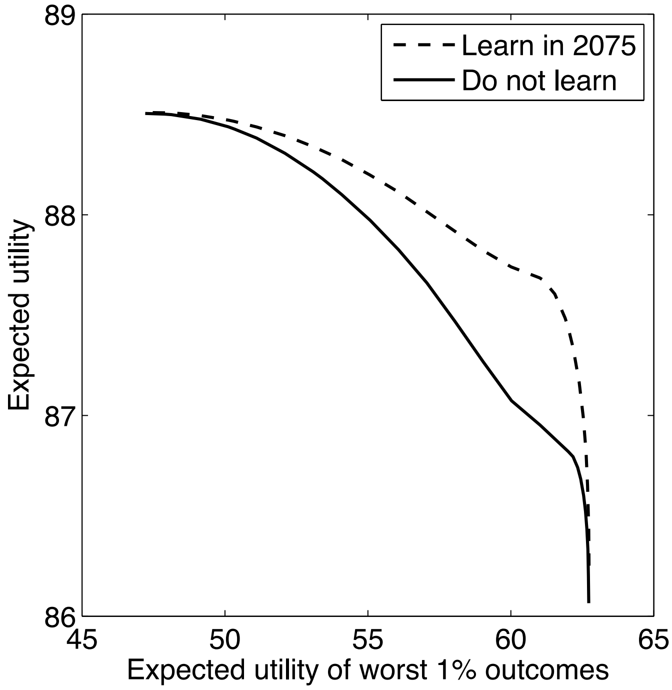

Fig. 9 The tradeoff between expected utility ($E[W]$) and expected value of the lowest 1% of utility ($E[W_{0.01}]$) for the cases where we do not learn during the considered time-horizon whether the MOC is sensitive to warming (solid), and when we learn about this sensitivity in the year 2075 (dashed). Preferred strategies for each case are determined using both the limited degree of confidence and safety first criteria with results for each case collated onto single curves. Both the climate threshold and parametric uncertainty are considered, and utilities have been rescaled to 100% of the business-as-usual (BAU) range.

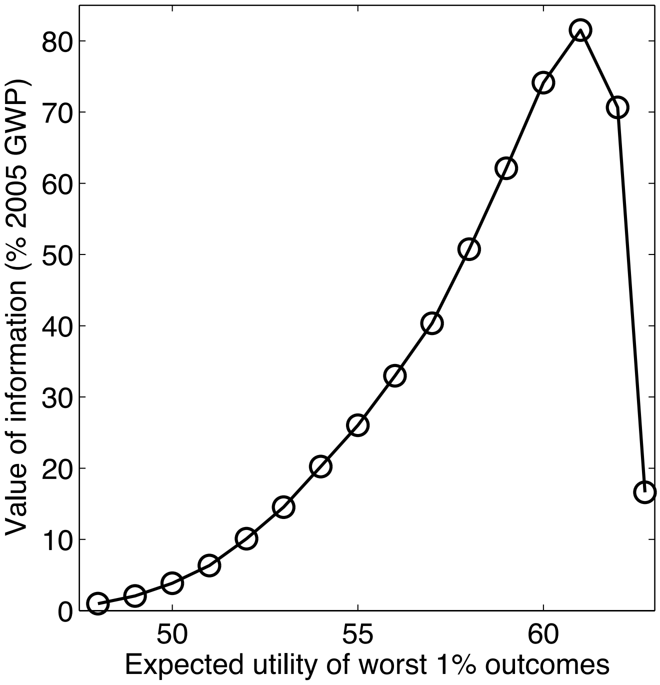

Fig. 10 The economic value of learning about the Meridional Overturning Circulation (MOC) sensitivity in 2075. We consider the effect of constraining the expected value of the worst 1% of utilities ($E[W_{0.01}]$) to a range of values using the safety first criteria.